```
@inproceedings{scotzniovsky2026large,
  author = {Luca Scotzniovsky and Nicholas C. Orndorff and Andrew H. Fletcher and Jason
    Kao and John T. Hwang},
  title = {Large-Scale Multidisciplinary Design Optimization of a Blended Wing Body
    Aircraft Using Mid-Fidelity, Multi-Point Analysis},
  booktitle = {AIAA AVIATION 2026 Forum},
  year = {2026},
  doi = {10.2514/6.2026-4247},
  url = {https://doi.org/10.2514/6.2026-4247}
}
```

Large-Scale Multidisciplinary Design Optimization of a Blended Wing Body Aircraft using Mid-Fidelity, Multi-Point Analysis

Luca Scotzniovsky*, Nicholas Orndorff†, and Andrew Fletcher‡
University of California San Diego, La Jolla, California, 92023

Jason Y. Kao§
Air Force Research Laboratory, Wright-Patterson AFB, Ohio, 45433

John T. Hwang¶
University of California San Diego, La Jolla, California, 92023

The blended-wing-body (BWB) aircraft configuration, like many other unconventional configurations, offers the potential for significant improvements in efficiency and performance. However, the lack of historical data and empirical sizing models presents a major challenge in conceptual design, especially in the early stages. Physics-based predictive models can help address this gap by enabling precise analysis across the range of operating conditions that the aircraft is expected to encounter. This paper demonstrates large-scale, multi-point multidisciplinary design optimization (MDO) of a generic transport BWB concept. We use a 3D panel code to capture aerodynamic performance across nine points in the range-payload space, a Reissner–Mindlin shell model to capture structural response, and empirical models to account for the remaining physics. The MDO framework is implemented using a graph-based modeling approach, which automates the construction of adjoint sensitivity analysis models for efficient gradient-based MDO. With this framework, we solve a series of six optimization problems, with up to 44 shape, structural, and control variables, enforcing different trim and static margin constraints. This study demonstrates the effectiveness of a comprehensive large-scale MDO approach with mid-fidelity, multi-point analysis to capture the critical aspects of the conceptual design problem for unconventional configurations such as the BWB.

I. Introduction

Interest in novel and unconventional aircraft concepts and technologies has surfaced in recent decades to address various initiatives in the aviation industry. These include configurations with improved aerodynamic efficiency [1–4] and novel propulsion technologies [4, 5] for green aviation, quiet supersonic aircraft for commercial operations [6], and electric vertical take-off and landing (eVTOL) aircraft to support the emerging urban air mobility (UAM) market [7]. The conceptual design stage of these aircraft configurations often requires significant fundamental ground-level work due to the lack of historical data or models. This naturally leads to increased time and cost to develop, test, and certify these aircraft, as design trade studies cannot be conducted with existing data. Improved confidence in early-stage designs can avoid costly late-stage redesigns and reduce the time to first flight. This can be achieved by introducing more realism into the design process through higher-fidelity, physics-based, uncertainty-quantified analysis [8]. With modern-day technology, this vision is more realistic than before, and even a reality for some. For example, the recently demonstrated YQF-44A unmanned aerial vehicle, developed by Anduril, went from clean-sheet design to first flight in only 556 days *.

Large-scale multidisciplinary design optimization (MDO), the application of numerical optimization to high-dimensional design problems containing many disciplines, lends itself to aircraft design due to the complex, coupled nature of aircraft disciplines and systems. Adapted from Venkataraman et al., the complexity of aircraft design

*PhD Candidate, Department of Mechanical and Aerospace Engineering, lscotzni@ucsd.edu, AIAA Student Member
†PhD Candidate, Department of Mechanical and Aerospace Engineering, norndorff@ucsd.edu, AIAA Student Member
‡PhD Candidate, Department of Mechanical and Aerospace Engineering, afletche@ucsd.edu, AIAA Student Member
§Research Aerospace Engineer, Aerospace Systems Directorate, jason.kao.1@us.af.mil, AIAA Member
¶Associate Professor, Department of Mechanical and Aerospace Engineering, jhwang@ucsd.edu, AIAA Member
*The press release for this development can be found here

optimization problems can be divided into three main categories: simulation, problem, and optimization [9]. Simulation complexity refers to solver accuracy: fidelity and number of states. Problem complexity refers to the depth and comprehensiveness of the design problem to simulate practical aircraft operations: disciplines, subsystems, and design conditions. Optimization complexity refers to the optimization problem formulation, specifically the set of design variables, parameters, and constraints. The challenge in aircraft design is solving problems whilst simultaneously increasing the complexity along each axis.

Much of the existing work in aircraft MDO focuses on the simulation and optimization complexity. The state-of-the-art aircraft MDO demonstrations consider coupled high-fidelity physics, such as aeroelasticity [10, 11], aeroacoustics [12, 13], and aeropropulsive [14], using adjoint-based optimization and considering 10s to 100s of design variables. However, many of these studies use a single-point approach, as the high computational cost prohibits expanding the problem scope to include more design conditions or disciplines. There has been work with parallel high-fidelity multipoint optimization in the literature, such as the study conducted by Kenway et al., which explores multipoint aerostructural optimization of a transport aircraft across five cruise and two structural sizing conditions using Euler computational fluid dynamics (CFD) and a structural finite-element model [15]. Euler CFD with empirical viscous corrections was selected over RANS CFD due to the intractable computational cost. This deliberate choice highlights the tradeoff between simulation complexity and design problem comprehensiveness.

On the other hand, other recent aircraft design optimization studies focus on detailed mission analysis [16, 17] and trajectory optimization [18, 19] over high-fidelity coupled physics. These demonstrations go beyond a single-point approach by including detailed analysis and modeling performance across a design condition or an entire mission profile, but typically utilize low-fidelity empirical or physics-based models. This tradeoff between low-fidelity and higher-fidelity physics-based solvers, which can incur prohibitively high computational costs, is made to allow for significant complexity in the design problem.

There has been recent progress towards comprehensive large-scale MDO, which aims to model the operation of the entire aircraft across many on- and off-design conditions. This was first demonstrated by Sarojini et al. on the NASA lift-plus-cruise configuration, considering over 100 design variables [20]. Key disciplines, like aerodynamics, structures, acoustics, and propulsion, were considered using low-fidelity physics-based models. This work was later expanded on by Ruh et al. by modeling a full mission profile across 20 design conditions with close to 400 design variables [21]. This work also introduced stability analysis and higher-fidelity acoustics models. CSDL [22], a recently developed graph-based modeling paradigm that automates adjoint-based derivative computation, significantly enabled these groundbreaking demonstrations. A major limitation of existing comprehensive large-scale MDO demonstrations is that key disciplines are modeled with low-fidelity methods. Despite that, existing work suggests that comprehensive large-scale MDO for aircraft design can facilitate the exploration of a large design space to capture important high-level trends, especially for novel and unconventional aircraft configurations.

The blended wing body (BWB) configuration has gained significant interest in recent decades as a novel aircraft configuration to address aviation emissions problems [1]. Compared to standard tube-and-wing configurations, BWBs have the potential to significantly improve fuel burn and aerodynamic efficiency. These benefits stem from the centerbody lift generation and reduced skin friction and interference drag from the wetted-area-to-volume ratio and smooth geometric span transition, respectively [23]. However, due to the absence of a tail, the mechanisms for achieving trim and stability are more tightly coupled to the aerodynamic performance than a traditional tube-and-wing configuration. Control surface deflection can potentially negatively impact spanwise load distributions and drag [24]. These challenges make it critical to introduce trim and stability considerations early on in the design process [23, 25]. Other challenges include integration of the aft-mounted propulsion system [24]. Early studies by Liebeck and Wakayama explored BWB optimization using WingMOD, an aircraft sizing and operation code that uses vortex lattice and monocoque beam analysis for coupled aeroelastic analysis [1, 26, 27]. These studies adopted a comprehensive design approach across many disciplines and flight conditions, opting for lower-fidelity methods [27]. In turn, due to the rapid turnaround time, many important design trends between aerodynamics, structures, and stability were captured. More recent BWB design studies focus heavily on the aerodynamic performance, showcasing adjoint-based aerodynamic shape optimization using Euler CFD [23, 28] and RANS CFD [29–31]. The high-fidelity flow physics is necessary to capture critical physics, such as compressibility and viscous effects, during nominal cruise operation. However, many of these studies incur prohibitive computational cost, preventing multidisciplinary considerations or a comprehensive multipoint approach consisting of many operating conditions. Other studies in the literature assess more of the problem and design complexity by conducting conceptual design studies of overall performance, addressing areas such as airframe integration and geometric layout. Brown et al. utilized zeroth-order models for mass, drag, and stability with detailed fuselage planform shaping and cabin layout models to assess the impact on total mass, aerodynamic efficiency, and fuel burn [2]. Ahuja et al. explored the impact of airframe configurations using FLOPS and CFD-based aerodynamic polars to assess key outputs such as fuel burn, ramp weight, and lift-to-drag ratio [32]. However, the existing BWB studies do not simultaneously model a comprehensive set of operating conditions with enough fidelity to capture the relevant physics and consider a large set of design variables and constraints.

In this work, we address this gap by demonstrating large-scale MDO on a generic transport blended wing body (BWB) concept. We model multiple design conditions, including nine points in the range-payload space in cruise. We employ a 3D panel code to capture aerodynamic performance, a Reissner–Mindlin shell model for structural analysis, and a series of empirical methods and corrections to account for missing physics and additional design conditions. The MDO framework is developed with the graph-based modeling framework, CSDL, to automate adjoint-based sensitivity analysis.

The remainder of this paper proceeds as follows. Section II provides basic information about the blended wing body configuration used in this work. Section III outlines our proposed computational framework, summarizing the details of our tools, methods, and implementation. Section IV presents a detailed description of our optimization problem formulation, including the design variables and constraints. Section V presents results for numerical experiments conducted before optimization to assess and improve the overall computational cost of the MDO demonstration. Section VI discusses the results of our comprehensive high-dimensional MDO demonstration of the BWB configuration. Finally, section VII summarizes the findings, takeaways, and future research directions.

II. Blended Wing Body Configuration

In this section, we describe the basic info about our blended wing body (BWB) configuration. The BWB aircraft considered in this paper is sized similarly to the smaller 747 variants, with a payload capacity of approximately 160,000 lb and a maximum takeoff weight of 350,000 lb. The aircraft is powered by two GE CF6-80A turbofans mounted on the aft centerbody, each rated at 48,000 lb of maximum sea-level thrust. To improve longitudinal stability, reflex is applied to the centerbody section to trim the large pitching moment without requiring significant washout in the outer wing section. The centerbody uses a symmetric NACA 0015 airfoil, which transitions to a transonic BACXXX airfoil for the outer wing section. The outer mold line (OML) geometry is designed in NASA's OpenVSP software [33], and the internal structure (ribs and spars) is added to the OML using Engineering SketchPad [34]. The baseline aircraft is sized for a nominal 2,500 nm mission using the semi-empirical methods of Raymer [35], with correction factors applied to better capture the aerodynamic and structural characteristics unique to the BWB configuration. The baseline aircraft weights are compared to a Boeing 737 and Boeing 747 in Table 1. A three-view diagram with dimensions is shown in Figure 1a. The BWB planform configuration is shown against the Boeing 737 and Boeing 747 configurations in Figure 1b.

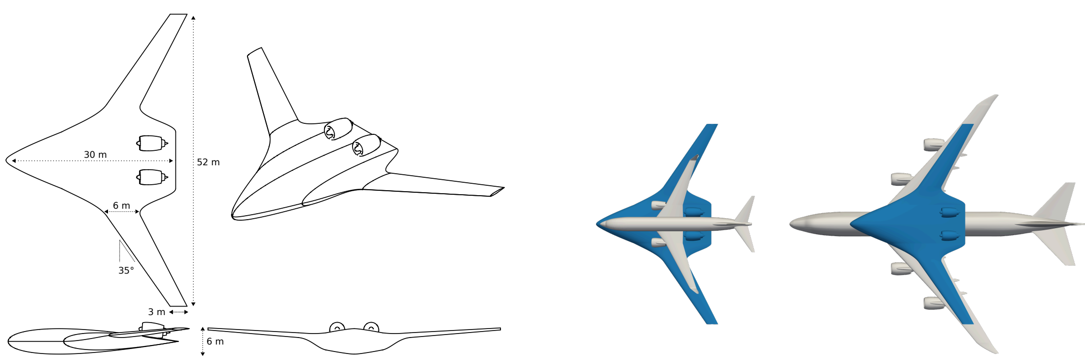

(a) Three-view diagram of the baseline BWB configuration with dimensions.

(b) Planform comparisons of the blended wing body configuration vs. the Boeing 737 and 747.

Fig. 1 Our baseline BWB configuration for our MDO demonstration is shown here, with original dimensions and the size relative to existing aircraft.

III. Computational Framework

In this section, we outline the details of the computational framework developed to demonstrate MDO of the blended wing body (BWB) configuration. We summarize our physics-based solvers, empirical methods, and computational tools.

Table 1 Comparison of BWB sizing and weight parameters against the Boeing 737 and 747 configurations.

Aircraft	Wingspan (ft)	Wing area (ft^2)	TOGW (lbs)	Empty weight (lbs)	Payload weight (lbs)	Fuel weight (lbs)
747	211	5500	875000	403500	171500	300000
737 (M7)	117.83	1344	177000	91300	39800	45900
BWB	142.78	3234	331522	126636	160000	44885.25

A. Geometry Parameterization

We represent the BWB geometry as a set of B-spline surfaces. To propagate physical changes to the BWB geometry, we use free-form deformation (FFD), a technique that parametrizes geometry changes rather than the geometry itself [36, 37]. FFD works by embedding a geometry into a bounding volume to obtain its parametric representation within the volume. The geometry adheres to the parametric coordinates within the bounding volume and changes accordingly when the bounding volume deforms. The FFD volume is represented as a series of B-splines; this work incorporates quadratic B-splines to propagate nonlinear variations to the geometry. This geometry parameterization approach also lends itself to updating surface meshes. Similar to FFD, any surface mesh or point cloud can be projected onto the geometry to obtain a parametric representation of the grid. Examples of the FFD bounding volume and mesh projections onto the B-spline surface representation of the geometry are shown in Figures 2a and 2b, respectively. Any geometric changes applied to the FFD volume get propagated to the geometry and thus to the projected mesh. The geometry parameterization approach here is implemented within an in-house software called lsdo_geo [38] †.

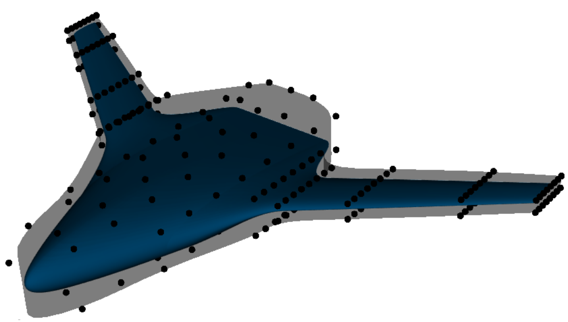

(a) The BWB geometry embedded within an FFD volume.

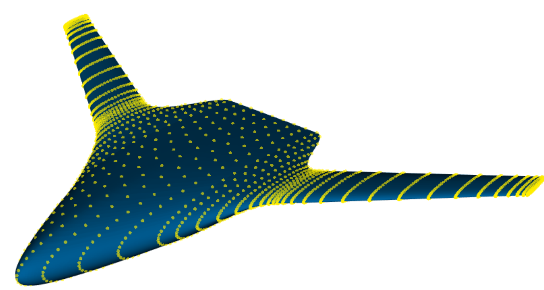

(b) A mesh projected onto the geometry.

Fig. 2 Examples of the geometry b-spline representation with FFD and point projections. FFD enables geometry parameterization in a smooth and continuous fashion, and the point projections enable a simple way to propagate geometry changes to meshes.

B. Aerodynamics

We utilize a 3D low-order panel method to analyze the flow field around the aircraft; we use VortexAD, a fast and memory-efficient panel code with automated adjoint-based sensitivities [39]. Panel methods are a class of numerical methods that compute flow properties around arbitrary three-dimensional bodies. Based on the assumptions of linearized potential flow theory, the flow field is assumed to be incompressible, inviscid, and irrotational. Applying these assumptions to continuity, our governing equation in linearized potential flow becomes

$$\nabla^2\Phi = 0, \tag{1}$$

where Φ is the velocity potential. The irrotational condition relates the velocity potential to the velocity as

$$\vec{U} = \nabla\Phi = \nabla\Phi^* + \nabla\Phi_\infty, \tag{2}$$

† https://github.com/LSDOlab/lsdo_geo

where $\vec{U}$ represents the total velocity vector. Discretizing our surface using a series of doublets and sources, and applying a zero perturbation potential boundary condition, our governing equation reduces to

$$A\vec{\mu} + B\vec{\sigma} = 0, \tag{3}$$

where the vectors $\vec{\mu}, \vec{\sigma} \in \mathbb{R}^n$ represent the vectorized doublet and source strengths across the mesh panels. The matrices $A, B \in \mathbb{R}^{(n,n)}$ represent the induced potential interactions between the mesh panels. With a low-order formulation, we treat the singularities as constant-strength on a panel. The source strengths are predefined as

$$\sigma = -\vec{U}_\infty \cdot \vec{n}, \tag{4}$$

where $\vec{U}_\infty$ is the free-stream velocity vector. After solving the linear system in Equation 3, the total velocity at the collocation points as

$$\vec{Q} = \vec{U}_\infty + \nabla\mu + \sigma\mathbf{n}, \tag{5}$$

where **n** represents the outward unit normal vector at the collocation point. The second and third terms correspond to the local change in velocity induced by the lifting body; the source term acts normal to the surface, and the doublet gradient acts tangent to the surface. The pressure coefficient at each panel, derived from the steady, linearized potential flow form of the Navier-Stokes momentum equation, is computed using

$$C_{p_i} = 1 - \frac{||Q||^2}{||U_\infty||^2}, \tag{6}$$

where the subscript i labels this as the incompressible pressure coefficient. The pressure coefficient can then be integrated to compute aerodynamic forces and moments. To capture compressibility, we apply the Prandtl-Glauert correction to the incompressible pressure coefficient:

$$C_{p_c} = \frac{C_{p_i}}{\sqrt{1 - M_\infty^2}}, \tag{7}$$

where M_∞ represents the freestream Mach number, and subscripts c represent the compressible pressure coefficient.

C. Structures

For structural analysis, we use a Reissner-Mindlin (RM) shell model [40]. The RM shell model is a fully nonlinear six-parameter (three displacements, three rotations) shell model based on first-order shear deformation theory. The shell solver is written using FEniCSx, an open-source PDE solver library that provides adjoint-based sensitivities. We interface to FEniCSx through FEMO, an open-source framework that facilitates the usage and implementation of FEniCSx finite-element models [41]. Written using CSDL, FEMO automates the data transfer for primal and adjoint model evaluations to the central optimization framework. A sample evaluation of the thickness distribution and stress field with the computational mesh is shown in Figure 3.

D. Load Transfer

To map aerodynamic loads to the structural mesh, we use Solver-Independent Field Representation (SIFR) method [42]. SIFR acts as a modular interface between solvers to map solution fields across disciplines. SIFR represents the geometry as a set of B-splines on which fields can be defined. This reduces the effort to develop custom load transfer methods across discipline solvers because the fields are mapped to some continuous representation of the geometry, typically the outer mold line. As a result, the solvers communicate only with SIFR. For the feedforward aerostructural model, the aerodynamic loads from the panel method are mapped to the geometry and then mapped onto the structural mesh to compute displacements and stresses. The reader is encouraged to explore the original paper by Warner et al. [42] for further details on the SIFR methodology.

E. Constraint Aggregation

To design a feasible structure, constraints on the stress field must be introduced. However, it is infeasible to constrain each individual value in the stress field. We use a filtered constraint aggregation approach using a smooth rectifier to ensure that the stress field does not exceed the yield stress of our material properties. Given a material yield stress σ_y, we define our stress violation as

$$SV = \frac{\sigma}{\sigma_y} - 1, \tag{8}$$

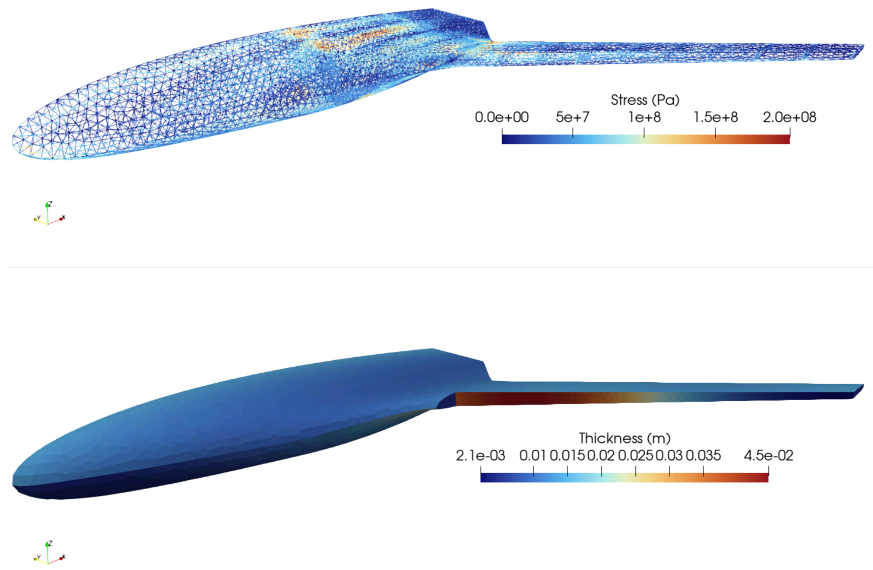

Fig. 3 Example evaluation of the RM shell model on the BWB configuration. The upper solution shows the stress field over the computational mesh, while the lower graphic shows the thickness distribution.

where values below zero indicate that the stress field is below the yield stress. Rather than constraining direct stress values (which are $O(10^6) - O(10^8)$), we normalize by the yield stress to work with numbers closer to $O(1)$. Yield stress violations will only occur for SV values above 0. As a result, we are not interested in values of SV below zero, as they no longer violate the maximum material stress. We evaluate the maximum $\max(x, 0)$ in a smooth and differentiable manner to effectively filter out the non-violating stress terms. We use a monotonic piecewise C2 function from Hollis et al. [43] to smooth the gradients of the maximum function around 0. The smooth maximum function becomes:

$$\max_s(0, x) = \begin{cases} 0 & x \leq 0 \\ 3x^5 - 8x^4 + 6x^3 & 0 < x < 1 \\ 1 & x \geq 1 \end{cases}, \tag{9}$$

where the region $0 < x < 1$ is smooth at the boundaries up to the second derivatives. Once the stress violation field has been filtered, we directly integrate over our structural domain:

$$SVC(SV) = \int_\Omega \max_s(0, SV) d\Omega, \tag{10}$$

where $SVC > 0$ indicates that there is at least one element in the domain that is above the yield stress. A direct sum (or integration) is used here to avoid the poor scaling of derivatives coming from p-norms and log-sum-exponential functions. Although the maximum stress is never explicitly computed, we are interested in constraining the entire stress field to be below the yield stress. Upon conducting numerical experiments, we found that this method slightly underpredicts the stress constraint violation. We introduce a 1.2 scaling factor to reach a conservative estimate of the stress constraint violation. The max stress constraint in our optimization problem is formulated such that $SVC \leq 0$.

F. Longitudinal Static Stability

Maintaining stability in-flight is a key aspect of aircraft operation. Unlike traditional tube-and-wing aircraft, tail-less aircraft, like BWB or flying wings, cannot leverage the benefits of an extra lifting surface with a large moment arm to maintain longitudinal stability or meet moment trim. Rather, these aircraft types rely fully on planform shape, airfoil reflex, and on-body trim surfaces. In this work, we assess longitudinal static stability behavior using static margin. The static margin is defined as

$$K_n = -\frac{C_{M\,\alpha}}{C_{L\,\alpha}}, \tag{11}$$

where $C_{M\,\alpha} = \frac{\partial C_M}{\partial \alpha}$ and $C_{L\,\alpha} = \frac{\partial C_L}{\partial \alpha}$ are the partial derivatives of the moment coefficient evaluated at the center of gravity and the lift coefficient with respect to the pitch angle. We use a finite difference approach to estimate these derivatives by perturbing the pitch angle. We use a step size of 0.1°. We expect to operate in the linear region of the lift and moment coefficients, so we can use a relatively large step size to estimate sensitivities to pitch angle.

G. Empirical Models and Corrections

To model high-level operation and capture design trends of the BWB influenced by various key disciplines, we use the Breguet range equation to capture tradeoffs in steady, level flight during cruise. The Breguet range equation can be written as

$$R = \frac{1}{TSFC}\frac{V}{g}\frac{L}{D}\ln\frac{W_1}{W_2}, \tag{12}$$

where W_1 and W_2 represent the initial and final cruise aircraft weights respectively, $TSFC$ is the thrust-specific fuel consumption, V is the cruise velocity, and $\frac{L}{D}$ is the lift-to-drag ratio. In MDO, this is commonly used due to the interactions between aerodynamics ($\frac{L}{D}$), structures (W), and propulsion ($TSFC$). Note that the fuel burn in cruise, denoted W_f, is directly related: $W_2 = W_1 - W_f$.

The lift-to-drag ratio $\frac{L}{D}$ plays a significant role in the Breguet range equation because it represents the aerodynamic efficiency of the aircraft, which heavily dictates the aircraft's fuel efficiency. Although our panel method will capture lift reasonably well, it only captures lift-induced drag. As a result, we will grossly overestimate the aerodynamic efficiency without empirical corrections to capture other forms of drag. Most notably, we require models for parasite drag and wave drag. To estimate parasitic drag, we use a fully-turbulent flat-plate skin friction coefficient combined with a form factor from Raymer [35]. The skin friction coefficient models viscous drag, while the form factor captures the effects of pressure drag. The skin friction coefficient is defined as

$$C_f = \frac{0.455}{(\log_{10} Re)^{2.58}(1 + 0.144M^2)^{0.65}}, \tag{13}$$

where Re is the Reynolds number and M is the mach number. The form factor is represented as

$$FF = \left[1 + \frac{0.6}{(x/c)_m}\frac{t}{c} + 100\left(\frac{t}{c}\right)^4\right]\left[1.34M^{0.18}(\cos\Lambda_m)^{0.28}\right] \tag{14}$$

where $(x/c)_m$ is the chordwise location of airfoil maximum thickness, and Λ_m is the sweep at the max-thickness line. Unlike a traditional tube-and-wing transport aircraft, the BWB is shaped like a flying wing; as a result, we can neglect typical interference factors between lifting and non-lifting bodies, like fuselages. Additionally, the form factor in Equation 14 is valid across the entire BWB wingspan due to its shape. The skin friction coefficient and form factor are used to compute the parasitic drag coefficient:

$$C_{D0} = C_f \frac{S_{wet}}{S_{ref}} FF, \tag{15}$$

where S_{wet} is the aircraft wetted area and S_{ref} is the reference area, taken to be the planform. These two area terms are computed directly using the geometry. We use a strip-theory approach to compute total parasitic drag by estimating C_{D0} along spanwise sections to capture changes in the local Reynolds number due to variations in chord length. The panel method also does not capture drag due to shock wave formation, known as wave drag. We estimate the wave drag coefficient using an empirical formulation from Raymer [35]:

$$C_{Dw} = 20(M - M_{cr})^4, \tag{16}$$

where the critical mach number M_{cr} is defined as

$$M_{cr} = M_{DD} - \left(\frac{0.1}{80}\right)^{1/3}, \tag{17}$$

where M_{DD} is the drag divergence number. We compute M_{DD} using the sweep-corrected Korn equation as described by Malone et al. [44]:

$$M_{DD} = \frac{\kappa_A}{\cos\Lambda} - \frac{t/c}{\cos^2\Lambda} - \frac{C_L}{10\cos^3\Lambda}, \tag{18}$$

where $\kappa_A s$ is the airfoil technology factor and Λ is the wing sweep angle. C_L is computed from the aerodynamic analysis of each mission, and t/c is updated accordingly from the geometry.

In addition, we model takeoff and landing missions within the optimization problem. Rather than simulating these segments directly, we utilize simplified formulations for landing and takeoff lengths from Anderson [45], as shown below:

$$\begin{aligned} S_{TO} &= \frac{1.44W^2}{g\rho S C_{Lmax} T}, \\ S_{LA} &= \frac{1.69W^2}{g\rho S C_{Lmax}(D + \mu_r(W - L))}, \end{aligned} \tag{19}$$

where weight, area, lift, and drag are computed upstream from other disciplines.

H. Graph-Based Modeling and Optimization Framework

We develop our BWB MDO demonstration framework using CSDL, a graph-based modeling paradigm that automates adjoint-based sensitivity analysis [22]. Our physics-based solvers, unless specified, are also developed using CSDL. Graph-based modeling offers vectorization to improve the forward evaluation run time. We utilize CSDL's vectorization capabilities to analyze mission segments simultaneously in parallel. We interface with the optimizer using modOpt, an open-source optimizer development library that offers various performant and educational optimizers [46]. For our optimization demonstration, we use PySLSQP, a wrapper for the SLSQP optimizer from SciPy [47]. PySLSQP offers a more transparent version of SLSQP with additional functionalities and post-processing analysis tools.

IV. Problem Description

In this section, we outline the problem formulation for our MDO demonstration of the blended wing body (BWB) configuration. Our objective is to optimize a generic transport BWB configuration across many cruise conditions. Each cruise condition is designed to represent a different operating point in the flight envelope; to model this, we populate the range and payload space. Our objective function is the weighted fuel burn across all missions:

$$W_f = \Sigma_i^n w_i W_{f\,i}, \tag{20}$$

where n is the number of cruise conditions, $W_{f\,i}$ represents the per-mission fuel burn. We treat condition 5 as our nominal condition, as this point represents the median payload and range. The mission weights w_i are determined based on the expected frequency of a given mission, where the largest weight is applied to the nominal condition. The full list of cruise design conditions, varying points in the range-payload space, is shown in Table 2. Next, we include two maneuver conditions to size the aircraft structure. We consider a 2.5g pull-up and -1g pull-down maneuver at sea-level. In addition, we include a stability condition by perturbing the pitch angle at the nominal (median) cruise condition. We compute static margin according to equation 11. We use the finite-difference method to estimate the moment and lift coefficient derivatives, perturbing the pitch angle by 0.1 degrees. Finally, we include simple takeoff and landing models.

A. Design Variables

We model a large set of design variables related to the operation and overall shape of the aircraft. We utilize pitch angle and sectional twist as our inflow design variables We assign a pitch angle to each cruise condition and maneuver condition, as it is the primary driver behind force balance. We include twist as a design variable along the wing. For this study, the centerbody twist has deliberately been left out of the problem to maintain a uniform fuselage shape across the span. Our methods are not capable of capturing penalties for twisted fuselages or the effects of separation and stall from high inflow angles. As a result, we also limit the pitch angle to tight bounds to avoid an unreasonable deck angle.

In addition to inflow variables, we include a set of planform variables. First, we assign three span variables to control the width of the centerbody, the transition region, and the outer wing. Next, we assign chord design variables to the centerbody and the wing. We treat the transition region as a blending region between the centerbody and outer

Table 2 For our multipoint approach, we devised a set of nine cruise missions for the BWB MDO demonstration, populating the range-payload space. We fix the speed to Mach 0.7 and the altitude to 30000 feet.

Mission	Payload (lbs)	Range (nmi)	Mach	Altitude (ft)
1	160000	2500	0.7	30000
2	160000	2000	0.7	30000
3	160000	1500	0.7	30000
4	100000	2500	0.7	30000
5	100000	2000	0.7	30000
6	100000	1500	0.7	30000
7	50000	2500	0.7	30000
8	50000	2000	0.7	30000
9	50000	1500	0.7	30000

wing, and intentionally omit any additional design variables in this region. We include wing sweep to address wave drag, moment trim and stability. In the centerbody, we include a trailing-edge airfoil reflex angle. We treat this design variable like a static elevator deflection that is fixed to the geometry. We do not include active control surface deflection for each mission in this demonstration. Although this limits the design scope, it provides time and memory savings, as the geometry and mesh update are not reevaluated for each condition. More information on the cost savings is provided in Section V. We also include design variables to control the structural mesh. We include 16 thickness b-spline coefficients as design variables, which control the thickness of the spars, ribs, and the top and bottom skin. Finally, we include a chord-wise engine placement design variable to let the optimizer select engine location. This variable will become active with the stability constraint, as the engine placement plays an important role in weight distribution. Our framework has physical limitations due to the nature and assumptions of the methods. To produce reasonable results, we want to prevent the optimizer from entering regions of the design space that our models cannot sufficiently penalize. For example, we tightly bound the pitch and twist variables because we do not have reasonable models to predict the penalty associated with stall and separation. On the other hand, we set larger bounds for some design variables in which our models can adequately penalize. This includes planform variables, as we incorporate reasonable weight and structural penalties that the optimizer will naturally try to avoid. The geometry design variables are depicted visually in Figure 4.

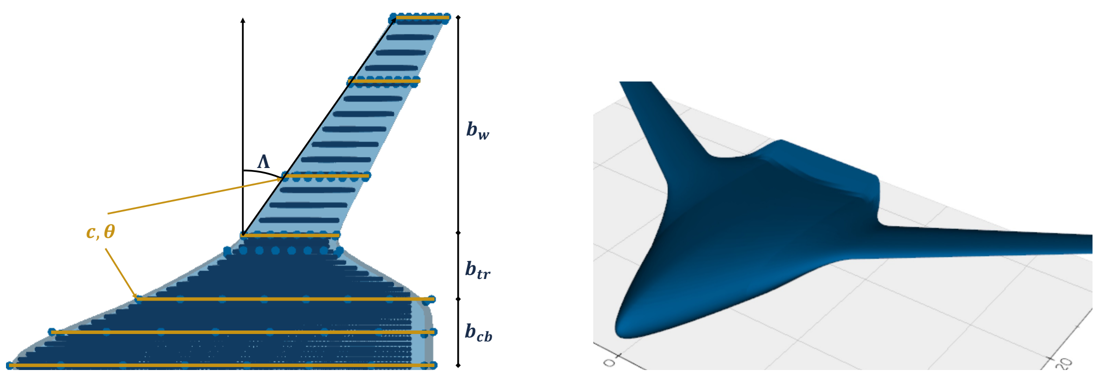

(a) We consider many design variables related to the BWB geometry.

(b) The centerbody reflex angle bends the trailing edge to generate pitching moment.

Fig. 4 We assign a large number of design variables relative to the geometry, such as chord c, twist θ, wing sweep Λ, sectional spans b_{cb}, b_{tr}, b_w, and the centerbody reflex angle δ_r.

B. Constraints

We assign a vertical force balance constraint to each cruise and maneuver condition. The primary driver of this constraint is the pitch angle. For the two maneuver conditions, we introduce max stress constraints. We use Aluminum 2024 for the structural material with a yield stress of 324 MPa. To ensure sufficient operation of the aircraft in the maneuver conditions, we apply a 1.5 safety factor as a buffer to the max stress calculation. We enforce moment trim and static margin at the nominal median condition in the payload-range space. A more appropriate approach to modeling stability is to constrain the static margin at the MTOW and empty weight missions. However, we are unable to satisfy these constraints at more than one condition without active trim surface control. We acknowledge this as an aspect for future work.

We introduce a series of constraints to prohibit the optimizer from generating unrealistic geometries. First, we introduce two volume constraints for the BWB. First, we apply a centerbody volume constraint based on the 747-400 fuselage volume; we apply a generous 1.5 safety factor as a conservative measure to not undersize the aircraft. We utilize this as a constraint due to the similar max payload weight. Next, we apply a fuel volume constraint in the transition region. We assume the transition region has 50% usable volume for fuel storage and apply that as a lower bound for our transition region volume. We also include a centerbody trailing edge sweep constraint to avoid abnormal and oblong centerbody shapes. We anticipate that the optimizer will attempt to lengthen the centerline to generate a long moment arm for the moment trim constraint. In addition, we apply an engine spacing constraint to the centerbody. This constraint avoids the centerbody from becoming too thin.

Finally, we include takeoff and landing length constraints using the formulation from Section III. We enforce that the takeoff and landing lengths are under 10000 feet. We enforce these constraints using the Telluride airport altitude of 9000 feet.

Table 3 Our MDO problem formulation considers 44 design variables to control operational and geometry variables. We assign 21 constraints related to trim, stability, structural stress, volume, and landing and takeoff lengths. The three sectional span terms correspond to the centerbody, transition, and wing, respectively. The chord distribution bounds are set relative to the baseline distribution c_0.

	Variable	Description	Quantity
minimize	f	Weighted fuel burn	**1**
with respect to	$-5 \leq \alpha \leq 3$	pitch angle (°)	11
	$(3, 3, 8) \leq (b_{cb}, b_{tr}, b_w) \leq (9, 8, 22)$	sectional span (m)	3
	$15 \leq \Lambda \leq 50$	wing sweep (°)	1
	$-5 \leq \theta \leq 5$	wing twist distribution (°)	4
	$0.5c_0 \leq c \leq 1.5c_0$	chord distribution (m)	7
	$-5 \leq \delta_r \leq 10$	centerbody reflex (°)	1
	$1 \leq t_{struct} \leq 100$	shell thickness coefficients (mm)	16
	$0.6 \leq ecf \leq 0.85$	engine chord placement	1
Total design variables:			**44**
subject to	$L = W$	force balance	11
	$1.5\sigma_m \leq \sigma_{yield}$	max stress	2
	$C_{My} = 0$	nominal moment trim	1
	$K_n = SM$	nominal static margin	1
	$\Lambda_{TE} = 0$	centerbody TE sweep	1
	$\Delta y_e > 0$	engine collision spacing	1
	$L_{TO} \leq 10000\text{ft}$	max takeoff length	1
	$L_{LA} \leq 10000\text{ft}$	max landing length	1
	$V_{cb} \geq 1.5V_{cargo}$	payload volume	1
	$V_{tr} \geq 2V_f$	fuel volume	1
Total constraints:			**21**

C. Optimization Problem Formulations

Table 3 summarizes the design variables and constraints of our BWB MDO demonstration. To build up to the full-scale problem, we solve a series of smaller optimization problems. These subscale optimization problems are described in Table 4. We divide our set of optimization problems into three groups: baseline, planform and stability. The baseline optimization problem serves to find an optimal structural thickness distribution and pitch angle across

Table 4 We solve a series of optimization problems to build up to our full-scale problem. The breakdown for each sub-problem is shown here.

		Baseline	Planform		Stability		
			No Trim	Trim	Loose	Relaxed	Tight
		P0	P1	P2	P3	P4	P5
Design variables	α	✓	✓	✓	✓	✓	✓
	t_{struct}	✓	✓	✓	✓	✓	✓
	θ		✓	✓	✓	✓	✓
	c		✓	✓	✓	✓	✓
	b		✓	✓	✓	✓	✓
	Λ		✓	✓	✓	✓	✓
	δ_r			✓	✓	✓	✓
	ecf				✓	✓	✓
Constraints	$L = W$	✓	✓	✓	✓	✓	✓
	$1.5\sigma_m \leq \sigma_{yield}$	✓	✓	✓	✓	✓	✓
	Volume constraints		✓	✓	✓	✓	✓
	L_{TO}, L_{LA}		✓	✓	✓	✓	✓
	$\Lambda_{TE} = 0$		✓	✓	✓	✓	✓
	$\Delta y_e > 0$		✓	✓	✓	✓	✓
	$C_{My} = 0$ (nominal)			✓	✓	✓	✓
	$K_n = SM$ (nominal)				0%	3%	5%

the cruise missions that can satisfy vertical force balance and material stress limitations. From there, we solve a pair of optimization problems with planform variables. We now include the geometry variables to assess the influence of geometry changes on the BWB performance. To assess the impact of the moment trim constraint, we first solve an untrimmed problem. Finally, we move to the stability-constrained problems. We solve for BWB configurations that achieve longitudinal static stability with incrementally tighter static margin constraints. For these problems, we use a warm-starting approach to accelerate optimization convergence. Because we incrementally build up to the tight 5% stability constraint, we expect the stability-constrained optimal designs to exhibit monotonic geometric trends as we tighten the static margin constraint. For the remainder of the paper, we will often refer to the optimization cases with their shorthand names: P0, P1, etc., for simplicity.

Figure 5 shows the optimization workflow, using methods described in Section III. The 11 evaluations of the panel method are used to compute aerodynamic performance parameters for the structural analysis and fuel burn computation using the Breguet range equation 12.

V. Numerical Experiments of the MDO Framework

In this section, we present results for two numerical experiments conducted prior to the BWB MDO demonstration. First, we assess the computational cost benefits and the aerodynamic performance error from fixing the linear system matrix of the panel method. Next, we present results on numerical experiments assessing the memory and time scaling of the MDO demonstration with respect to the number of design conditions.

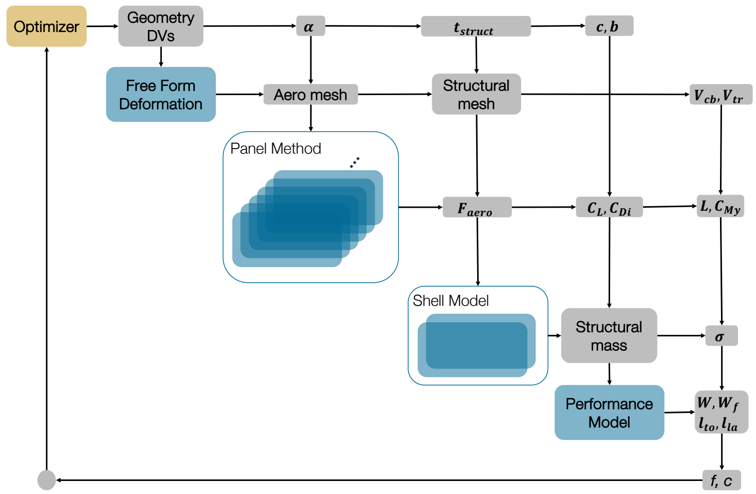

Fig. 5 XDSM showing the workflow of the optimization problem

A. Experiments on Fixing the Panel Method Wake Angle

The assembly of the linear system matrix A in Equation 3 incurs the highest computational cost within the panel method. Linear system assembly exhibits quadratic scaling relative to grid size because the matrices A and B represent the induced potential of each panel by each panel. Although we can limit the memory cost of $B\vec{\sigma}$ using batched matrix-vector products, the assembly of matrix A limits the capability of solving problems with large grids. In addition, this limits the tractability of solving problems with many design conditions, as each matrix A is dependent on the inflow angle through the Kutta condition. The Kutta condition relates the wake doublet strength to the trailing edge of the body as

$$\mu_w = \mu_u - \mu_l, \tag{21}$$

where μ_w represents the wake doublet, and the subscripts u, l represent the corresponding trailing edge elements on the upper and lower surface, respectively. By recomputing the matrix A for each design condition, the memory and time scale with the expensive matrix assembly. For this demonstration, we assume that the wake angle has a minimal influence on the overall solution. With this assumption, the linear system matrix A can be reused across conditions. The Kutta condition is still enforced, but the wake propagates without influence from the inflow angle.

With this adjustment, we alleviate the prohibitive memory and time cost of the panel method to re-evaluate each linear system matrix for each design condition. Figure 6 highlights the forward and adjoint time and memory cost as a function of the number of missions. We see that there is an upfront cost for one mission; this is representative of the

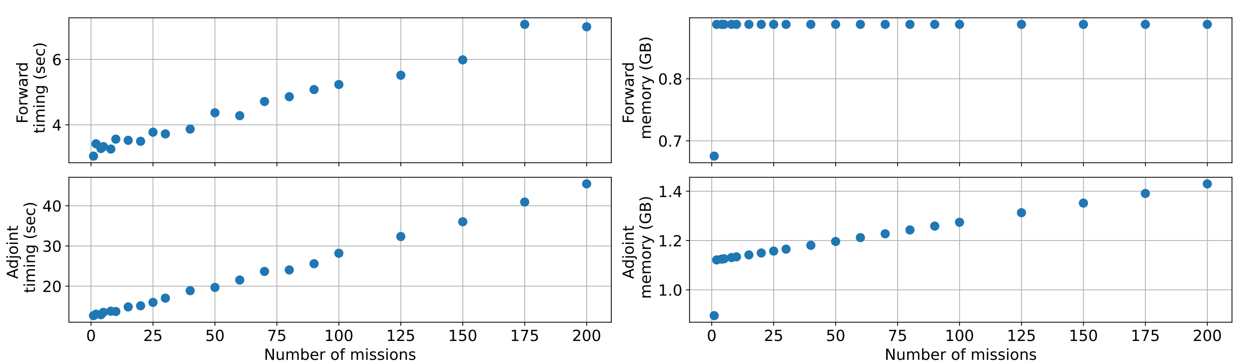

Fig. 6 Scaling of forward and adjoint evaluation time and memory vs. number of cruise conditions for the panel method. We see a steep upfront cost of one model evaluation for the AIC matrix generation. However, each subsequent evaluation is low-cost in time and memory.

expected cost of one linear system matrix generation. However, we see that additional missions require far less time and memory to solve, as the remaining computations within the panel method incur an inconsequential cost. Each additional mission incurs an extra 0.02 seconds and 0.17 seconds for forward and adjoint evaluations, respectively. When compared to the upfront cost of linear system assembly, which is around 3 seconds for the forward evaluation and 15 seconds for the adjoint evaluation, this benefit is significant.

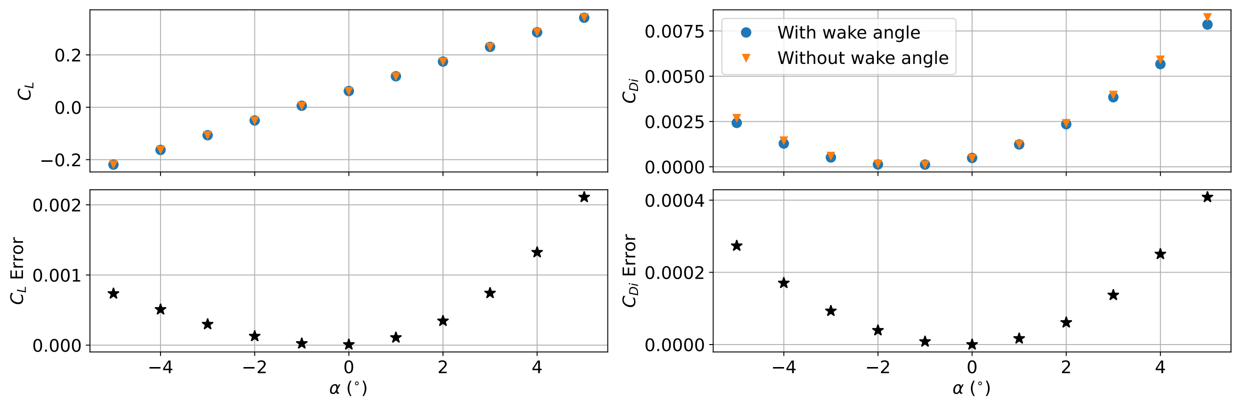

Fig. 7 We assessed the error on C_L and C_{Di} when neglecting the wake deflection angle in the panel method, using the default panel method as our reference. Within the bounds of pitch angle for our demonstration, we see errors of around 1% and 5% for C_L and C_{Di}, respectively.

As mentioned previously, the validity of this approach relies on the assumption that the wake angle plays an insignificant role. To test this hypothesis, we conducted a sweep over the angle of attack to assess the error of C_L and C_{Di} when neglecting the inflow angle on the wake, as shown in Figure 7. Our reference for the error is the baseline panel method with no adjustments to the linear system matrix. We see reasonable agreement between the aerodynamic coefficients when we ignore the angle of attack influence on the wake; the maximum errors of C_L and C_{Di} are around 1% and 5%, respectively. We accept these errors as part of a tradeoff, given that this technique provides significant benefits to computation time and memory.

B. Optimization Problem Scaling with Design Conditions

To assess the scaling of the entire optimization problem, we conducted numerical experiments on the graph size, compile time, and total memory of the optimization problem. Compile time refers to the initial setup time required for CSDL to assemble the computational graph, which represents the sequence of operations that define the forward evaluation and adjoint derivative computation. This is also reflected in the number of forward and adjoint graph nodes. We assessed these outputs by testing i^2 missions for integers $i \in [1, 10]$. The data for the graph assembly versus the number of missions is shown in Figure 8. We see that the overall operation count, except for the one mission case, does not vary as we increase the number of missions. This is a direct benefit of vectorization, made possible through CSDL. Graph size, among other factors, also contributes to memory usage and compile time. As missions are varied, we also see minimal variation in memory usage. The high upfront memory cost is attributed to the assembly of other aspects of the MDO framework, such as the costly FEM evaluation for the structural sizing conditions. However, the memory cost of subsequent forward and adjoint evaluations, despite going up to 100 conditions, is insignificant. Each mission involves an evaluation of the panel method and the performance model to compute high-level optimization variables. We see the benefits of the wake angle independence in the panel method formulation, as described in the results of the previous numerical experiments in Section V.A. The compile and run time, on the other hand, exhibits unfavorable exponential scaling as the number of missions grows. This scaling is driven by the vectorized vertical force balance constraint, which expands as the number of missions increases. Each trim constraint, however, is independent of the other missions. As a result, we incur redundant derivative evaluations across missions; the number of redundant derivative computations scales roughly quadratically with the number of missions. Although the compile time is manageable for cases under 40 to 50 missions, there is potential to improve the efficiency of our MDO framework by removing the redundant, empty derivative evaluations. These trends will help inform how far we can expand our optimization problem formulation to reach an unprecedented level of complexity for future MDO demonstrations.

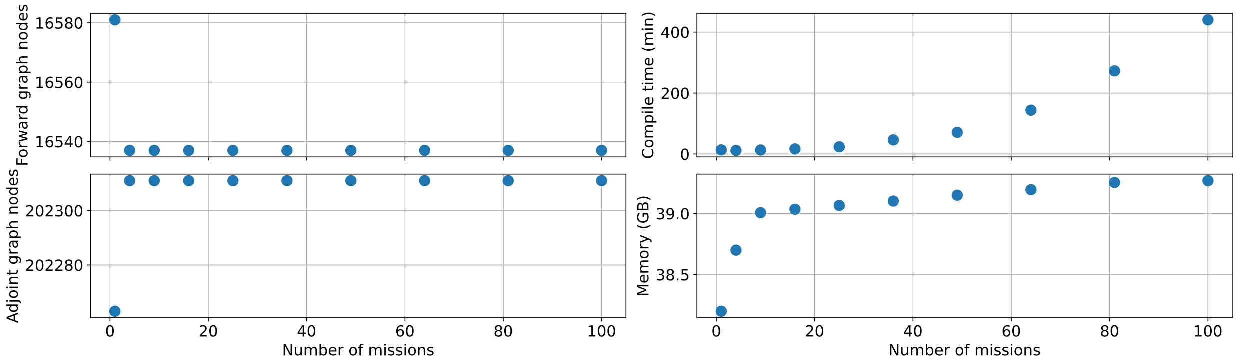

Fig. 8 Scaling of the optimization problem compile stage vs. number of cruise conditions. In our framework, the number of forward and adjoint graph nodes is independent of the number of missions. In addition, the memory scaling is favorable with each additional mission. The compile and run time sees unfavorable exponential scaling.

VI. MDO Demonstration on a Blended Wing Body Configuration

In this section, we discuss the results of the 6 optimization cases. Due to the modularity of our MDO framework, we can set up and run each optimization case with minimal difficulty. The efficiency of our MDO framework also leads to a reasonable turnaround time for each optimization case, allowing us to study many optimized designs in a short time. First, we compare high-level outputs from the optimization results. These are shown in Table 5. The baseline optimization problem allowed pitch and structural thicknesses to vary to reach an initial feasible design. Although there are no constraints set on moment trim or static margin, it is important to note the trends of the initial configuration to understand how the optimizer will drive the design towards feasibility in future cases. With the negative moment coefficient, we see that the BWB configuration has a natural nose-down tendency. Although the optimized baseline configuration is considered stable, there is a need for a counter-moment mechanism to oppose the nose-down tendency. The initial introduction of planform variables from P0 to P1 allows for a significant improvement to overall performance, with a 13.6% reduction in the weighted fuel burn and a 25.3% reduction in maximum takeoff weight. Although there is a less-severe nose-down tendency for the P1 converged design, the configuration becomes unstable. The reduction in the planform shape likely shifted the neutral point too far forward with minimal movement of the center of gravity (CG). When introducing a moment constraint in case P2, there are minimal differences to the converged designs. The moment constraint penalizes the weighted fuel burn by a small fraction. Although the additional moment trim constraint leads to a less optimal solution than in case P1, it is important to note that the max-payload, max-range condition of P2 operates more efficiently than P1. This highlights the importance of a multipoint design approach over a single-point approach. In cases P3 to P5, where stability constraints are introduced and slowly tightened, we begin to see significant increases in the overall weights and hindrances in aerodynamic performance. The need to trim and stabilize the aircraft reduces the lift-to-drag ratio and necessitates a higher lift coefficient to satisfy trim. One key factor shared by the results of each optimization case is the wing loading; typical tube-and-wing configurations experience a wing loading in excess of 100

	Overall outputs			MTOW mission				Nominal condition		
Case	W_f (lbs)	Structural weight (lbs)	W/S (lbs/ft^2)	TOGW (lbs)	W_{f1} (lbs)	L/D	C_L	C_M	Static margin (%)	Time (hrs)
P0	33319	81663	69.3	393802	47244	27.61	0.323	-0.042	11.1	0.74
P1	28783	34010	50.3	294207	43669	21.96	0.234	-0.005	-5.6	7.40
P2	29370	32055	47.6	290012	43432	21.75	0.221	0	-6.2	10.21
P3	36161	50103	47.7	334717	50607	21.53	0.222	0	0	11.25
P4	44538	67091	51.1	381349	61153	20.19	0.238	0	3	16.01
P5	60331	117460	60.2	506328	81354	20.14	0.280	0	5	9.38

Table 5 Summary of high-level results across the various optimization cases.

lbs/ft^2. We can attribute the lower wing-loading of the BWB to the structural efficiency and aerodynamic performance, most notably again from the lifting shape of the centerbody. The weight breakdown for each optimization case is also shown in Figure 15 of the Appendix section IX. In addition to stability being a necessary aspect of safe and reasonable flight operations, the stability-constrained solutions converge to more reasonable weight fractions, according to existing studies done in the literature [32]. Without stability constraints, the design does not experience performance penalties that mimic aircraft operation. As a result, the designs converge to unrealistic payload weight fractions in excess of 50% with structural weight fractions close to 10%. Noticeable weight penalties appear as the planform is sized up to meet moment trim and stability constraints. As the stability constraint is tightened, the empty weight and structural weight fractions approach reasonable ratios of roughly half and a quarter of the total weight, respectively. We also see reasonable solve times for each optimization case. Most of the optimization cases run quickly enough to converge overnight. As expected, the overall optimization run time increases with problem complexity. Although warm-starting is used to solve the stability-constrained problems P3 to P5 using the previously converged solution, the run times remain relatively unchanged due to the competing constraints. The static margin constraints tend to diminish the lift and aerodynamic efficiency, which negatively impacts the trim constraint and the objective. The noticeable increase in the case P4 run time is attributed to the static margin constraint scaling from the warm-started solution. The optimizer struggled to exit the local design variable space for the 0% static margin case to tighten the stability constraint. Overall, this data highlights the efficiency of our MDO framework to conduct full-scale optimization studies with reasonable convergence time.

We can also compare the optimized results against existing aircraft that operate at similar ranges and payload weight, such as the Boeing 737-M7 and 747-400. The BWB is optimized for payload weights similar to the 747 at flight ranges closer to that of the 737; we anticipate the optimized fuel burn and total MTOW to converge somewhere in between. This trend is shown in Figure 9a. Despite carrying a similar payload weight at MTOW, the optimized BWB designs operate well below the MTOW of the 747, even with the tightest stability constraint. This difference is likely due to the design range of the BWB in this study. We see similar fuel burn between the 737 and optimized BWB designs, despite significantly higher payload weight for the BWB. This highlights the aerodynamic efficiency benefits of the BWB, with the additional centerbody lift generation and reduced drag penalties from component interference and large wetted-area-to-volume ratios prevalent in traditional tube-and-wing aircraft.

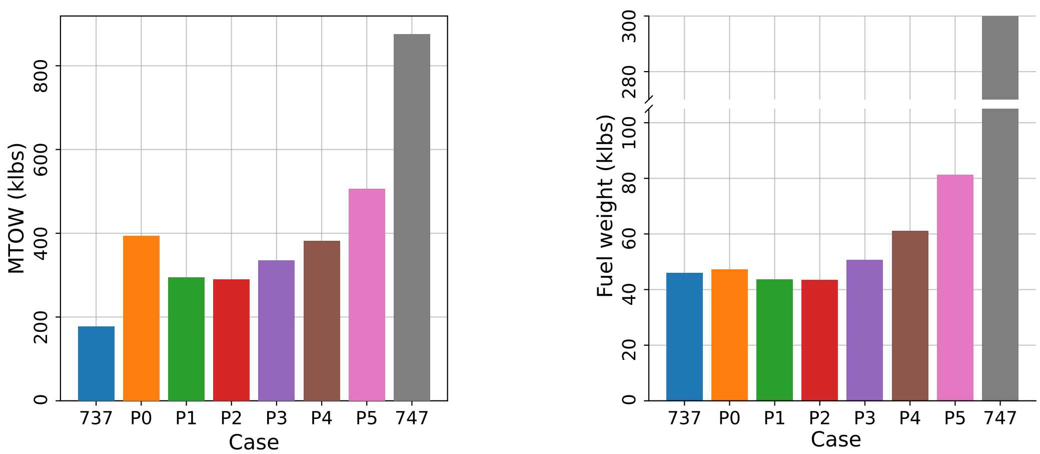

(a) Comparison of MTOW **(b) Comparison of fuel burn at MTOW.**

Fig. 9 Comparison of properties at MTOW for the six optimization cases to the Boeing 737 and 747 configurations.

The planform shape is a key aspect of the coupling between aerodynamics and stability. The optimized planform geometries for each case are shown in Figure 10. The initial introduction of planform variables from P0 to P1 allows for a significant improvement to overall performance by first sizing down the aircraft to reduce weight. In the absence of moment and stability constraints, the optimizer can drastically reduce the planform area to satisfy vertical force balance and minimize drag penalties, due to the significant centerbody lift generation. With the introduction of a moment constraint in case P2, there are minimal differences to the converged designs, shown in Figure 10e. The centerbody

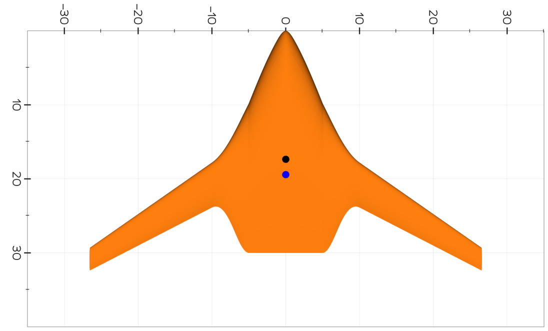

(a) P0: The baseline configuration

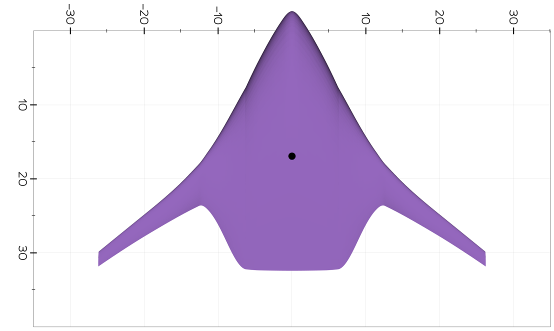

(b) P3: Loose stability design (SM = 0%)

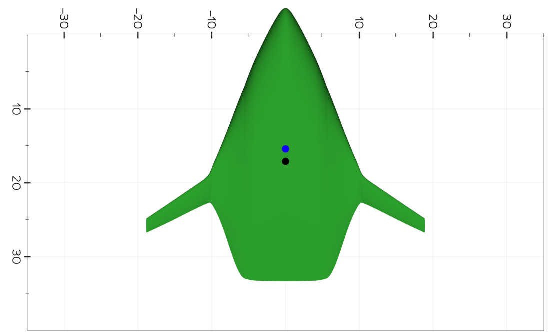

(c) P1: Planform design without moment trim

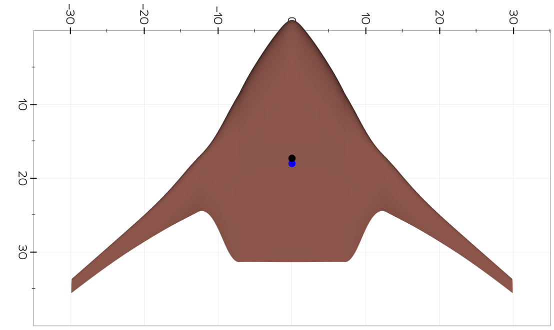

(d) P4: Relaxed stability design (SM = 3%)

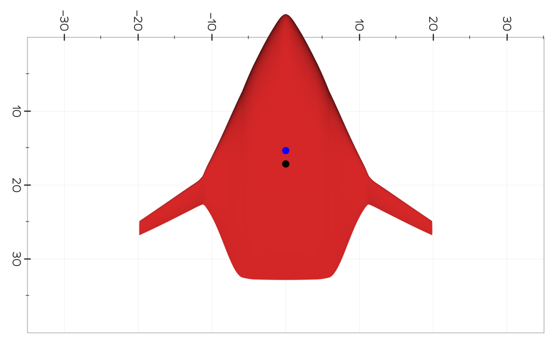

(e) P2: Planform design with moment trim

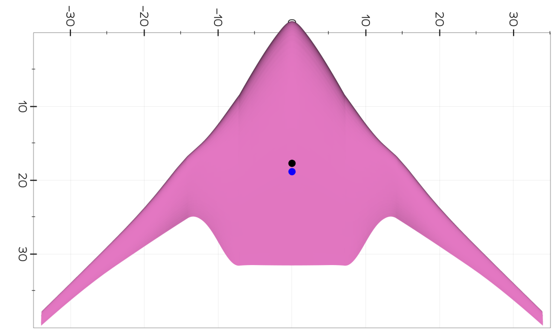

(f) P5: Tight stability design (SM = 5%)

Fig. 10 Comparison of the optimized planform geometries for the six optimization cases. The black dot represents the aircraft's center of gravity, and the blue dot represents the aircraft's neutral point.

airfoil reflex is the main driver to satisfy moment trim. To counteract the centerbody lift reduction from the trailing-edge airfoil reflex, the optimizer increases the overall wingspan to produce lift at the outer wing sections. As we see, these cases converge to somewhat unreasonable solutions due to the lack of necessary physical constraints within the problem. Cases P3 to P5 highlight the significance of the coupling between stability and aerodynamic performance. We begin to see large increases in the planform geometry, as seen in Figures 10b, 10d, and 10f compared to the solutions with no stability constraints in Figures 10c and 10e. With no tail, the BWB can utilize two mechanisms to trim moments and stabilize: the centerbody trailing-edge airfoil reflex, as mentioned previously, and wing twist coupled with the wing sweep. We see the second mechanism play a noticeable role in achieving stability. As the nominal condition static margin constraint is tightened from a loose tolerance of 0% up to 5%, the geometry exhibits more severe sweep and washout at the wing. This also directly hinders the aircraft's aerodynamic efficiency, as shown in Table 5.

These trends are also evident in the pressure coefficient distributions. Figure 11 shows a top-half view of the pressure coefficients over the optimized geometry. The pressure coefficients exhibit similar distributions in the centerbody region. This is because the centerbody utilizes a symmetric airfoil at a relatively consistent pitch angle across all optimization cases. Additionally, no twist has been applied to the centerbody section. For the cases where stability constraints are enforced, there is a noticeable trend in the spanwise variation of the pressure distributions. In figures 11a, 11b, and 11c, the entire outboard wing section produces lift, as seen from the blue low-pressure zones. The outboard wings are contributing to the vertical force balance constraints. However, we see a different trend with the stability-constrained solutions in figures 11d, 11e, and 11f. The low-pressure zone is prominent in the transition region, between the centerbody and the outboard wing. This low-pressure quickly fades towards the wing tip as the wing is swept back. A tip-loaded wing at a severe sweep angle significantly contributes to a nose-down moment tendency. Although this can be counteracted with more airfoil reflex at the centerbody, this combination is likely to reduce aerodynamic efficiency, negatively impacting the fuel burn and overall weight.

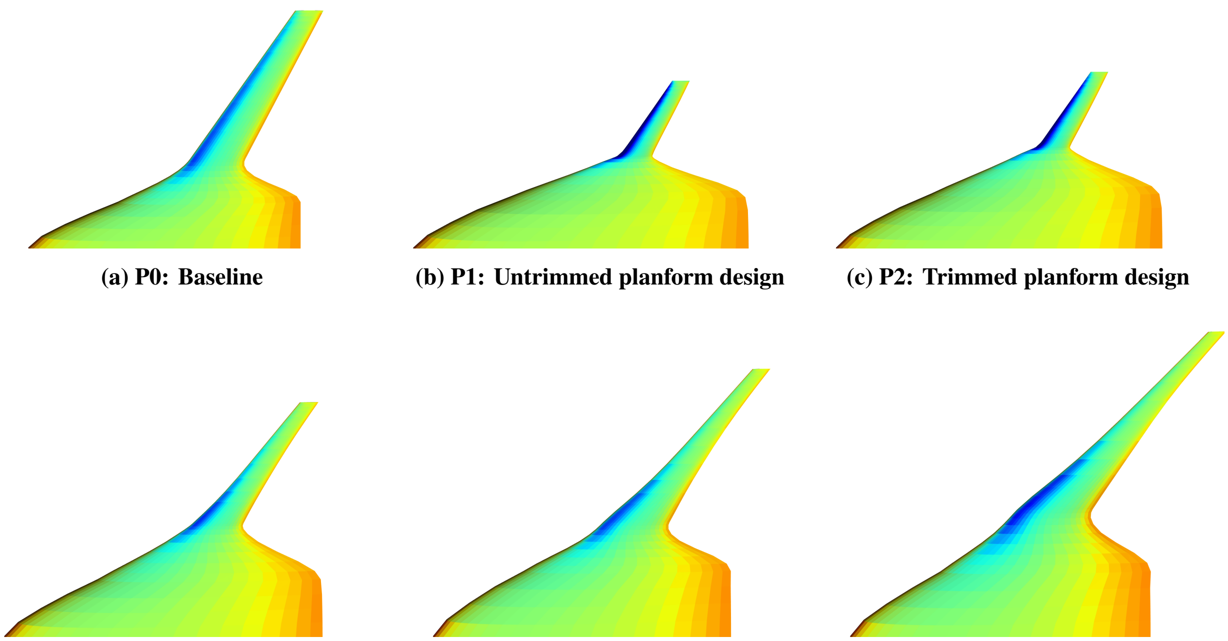

(a) P0: Baseline **(b) P1: Untrimmed planform design** **(c) P2: Trimmed planform design**

(d) P3: Loose stability (SM = 0%) **(e) P4: Relaxed stability (SM = 3%)** **(f) P5: Tight stability (SM = 5%)**

Fig. 11 Comparison of the pressure coefficients of the upper surface of each optimized design. The dark blue zones highlight low-pressure zones, while red denotes stagnation. The pressure coefficients are computed using the panel method.

As previously mentioned, the tight coupling of the aerodynamic performance and stability behavior is a major driver of the BWB configuration. The mechanics of this coupling can be further analyzed by assessing spanwise and sectional properties. The central geometry parameterization simplifies the evaluation of geometry parameters and flow-field properties using the panel method to assess these trends. Figure 12 shows the spanwise lift distribution and inflow angle distribution for the three stability-constrained optimization cases. The inflow angle represents the sum of the sectional geometric twist and the operating pitch angle. The results presented here are for the nominal cruise condition of each optimization case. The small vertical markers along each curve represent the boundaries of the three main wing sections. In the stability-constrained designs, the transition region and root section of the wing contribute significantly to the overall lift generation. The sectional lift plateau towards the center line (spanwise section of zero) can be explained by numerous factors: the centerbody airfoil reflex needed for a nose-up moment through downlift generation, the symmetric centerbody airfoil, and the fixed geometric centerbody twist distribution combined with a moderate pitch angle. The aft centerbody airfoil reflex, despite reducing centerbody lift, cannot be sacrificed as it is the driving force behind reaching moment trim. As a result, the optimizer increases the geometric twist at the transition region and wing root to generate sufficient lift and meet the vertical force balance constraint. However, the geometric twist quickly drops to avoid excess lift production towards the outward wing section. This serves to produce minimal nose-down tendency and help with

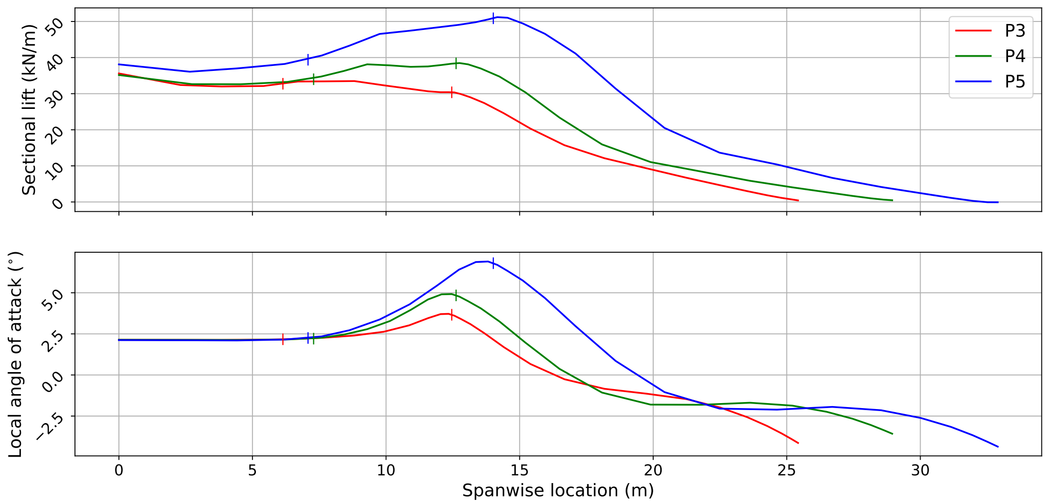

Fig. 12 The spanwise lift and local angle of attack distributions highlight regions of notable lift generation or lift reduction. The vertical bars represent the boundaries between the centerbody, transition, and wing regions. The tendency to impose severe washout at the wing tip suggests a lift taper to prevent further nose-down moment.

stability and moment trim. It is important to note that the airfoils used in the outboard wing section have a noticeable camber. To generate minimal lift, the optimizer introduces severe washout to produce a negative inflow angle.

The sectional pressure coefficient distributions and airfoil shapes provide additional insight about the tight coupling between the aerodynamic and stability behavior, as shown in Figure 13. In the 4x4 grid plot, each column represents a different spanwise section; η_s refers to the normalized spanwise section. The first row shows the sectional airfoil thickness to chord ratio; the x-axis represents the normalized chord location η_c. The three rows below each depict sectional pressure coefficient distributions for the three stability-constrained optimization cases; these cases are labeled on the right margin. The upper and lower surface pressure coefficients are plotted separately to emphasize pressure inversion. We evaluate properties for $\eta_s = [0.15, 0.4, 0.55, 0.95]$ to inspect variations in sectional pressure at key sections along the span.

At the spanwise section $\eta_s = 0.15$, there is a relatively small pressure differential between the upper and lower surfaces towards the leading edge. We see a noticeable pressure inversion towards the trailing edge, where downlift is locally produced. This is a key indicator that the centerbody airfoil reflex is driving the moment trim constraint. As the static margin is tightened to 5% for P5, the trailing edge pressure inversion becomes more pronounced, highlighting more centerbody airfoil reflex to maintain moment trim and stability. The transition and wing root sections at $\eta_s = 0.4$ and 0.55 have a large pressure differential between the upper and lower surfaces of the airfoil, indicating the significant lift generation at this spanwise section. The airfoil shapes in this region are a blend between the symmetric centerbody airfoil and the cambered wing airfoil. The optimizer is exploiting both a geometric twist and a slight natural camber to generate lift. Towards the outboard section of the wing at $\eta_s = 0.95$, a severe washout angle is used to minimize lift production; this is seen in the airfoil shape in Figure 13 and the inflow angle distribution in Figure 12. The flow at the forward section of the upper airfoil surface is stagnated. Beyond the leading edge, there is minimal pressure differential across the chord. These trends substantiate the reduced lift production towards the wing tips to satisfy moment trim.

In the stability-constrained optimization cases P3 to P5, the static margin is constrained only at the nominal cruise condition, operating at the midpoint of the range and payload space. This work does not include active trim surface deflection to trim moments and thus does not apply a moment trim and stability constraint at each design condition. Variations in range and payload weight will inevitably influence the weight distribution of the aircraft, altering the CG and subsequently the static margin. As a result, the static margin may become negative for some combination of range and payload weight. We explored the relationship between gross weight and static margin across the nine different cruise missions for the three stability-constrained problems, as shown in Figure 14. Variations in color (red, green, and blue) represent different cruise range values. Variations in the plot symbols (circle, star, and triangle) represent different

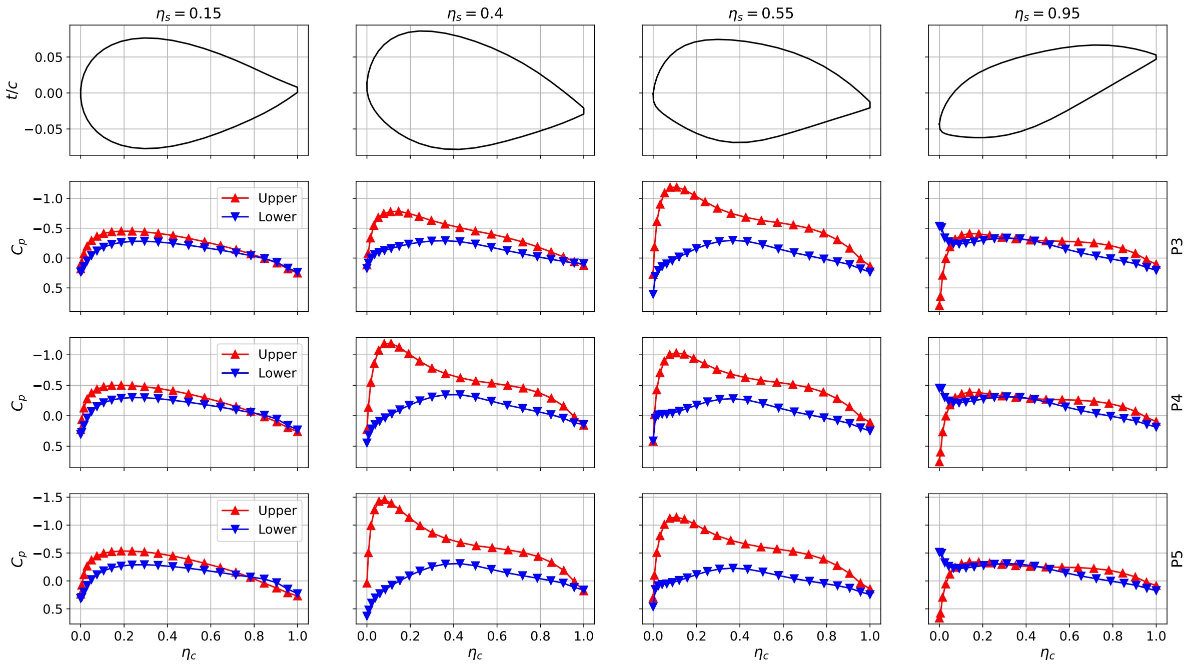

Fig. 13 The sectional airfoils and pressure distributions provide further insight into the spanwise lift variation. The pressure inversion at the centerbody trailing edge highlights the impact of the reflex angle to satisfy moment trim. Large pressure deltas in the mid-wing section around transition suggest significant lift generation. Meanwhile, the stagnated tip section generates close to zero lift, despite having a cambered airfoil.

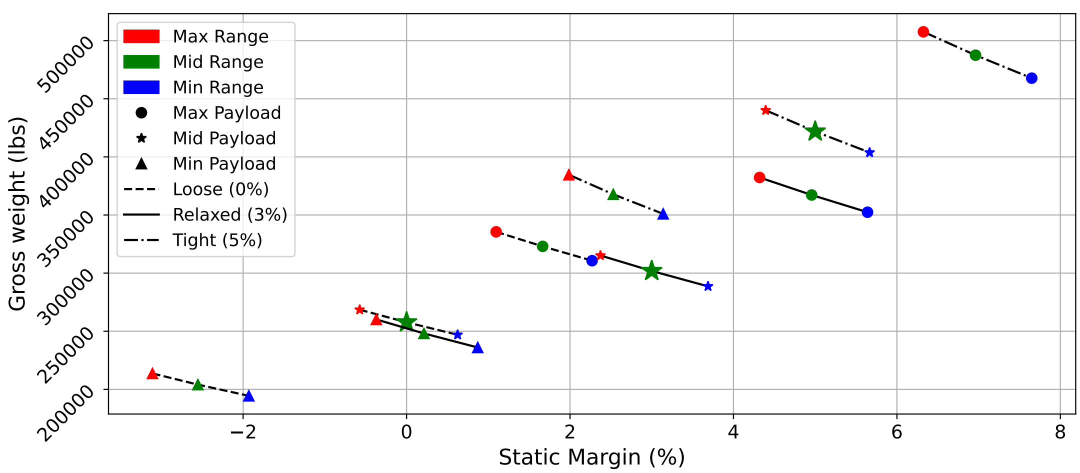

Fig. 14 Static margin vs. gross weight evaluation across all cruise design conditions for the three stability cases P3, P4, and P5. The large green star denotes the nominal condition in which the static margin constraint is enforced. Longitudinal static stability is satisfied for each multipoint analysis only in the tight 5% nominal static margin optimization case.

payload weights. The differing line styles represent the optimization case (P3, P4, and P5), each with a different nominal condition static margin constraint. For a given line style, payload weight increases lead to larger static margin, as the payload CG is in front of the aircraft CG. Meanwhile, increases in range lead to a decrease in static margin. The fuel tanks are stored in the transition region of the BWB, positioned behind the CG. Any further weight additions behind the CG will shift it further aft, potentially behind the neutral point. Of the three stability-constrained optimization cases, we find that only the tight 5% nominal condition static margin case is stable at each point in the payload-range space. This further highlights the benefit of a multipoint approach to capture multiple points in the design space; a single-point approach would not be able to quantify changes to stability properties at other trimmed conditions. In this work, we take a static approach to the multipoint mission segment modeling. As a result, we do not model changes to the weight distribution in cruise as fuel is burned. However, these trends provide insight into how stability properties vary as weight is lost. Along a payload weight isoline, we can effectively simulate the effects of fuel burn in cruise by moving from the maximum range point (in red) to the minimum range point (in blue). Along this curve, we see an increase in static margin. Given that the fuel storage space also sits aft of the BWB CG, we expect that weight reductions due to fuel burn during cruise will lead to a more stable aircraft.

VII. Conclusion

In this work, we have presented a demonstration of large-scale, multi-point, multidisciplinary design optimization (MDO) of a blended wing body (BWB) configuration with adjoint-based sensitivity analysis. We developed an efficient graph-based MDO framework tailored towards simulating a large number of on- and off-design conditions. We applied a mid-fidelity panel method, a high-fidelity shell model, and numerous empirical models to assess BWB performance across many design conditions.

We conducted numerical experiments on the panel method with a constant wake-propagation angle to assess computational cost and error. We see errors of 1% and 5% for the lift and the induced drag coefficients, respectively. Although these errors are not trivial, we deem it acceptable to capture reasonable aerodynamic trends for the sake of this demonstration. We then conducted numerical experiments to assess the memory and time cost scaling versus the number of cruise missions of our MDO framework. The favorable memory scaling of the optimization problem provides potential to scale up to an even larger number of cruise conditions. Overall, these experiments provide insight for future studies to expand our MDO problem complexity. With our MDO framework, we successfully optimized a BWB configuration across a series of cases, up to a tightly-bound static margin constraint of 5%. Our graph-based MDO framework enabled efficient optimization studies by using vectorization and adjoint-based sensitivity analysis. In addition, the use of CSDL and modOpt simplified warm starting, allowing a sequential approach to solve optimization problems of increasing complexity. As a result, we were able to quantify significant trends and tradeoffs between the tightly-coupled aerodynamic performance and stability metrics of the BWB across various static margin constraints. Despite the reduced fidelity of the models, we demonstrate the ability of our MDO framework to efficiently capture significant design trends across a high-dimensional input space with a multipoint analysis approach.

We identify several directions for future work. First, we aim to introduce airfoil shape variables to exploit additional design improvements to the BWB configuration. This includes airfoil reflex on the wing sections, as well as thickness effects across the entire geometry. To develop a more comprehensive design problem, we also aim to include climb and one engine inoperative (OEI) conditions to perform engine sizing. To further assess trim behavior, we aim to introduce active trim surface deflections across all design conditions. We are interested in conducting configuration-level design studies related to quantities such as engine count, engine placement, control surface locations, and including more on- and off-design conditions. With the addition of more cruise segments, we can explore the BWB design from a design under uncertainty perspective. Finally, we aim to introduce higher-fidelity tools and methods, such as CFD, to more accurately capture compressibility effects and operate at higher Mach numbers. We also aim to include geometric non-interference constraints as a higher-fidelity method to constrain payload shape and configuration.

VIII. Acknowledgments

The authors would like to acknowledge the Multidisciplinary Science and Technology Center of the Aerospace Systems Directorate, Air Force Research Laboratory for funding and supporting this effort through the Collaborative Center for Design and Research of Interdisciplinary Systems program. Distribution A: Approved for public release; distribution is unlimited. AFRL-2026-2042.

IX. Appendix

Here we have a comparison of the weight breakdowns across optimization cases. With too much freedom on the planform variables and no performance penalties to meet stability, the optimizer converges to unrealistic weight distributions, where payload exceeds 50% of the total weight, and the structure accounts for just over 10%. The use of tight stability constraints forces the optimizer to expand the planform shape, introducing a structural penalty for the increased planform area and loads on the lifting body. Subsequently, we see more reasonable weight fractions, as shown in Figure 15f.

References

[1] Liebeck, R. H., "Design of the blended wing body subsonic transport," *Journal of aircraft*, Vol. 41, No. 1, 2004, pp. 10–25.

[2] Brown, M., and Vos, R., "Conceptual design and evaluation of blended-wing body aircraft," *2018 AIAA Aerospace Sciences Meeting*, 2018, p. 0522.

[3] 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., et al., "Boundary layer ingestion benefit of the D8 transport aircraft," *AIAA journal*, Vol. 55, No. 11, 2017, pp. 3693–3708.

[4] Harrison, N. A., Gatlin, G. M., Viken, S. A., Beyar, M., Dickey, E. D., Hoffman, K., and Reichenbach, E. Y., "Development of an efficient M= 0.80 transonic truss-braced wing aircraft," *AIAA Scitech 2020 Forum*, 2020, p. 0011.

[5] Borer, N. K., Derlaga, J. M., Deere, K. A., Carter, M. B., Viken, S., Patterson, M. D., Litherland, B., and Stoll, A., "Comparison of aero-propulsive performance predictions for distributed propulsion configurations," *55th AIAA aerospace sciences meeting*, 2017, p. 0209.

[6] Shah, G., "NASA Quesst Mission Overview," *LAX Community Noise Roundtable Meeting*, 2024.

[7] Johnson, W., and Silva, C., "NASA concept vehicles and the engineering of advanced air mobility aircraft," *The Aeronautical Journal*, Vol. 126, No. 1295, 2022, pp. 59–91.

[8] Kolonay, R. M., "MDAO - An Air Force Perspective," , 2021. URL `https://www.phoenix-int.com/wp-content/uploads/2021/01/Kolonay_PIMDAOEvent2021Rev3.pdf`.

[9] Venkataraman, S., and Haftka, R. T., "Structural optimization complexity: what has Moore's law done for us?" *Structural and Multidisciplinary Optimization*, Vol. 28, No. 6, 2004, pp. 375–387.

[10] Gray, A. C., Riso, C., Jonsson, E., Martins, J. R. R. A., and Cesnik, C. E. S., "High-Fidelity Aerostructural Optimization with a Geometrically Nonlinear Flutter Constraint," *AIAA Journal*, Vol. 61, No. 6, 2023, pp. 2430–2443. https://doi.org/10.2514/1.J062127, URL https://doi.org/10.2514/1.J062127.

[11] McDonnell, T. G., "Gradient-Based Optimization of Highly Flexible Aeroelastic Structures," Ph.D. thesis, 2023. URL https://www.proquest.com/dissertations-theses/gradient-based-optimization-highly-flexible/docview/2866082369/se-2.

[12] Economon, T., Palacios, F., and Alonso, J., "A coupled-adjoint method for aerodynamic and aeroacoustic optimization," *12th AIAA Aviation Technology, Integration, and Operations (ATIO) conference and 14th AIAA/ISSMO multidisciplinary analysis and optimization conference*, 2012, p. 5598.

[13] Fabiano, E., and Mavriplis, D., "Adjoint-based aeroacoustic design-optimization of flexible rotors in forward flight," *Journal of the American Helicopter Society*, Vol. 62, No. 4, 2017, pp. 1–17.

[14] Abdul-Kaiyoom, M. A. S., Yildirim, A., and Martins, J. R. R. A., "Coupled Aeropropulsive Design Optimization of an Over-wing Nacelle Configuration," *Journal of Aircraft*, Vol. 0, No. 0, 0, pp. 1–23. https://doi.org/10.2514/1.C037678, URL https://doi.org/10.2514/1.C037678.

[15] Kenway, G. K. W., and Martins, J. R. R. A., "Multipoint High-Fidelity Aerostructural Optimization of a Transport Aircraft Configuration," *Journal of Aircraft*, Vol. 51, No. 1, 2014, pp. 144–160. https://doi.org/10.2514/1.C032150, URL https://doi.org/10.2514/1.C032150.

[16] Adler, E. J., and Martins, J. R. R. A., "Efficient Aerostructural Wing Optimization Considering Mission Analysis," *Journal of Aircraft*, 2022. https://doi.org/10.2514/1.c037096.

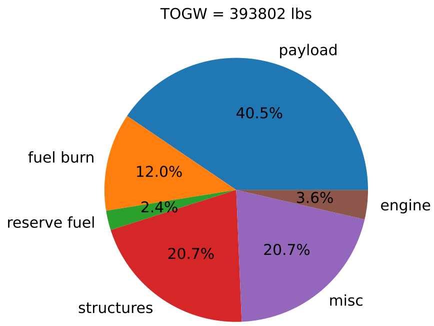

(a) P0: The baseline configuration

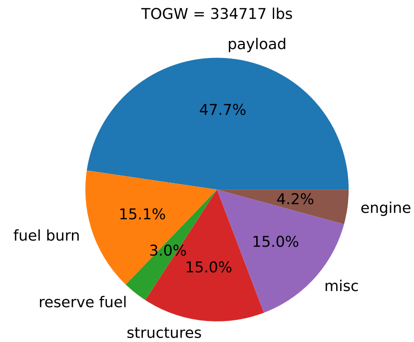

(b) P3: Loose stability design (SM = 0%)

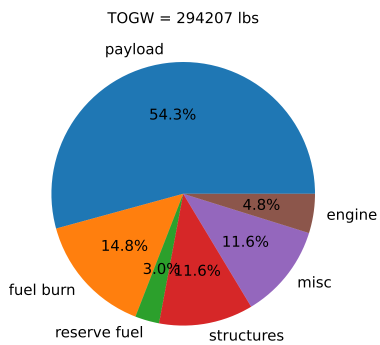

(c) P1: Planform design without moment trim

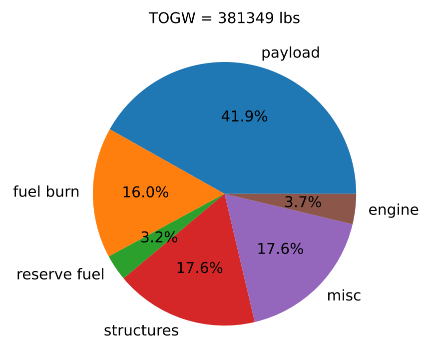

(d) P4: Relaxed stability design (SM = 3%)

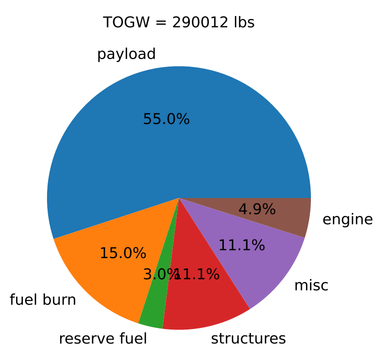

(e) P2: Planform design with moment trim

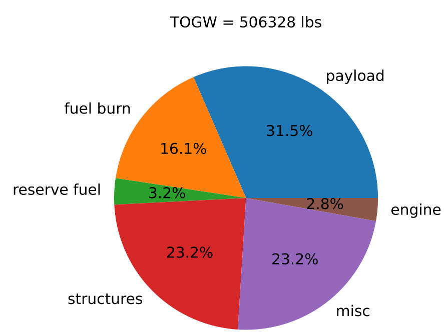

(f) P5: Tight stability design (SM = 5%)

Fig. 15 Weight breakdowns across each optimization problem. The weight breakdowns become more reasonable as the stability constraints are introduced

[17] Kao, J., Hwang, J., Martins, J. R., Gray, J. S., and Moore, K. T., "A modular adjoint approach to aircraft mission analysis and optimization," *56th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, 2015, p. 0136.

[18] Falck, R. D., Chin, J., Schnulo, S. L., Burt, J. M., and Gray, J. S., "Trajectory optimization of electric aircraft subject to subsystem thermal constraints," *18th AIAA/ISSMO multidisciplinary analysis and optimization conference*, 2017, p. 4002.

[19] Orndorff, N. C., Sarojini, D., Scotzniovsky, L., Gill, H., Lee, S., Cheng, Z., Zhao, S., Mi, C., and Hwang, J. T., "Air-taxi transition trajectory optimization with physics-based models," *AIAA SCITECH 2023 Forum*, 2023, p. 0324.

[20] Sarojini, D., Ruh, M. L., Joshy, A. J., Yan, J., Ivanov, A. K., Scotzniovsky, L., Fletcher, A. H., Orndorff, N. C., Sperry, M., Gandarillas, V. E., Asher, I., Chambers, J. T., Gill, H., Lee, S., Cheng, Z., Rodriguez, G., Zhao, S., Mi, C., Nascenzi, T., Cuatt, T., Winter, T. F., Guibert, A. T., Cronk, A., Kim, H. A., Meng, S., and Hwang, J. T., *Large-Scale Multidisciplinary Design Optimization of an eVTOL Aircraft using Comprehensive Analysis*, 2023. https://doi.org/10.2514/6.2023-0146, URL https://arc.aiaa.org/doi/abs/10.2514/6.2023-0146.

[21] Ruh, M. L., Fletcher, A., Sarojini, D., Sperry, M., Yan, J., Scotzniovsky, L., van Schie, S. P., Warner, M., Orndorff, N. C., Xiang, R., Joshy, A. J., Zhao, H., Krokowski, J., Gill, H., Lee, S., Cheng, Z., Cao, Z., Mi, C., Silva, C., Wolfe, L., Chen, J.-S., and Hwang, J. T., *Large-scale multidisciplinary design optimization of a NASA air taxi concept using a comprehensive physics-based system model*, 2024. https://doi.org/10.2514/6.2024-0771, URL https://arc.aiaa.org/doi/abs/10.2514/6.2024-0771.

[22] Gandarillas, V., Joshy, A. J., Sperry, M. Z., Ivanov, A. K., and Hwang, J. T., "A graph-based methodology for constructing computational models that automates adjoint-based sensitivity analysis," *Structural and Multidisciplinary Optimization*, Vol. 67, No. 5, 2024, p. 76.

[23] Qin, N., Vavalle, A., Le Moigne, A., Laban, M., Hackett, K., and Weinerfelt, P., "Aerodynamic considerations of blended wing body aircraft," *Progress in Aerospace Sciences*, Vol. 40, No. 6, 2004, pp. 321–343.

[24] Wakayama, S., and Kroo, I., "The challenge and promise of blended-wing-body optimization," *7th AIAA/USAF/NASA/ISSMO Symposium on Multidisciplinary Analysis and Optimization*, 1998, p. 4736.

[25] Larrimer, B. I., *Beyond tube-and-wing: the X-48 blended Wing-Body and NASA's quest to reshape future transport aircraft*, NASA, 2020.

[26] Wakayama, S., and Kroo, I., "Subsonic wing planform design using multidisciplinary optimization," *Journal of Aircraft*, Vol. 32, No. 4, 1995, pp. 746–753.

[27] Wakayama, S., "Multidisciplinary design optimization of the blended-wing-body," *7th AIAA/USAF/NASA/ISSMO Symposium on Multidisciplinary Analysis and Optimization*, 1998, p. 4938.

[28] Kuntawala, N., Hicken, J., and Zingg, D., "Preliminary aerodynamic shape optimization of a blended-wing-body aircraft configuration," *49th AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition*, 2011, p. 642.

[29] Lyu, Z., and Martins, J. R., "Aerodynamic design optimization studies of a blended-wing-body aircraft," *Journal of Aircraft*, Vol. 51, No. 5, 2014, pp. 1604–1617.

[30] Cui, Y., He, J., Li, Q., and Zhang, B., "Multi-Point Optimization Design of Blended Wing Body Based on Discrete Adjoint Method," *Aerospace*, Vol. 12, No. 5, 2025, p. 404.

[31] Reist, T. A., and Zingg, D. W., "Aerodynamic design of blended wing-body and lifting-fuselage aircraft," *34th AIAA Applied Aerodynamics Conference*, 2016, p. 3874.

[32] Ahuja, J., Perron, C., Rivera, R. B., Tai, J. C., and Mavris, D. N., "Comparison of blended wing body and tube-and-wing performance characteristics," *The Aeronautical Journal*, Vol. 129, No. 1337, 2025, pp. 1764–1790.

[33] McDonald, R. A., and Gloudemans, J. R., "Open vehicle sketch pad: An open source parametric geometry and analysis tool for conceptual aircraft design," *AIAA SciTech 2022 Forum*, 2022, p. 0004.

[34] Haimes, R., and Dannenhoffer, J., "The engineering sketch pad: A solid-modeling, feature-based, web-enabled system for building parametric geometry," *21st AIAA Computational Fluid Dynamics Conference*, 2013, p. 3073.

[35] Raymer, D., *Aircraft design: a conceptual approach*, American Institute of Aeronautics and Astronautics, Inc., 2012.

[36] Sederberg, T. W., and Parry, S. R., "Free-form deformation of solid geometric models," *Proceedings of the 13th Annual Conference on Computer Graphics and Interactive Techniques*, Association for Computing Machinery, New York, NY, USA, 1986, p. 151–160. https://doi.org/10.1145/15922.15903, URL https://doi.org/10.1145/15922.15903.

[37] Kenway, G., Kennedy, G., and Martins, J. R., "A CAD-free approach to high-fidelity aerostructural optimization," *13th AIAA/ISSMO multidisciplinary analysis optimization conference*, 2010, p. 9231. https://doi.org/10.2514/6.2010-9231, URL https://arc.aiaa.org/doi/abs/10.2514/6.2010-9231.

[38] Fletcher, A. H., and Hwang, J. T., "Implicit Nonlinear Geometry Parameterization for Multidisciplinary Design Optimization," *Proceedings of the ASME 2026 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference IDETC/CIE2026*, 2026. Accepted.

[39] Scotzniovsky, L., and Hwang, J. T., "A fast, memory-efficient panel method for large-scale multidisciplinary design optimization under uncertainty using graph-based modeling," *AIAA AVIATION FORUM AND ASCEND 2025*, 2025, p. 3021.

[40] Campello, E., Pimenta, P., and Wriggers, P., "A triangular finite shell element based on a fully nonlinear shell formulation," *Computational mechanics*, Vol. 31, No. 6, 2003, pp. 505–518.

[41] Xiang, R., van Schie, S. P., Scotzniovsky, L., Yan, J., Kamensky, D., and Hwang, J. T., "Automating adjoint sensitivity analysis for multidisciplinary models involving partial differential equations," *Structural and Multidisciplinary Optimization*, Vol. 67, No. 8, 2024, p. 146.

[42] Warner, M., Fletcher, A., van Schie, S. P., and Hwang, J. T., "A method for modular mdo model assembly via a solver independent field representation," *AIAA SciTech 2024 Forum*, 2024, p. 1802.

[43] Smith, H., and Norato, J., "Topology optimization of structures made of fiber-reinforced plates," *Structural and Multidisciplinary Optimization*, Vol. 65, No. 2, 2022, p. 58.

[44] Malone, B., and Mason, W. H., "Multidisciplinary optimization in aircraft design using analytic technology models," *Journal of Aircraft*, Vol. 32, No. 2, 1995, pp. 431–438.

[45] Anderson, J. D., and Bowden, M. L., *Introduction to flight*, Vol. 582, McGraw-Hill Higher Education New York, NY, USA, 2005.

[46] Joshy, A. J., and Hwang, J. T., "modopt: A modular development environment and library for optimization algorithms," *Advances in Engineering Software*, Vol. 213, 2026, p. 104084.

[47] Joshy, A. J., and Hwang, J. T., "PySLSQP: A transparent Python package for the SLSQP optimization algorithm modernized with utilities for visualization and post-processing," *arXiv preprint arXiv:2408.13420*, 2024.