DuctAPE: A steady-state, axisymmetric ducted fan analysis code designed for gradient-based optimization

Judd Mehr*, Nathan Lehnhof†, and Andrew Ning‡
Brigham Young University, 360 EB, Provo, UT, 84602

Ducted rotors are an intriguing option for the harvest and use of clean energy due to their potential benefits in both aerodynamic efficiency and reduced noise profiles compared to open rotor systems. Exploration of conceptual ducted rotor design in the context of novel and complex applications may be greatly aided by the use of gradient-based optimization. Despite ducted rotors having been relatively well studied for the last century, modern, gradient-based, optimization-ready, mid-fidelity analysis tools for ducted rotors are scarce. Building from existing ducted fan analysis methods and utilizing modern automatic differentiation tools, we have developed an optimization-ready analysis tool for low-Mach ducted rotors. In this work, we present our ducted rotor analysis code—*Duct*ed *A*xisymmetric *P*ropulsor *E*valuation (DuctAPE)—and showcase its suitability for gradient-based optimization. We show that DuctAPE matches the Ducted Fan Design Code within 1/2 percent, and matches experimental data within experimental uncertainties. We then show gradient-based optimizations for the aerodynamics of ducted rotors across a range of applications. We therefore show that DuctAPE is appropriate for use in both aerodynamic analysis and gradient-based optimization of ducted rotors. As such, DuctAPE is poised for immediate application in more comprehensive gradient-based multidisciplinary optimizations of electric ducted fans and similar applications.

Acronyms

AWE Airborne Wind Energy

DFDC Ducted Fan Design Code

DuctAPE Ducted Axisymmetric Propulsor Evaluation

DWT Ducted Wind Turbine

EDF Electric Ducted Fan

*Ph.D. Candidate, Department of Mechanical Engineering, juddmehr@byu.edu
†Undergraduate Research Assistant, Department of Mechanical Engineering
‡Professor, Department of Mechanical Engineering, aning@byu.edu, Associate Fellow AIAA.

I. Introduction

Growing public concern for the environmental effects of fossil fuels has induced rapidly increasing interest in the development of technologies for the harvest and consumption of clean energy. At the energy harvest end of the spectrum, wind energy technologies are continuing to evolve, despite many decades of optimization. On the other side of the spectrum, advanced air mobility (AAM) research has been focused on electric propulsion systems in recent years [1]. In both application categories, ducted rotors have been suggested as potentially beneficial both for aerodynamic improvements as well as noise reduction [2, 3]. As novel concepts are proposed for both wind energy and AAM applications, optimization plays an important role in design space exploration. Gradient-based optimization techniques are especially helpful when considering multiple disciplines and/or large numbers of design variables and constraints, due to the inherent scalability of such methods. It would therefore be desirable to have a low-cost ducted rotor analysis tool suitable for use in modern, gradient-based, multi-disciplinary optimization (MDO) and design for both electric ducted fan (EDF) and ducted wind turbine (DWT) applications.

Ducted fans have been designed and in use since the 1930s, starting with Luigi Stipa's "intubed" propeller aircraft seen in figure 1. Ludwig Kort followed shortly thereafter, patenting his nozzle enclosed marine propeller technology [4]. DWT concepts have also been proposed since the 1950's starting with a study by Lilley and Rainbird [5] and subsequent experimentation. Ducted rotor analysis methods can generally be divided into categories of direct or inverse methods; as well as 1-D, axisymmetric, or fully 3-D. Methods also range from low-fidelity to high-fidelity.

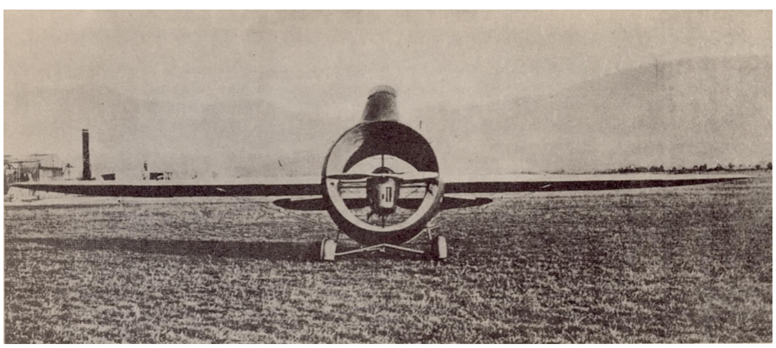

Fig. 1 Front view of the Stipa monoplane with venturi fuselage [6]—the first case of a rigorously studied ducted rotor.

Inverse approaches, in which the desired performance is prescribed and the required geometry is solved for, have been in use for many decades [7]. The success of inverse methods for design is highly dependent on (and limited to) the experience of the designer—unless optimization methods are used, such as Larocca who combined a throughflow inverse method with a multi-objective genetic algorithm to optimize a linear cascade as well as compressor stage [8]. Another interesting application of inverse design methods applied to ducted rotor optimization was done by Gomes et al.. They used an inverse method in a gradient-free differential evolutionary algorithm to optimize a self-rectifying impulse turbine blade for use in energy harvest in an oscillating water column [9]. Even if optimization methods are used, inverse methods are obviously not designed for direct analysis, limiting their flexibility.

Axisymmetric methods are quite common for both turbomachinery and DWT design and optimization using both direct and inverse methods. Persico and Rebay used TZFlow (a direct, axisymmetric, finite element/volume, throughflow solver for turbomachinery [10]) in conjunction with a meta-model enhanced genetic algorithm to optimize the discharge flow angles of the rotor and stator at three span-wise stations of a low-speed compressor stage [11]. They found that enhancing the genetic algorithm with a meta-model led to quick convergence with a Kriging meta-model outperforming a Feed-Forward Neural Net model for their optimization. Petrovic et al. performed a similar optimization using an inverse design, axisymmetric, throughflow model for a multi-stage axial turbine, adding the center body and casing geometries as design variables [12]. They used a constrained hybrid optimization algorithm that combines a genetic algorithm, the Nelder-Mead simplex method, simulated annealing, and the Davidon-Fletcher-Powell gradient search method. They found that the optimized design performed better than the initial configuration not only at the optimal point, but also across a wide range of off-design loading conditions.

We also find axisymmetric methods applied to DWT analysis and optimization. Oka et al. used an axisymmetric viscous meridional analysis based on the Navier-Stokes equations, representing blade loading as body forces. They also claim to be the first to apply genetic algorithms to the optimization of ducted wind turbines (both rotor and duct) [13]. Axisymmetric Reynolds-Averaged Navier-Stokes (RANS) models are also popular for the analysis and optimization of DWTs. Khamlaj and Rumpfkeil used axisymmetric RANS combined with blade element momentum theory for analysis and perform optimization using a multi-objective genetic algorithm [14]. They optimized chord and twist distributions, pitch angle, and duct shape and found that when optimizing them all together, they could decrease rotor thrust by roughly 10% while also increasing power coefficient by more than 10% relative to the baseline ducted turbine design. Bagheri-Sadeghi et al. used axisymmetric RANS combined with an actuator disk model for analysis and used pattern search methods for the optimization of duct geometry focusing on relative position of the rotor in the duct [15]. They found rotor performance to be relatively insensitive to the axial position of the rotor inside the duct, and found that the best performing operating conditions were very near separation.

Full three-dimensional RANS is also a common approach for both turbomachinery and DWT analysis and optimization, though it tends to be much more computationally expensive than axisymmetric methods. Wang et al. used an artificial neural net enhanced NSGA-II evolutionary algorithm, in combination with a 3-D RANS method, in a multi-objective optimization of the NASA Rotor 37 [16, 17]. With the neural net enhanced method, they saw a two order of magnitude decrease in required computational time; even so, computational time for their optimizations were on the order of days. Rahmatian et al. used a response surface method based on 79 3-D CFD analyses to create a surrogate model used with a genetic algorithm to optimize the geometry of a DWT flanged duct [18]. They found that the optimized duct achieves a higher power coefficient than both the nominal open and ducted rotor, and that the optimal max power coefficient is achieved at a higher tip speed ratio than the nominal cases' max power states.

As can be seen, a variety of optimization methods are also employed in the context of ducted rotor optimization, with gradient-free methods being most common. There are, however, several cases of gradient-based optimization methods being used to great success. Walther and Nadarajah used a discrete adjoint RANS solver with gradient-based optimization methods to minimize the entropy generation rate of a single stage transonic axial flow compressor [19]. With their gradient-based methods, they achieved an optimal solution in 18 iterations without the use of any surrogate models. In addition, they did not appear to encounter any issues with local optima, as many authors who have used gradient-free methods suggest is a common pitfall of gradient-based approaches. Papadimitriou and Giannakoglou developed an adjoint solver used in the optimization of a 3-D peripheral compressor blade cascade using a steepest descent method for the optimization. They found their gradient-based approach to be efficient, optimizing 28 design variables within 35 iterations [20]. Aranake and Duraisamy also used RANS combined with an actuator disk to analyze ducted wind turbines, but obtained gradients through finite difference methods [21].

Despite the long and rich history of ducted rotor design and optimization, there are few mid-fidelity ducted rotor analysis tools designed for use with modern, gradient-based optimization methods. It may also be noted that many turbomachinery methods are designed to model only a few stages, or perhaps only the internal geometries of such systems. It may be desirable for AAM (and is desirable for DWT) applications, however, to also have capabilities to model both the internal as well as external aerodynamics of a ducted rotor. One potentially viable approach would be to adjust blade element momentum theory (BEMT) for open rotors to the physics of a ducted rotor such as the work done by Stahlhut [22]. This approach would be relatively simple and suitable for the optimization of rotor blade geometry; unfortunately, some details concerning the duct and center body geometries are neglected when only using BEMT and therefore cannot be included in an optimization setting. Another methodology employed in some form or another by several works in the literature is to combine a vortex method with an actuator disk model [23, 24]. A good example of this approach is that taken by Bontempo and Manna, who couple the non-linear actuator disk method of Conway [26] with the panel method of Martensen [27]. While this method has been shown to work well for both ducted fans and DWTs [25, 28, 29], the actuator disk model for the rotor is somewhat limiting for direct optimization of the rotor blade geometry.

Yet another low-cost approach combines the previous concepts and uses a blade element based lifting-line model for the rotor, and a panel method for the duct and center body. Perhaps the best known code employing this methodology is the Ducted Fan Design Code (DFDC) [30] developed by Drela and Youngren, which has been compared with (and shown to match decently well to) experimental data and mid-fidelity methods [31]. DFDC sees active use in industry and is especially known for its inverse-design capabilities, which lends confidence to its applicability in a gradient-based optimization setting. DFDC has already been used in an optimization setting by Korondi et al. as part of building a multi-fidelity surrogate model for optimization of ducted propellers; though DFDC was not used directly in optimization in this case [32]. As it stands currently, DFDC is not automatically suitable for gradient-based optimization settings without modifications. In this work, we present a ducted rotor analysis code, based on the models implemented in DFDC, with differences and modifications making it suitable for use in modern gradient-based MDO applications to EDF, and a DWT cross-over application for Airborne Wind Energy (AWE)—for which ducted rotors have only (to the best of our knowledge at the time of publication) been considered for buoyant shell turbine applications [33].

The remainder of this work is organized as follows: In section II we establish the fundamental methodologies for our ducted rotor analysis, including details for modeling the duct and center body, the rotors, and wake for both propulsive and energy harvest operating conditions. We present verification of our code against DFDC and validation with respect to experimental data in section III. We then proceed to showcase the applicability of our analysis tool to gradient-based optimization in section IV by performing basic aerodynamic optimizations of an EDF and an AWE ducted rotor. We then conclude in section V as well as discuss expected and recommended future work.

II. Methods: Ducted Axisymmetric DuctedRotor Evaluation

In this section, we cover the individual components contained within the **Duct**ed **A**xisymmetric **P**ropulsor **E**valuation (DuctAPE) code we have developed for this work*. As an overview, DuctAPE combines an axisymmetric panel method to model the duct and center body, and a blade element lifting-line model combined with an axisymmetrically smeared wake model to model the rotors and wake. Since much of the underlying theory for DuctAPE is either well known, or overviewed in the DFDC theory document [30], we summarize the key concepts and forego overly detailed derivations.

A. Note on Nomenclature

For the content in this work, there are several components of electric ducted rotors to which we will refer frequently. Each of these components my have several names assigned in the literature depending on the discipline and even era in which they are used. For clarity, we define the terminology we will use throughout this work.

Starting from the center and working our way outwards, we call the body of revolution at the center of the electric ducted rotor the *center body*. We call the rotor a *rotor* to maintain generality as it could be either a propeller or wind turbine; the context will make it clear which applies. We call the end of the rotor blade connected to the center body the rotor hub; and we call the end of the blade the rotor tip. Finally, the annular airfoil rotated about the rotor and encasing the other components we call the *duct*. The inner surface (toward the axis of rotation) of the duct we call the duct casing, or simply *casing*; and the outer surface (away from the axis of rotation) of the duct we call the duct nacelle, or simply *nacelle*.

*`https://github.com/byuflowlab/DuctAPE.jl`

B. Duct and Center Body Model

The duct and center body are modeled in DuctAPE using an axisymmetric linear vortex panel method. As discussed by Lewis [34], an axisymmetric panel method can conveniently be developed nearly identically to a standard 2D (planar) panel method (see for example Katz and Plotkin [35]). The only required difference between the 2D and axisymmetric methods is the replacement of the 2D singularities with their axisymmetric counterparts. We therefore take the standard 2D panel method approach and develop a linear system of equations with the unknowns being the strengths associated with linearly distributed vortex panels (γ_j). In the axisymmetric case, the panels defining the boundary are more accurately described as axisymmetric bands, as shown in fig. 2a. Thinking of these bands as panels (as shown in fig. 2b), however, is a safe enough approach since (by the axisymmetric assumption) the unit ring vortex has no influence in the tangential direction ($\hat{\boldsymbol{e}}_\theta$ in fig. 2) [36].

(a) Axisymmetric band coordinate system.

(b) Panel representing an axisymmetric band; $\hat{\boldsymbol{e}}_\theta$ out of the page.

Fig. 2 Axisymmetric band and panel geometries as well as coordinate system definition.

We assemble our linear system of boundary integral equations in the typical manner, applying the Neumann boundary condition of no normal flow through the control points:

$$\sum_{j=1}^{N} \gamma_j \boldsymbol{G}_{ij} = -\left(\boldsymbol{V}_\infty + \boldsymbol{V}_{\text{ext}}\right) \cdot \hat{\boldsymbol{n}}_i \tag{1}$$

where $\boldsymbol{G}$ is comprised of the induced velocity normal to the ith control point due to the jth panel node, and γ_j is the strength of the linear vortex distribution associated with the jth panel node. $\boldsymbol{V}_\infty$ is the freestream velocity, and $\boldsymbol{V}_{\text{ext}}$ are the externally induced velocities, namely the velocities induced by the rotor and wake. To obtain the elements of $\boldsymbol{G}$, we numerically integrate the induced velocity per unit length across the panels using Gauss-Legendre quadrature. We choose to numerically integrate due to the difficulty in obtaining an analytic integral for the unit induced velocities per unit length which are defined to be [34, 37]:

$$v_z^\gamma = \frac{1}{2\pi r_o}\frac{1}{D_1}\left[\mathcal{K}(m) - \left(1 + \frac{2(\rho - 1)}{D_2}\right)\mathcal{E}(m)\right] \tag{2a}$$

$$v_r^\gamma = -\frac{1}{2\pi r_o}\frac{\xi/\rho}{D_1}\left[\mathcal{K}(m) - \left(1 + \frac{2\rho}{D_2}\right)\mathcal{E}(m)\right] \tag{2b}$$

where the superscript, γ, indicates a unit vortex induced velocity. In addition, $\mathcal{K}(m)$ and $\mathcal{E}(m)$ are complete elliptic integrals of the first and second kind, respectively, and

$$m = \left(\frac{4\rho}{\xi^2 + (\rho + 1)^2}\right) \tag{3}$$

$$\xi = \frac{z - z_o}{r_o} \tag{4}$$

$$\rho = \frac{r}{r_o} \tag{5}$$

$$D_1 = \left[\xi^2 + (\rho + 1)^2\right]^{1/2} \tag{6}$$

$$D_2 = \xi^2 + (\rho - 1)^2; \tag{7}$$

where (z, r) and (z_o, r_o) are the point being influenced and the point of influence, respectively.

We choose to use Gauss-Legendre quadrature rather than the adaptive Romberg integration methodology used in DFDC specifically for application to gradient-based optimization. By choosing a quadrature method with no adaptation, we can keep any integration error continuous. With an adaptive method, which has both an iteration limit and convergence criteria, the integration errors could potentially be noisy due to the lack of a consistent termination condition. In addition, Gauss-Legendre quadrature tends to be an efficient method, and we have found 8 quadrature points per panel to be sufficient. In the case of panel self-induction, we utilize a separation of singularity method to avoid integrating across a singularity when the point of influence is the panel's own control point.

We also apply the typical Kutta condition to the duct, to require flow to leave the trailing edge by setting the sum of the trailing edge node strengths to zero: $\gamma_1 + \gamma_N = 0$. By itself, this standard Kutta condition can lead to spurious spikes in surface velocity near the trailing edge. In order to increase the numerical robustness of the panel method, we employ an approach taken in DFDC and apply an additional, indirect Kutta condition by placing a control point just inside the interior of the duct trailing edge (see fig. 3) and define an associated unit normal oriented such that the unit normal is effectively in the direction of the bisection angle of the trailing edge panels. We apply the same boundary condition on this control point as the other control points in that we set the normal velocity to be zero through the control point.

A special consideration of axisymmetric, linear panel methods is that bodies of revolution will have at least one panel node on the axis of revolution (at the leading edge). Since the induced velocities of an elementary vortex increase

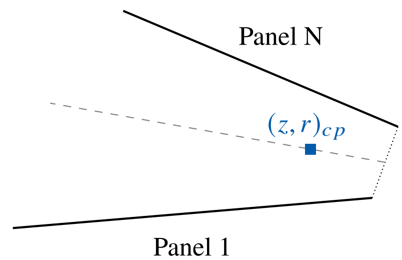

Fig. 3 For greater numerical robustness an additional control point (blue) is placed just inside the duct trailing edge (which may or may not be sharp).

with decreasing radius, a radius of zero leads to infinite velocities. In reality, the induced velocity from a zero radius vortex ring is zero. Therefore in our system, we need to prescribe the strengths of panel nodes on the axis of rotation to be zero strength. In order to achieve this, we take an approach similar to applying the Kutta condition: we simply add the equation $\gamma_{LE}^{cb} = 0$ to the system, where γ_{LE}^{cb} is the prescribed node strength for the center body leading edge. We could alternatively remove the unknown leading edge node strength from the system altogether, but our approach simplifies the bookkeeping of the code implementation. If the center body trailing edge is sharp, then we have an additional node on the axis of rotation and also need to prescribe its strength to zero.

C. Rotor Model

Rotors are modeled in DuctAPE using a blade element lifting-line approach. Specifically, we take the body, wake, and rotor induced velocities and calculate inflow angles and magnitudes at each blade element location and use airfoil lookup tables to determine the section lift and drag. From the section lift and drag, we determine the section circulation and approximate the section drag as discussed below.

1. Blade Element Circulation

To model the rotor blades axisymmetrically, we assume that they can reasonably be modeled as a lifting-line such that local blade element circulation can be expressed according to the Kutta-Joukowski theorem; and we define the blade element circulation as

$$\Gamma(r) = \frac{1}{2} W c c_\ell, \tag{8}$$

where the inflow magnitude, $W = |W_z + W_\theta|$, the blade element chord length, c, and the blade element lift coefficient, c_ℓ, are all functions of the radial position, r, along the rotor blade. The blade element lift coefficients are determined based on a lookup table taking in any combination of local inflow angle, β_1, and stagger angle, γ_{be}, or angle of attack, α, as seen in fig. 4; as well as local values for solidity, Reynolds number, and/or Mach number (the specific combination of inputs depending on the airfoil or cascade polar database used).

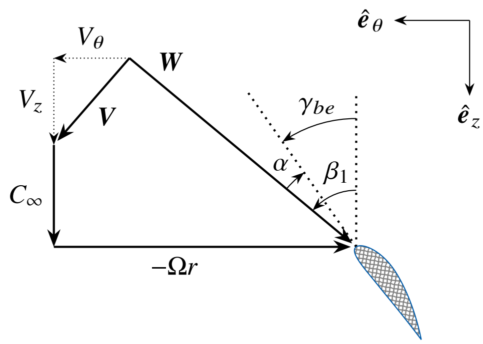

Fig. 4 Velocity decomposition with angles in the blade element frame. Where W is the blade element relative velocity, V is the induced velocity, Ωr is the rotor rotation rate, and C_∞ is the freestream velocity.

2. Rotor Profile Drag Approximation

We also assume that the rotor blade section profile drag can be approximated by the addition of source panels along the rotor blade. The inviscid approximation of the profile drag per unit length takes a similar form to the local circulation:

$$\Sigma = \frac{1}{2} W c c_d, \tag{9}$$

where c_d is the blade element drag coefficient obtained alongside the lift coefficient from a lookup table; and again each of the terms on the right hand side of eq. (9) are functions of the radial position along the blade. For an axisymmetric representation, we obtain the total source sheet strength per unit length for a given blade element by smearing the total source strength per unit span of all the rotor blades, B, around the non-dimensional circumference, 2π. Therefore the expression for the smeared rotor source strength per unit length along the blade is

$$\sigma = \frac{B}{4\pi} W c c_d. \tag{10}$$

3. Induced Axial, Radial, and Swirl Velocities

The induced axial and radial velocities (used in the rotor, wake, and body models) are calculated based on the unit induced velocities, $\hat{V}_z$ and $\hat{V}_r$, and the various panel strengths. In general terms, the axial and radial induced velocities are calculated as

$$V_z^A = \hat{V}_z^{AB} \zeta^B, \tag{11a}$$

$$V_r^A = \hat{V}_r^{AB} \zeta^B, \tag{11b}$$

where the superscript A indicates the object being influenced, the superscript B indicates the object doing the influencing, $\hat{V}_z^{AB}$ and $\hat{V}_r^{AB}$ are the axial and radial unit induced velocities from the nodes of object B to the control points of object A, and ζ indicates the node strengths associated with object B.

The swirl velocity induced by upstream rotor blades, V_θ, can be determined by applying Stokes' and Kelvin's theorems. If we define a control volume around a streamtube as shown in fig. 5, where the first curve is taken about all upstream rotors along a streamline, and the second curve is taken about the axis of rotation, only in the r-θ plane with radius such that the edge of the contour lies on the same streamline upon which the first curve lies (see the dotted line in fig. 5), we see by Kelvin's theorem (conservation of circulation), that the circulation due to the upstream rotors can be related to the tangential velocity downstream of the rotors through Stokes' theorem.

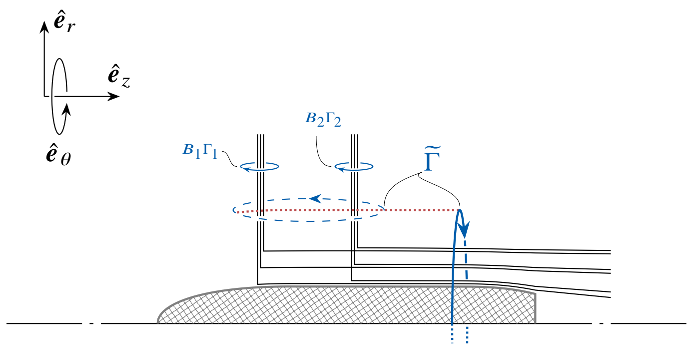

Fig. 5 Circulation is conserved between the dashed and solid contours, noting the red dotted line indicating the streamline on which the $\widetilde{\Gamma}$ contours align. The integral over the contour about the axis of rotation yields V_θ in terms of $\widetilde{\Gamma}$.

The expression for V_θ downstream from the rotor is then given by

$$V_\theta = \frac{\widetilde{\Gamma}}{2\pi r}, \tag{12}$$

where V_θ in our smeared, axisymmetric model is the circumferentially averaged swirl velocity induced by upstream rotors, and $\widetilde{\Gamma}$ is the net circulation contribution of all the blades of the upstream rotors:

$$\widetilde{\Gamma} = \sum_{i=1}^{N} B_i \Gamma_i. \tag{13}$$

For the self-induced case, the contour is placed at the rotor plane. This means that the rotor "sees" infinite trailing vortices from any upstream rotors, but only semi-infinite trailing vortices for itself. Thus the rotor experiences the full swirl induced by upstream rotors, but only half of its own swirl contribution:

$$(V_\theta)_{\text{self}} = \frac{1}{2\pi r}\left(\widetilde{\Gamma} + \frac{1}{2}B\Gamma\right), \tag{14}$$

where $B\Gamma$ here is the number of blades and circulation of the rotor itself.

D. Wake Model

The wake model in DuctAPE is an axisymmetrically smeared representation of the vortex filaments shed from the rotor blades. The axisymmetric vortex sheets are discretized into vortex panels like those used to model the duct and center body, with strengths dependent on the upstream rotor circulation and velocities induced on the wake. In turn the wake vortex panels induce velocities on the rotors and bodies, thereby coupling the wake to the rotors and bodies.

1. Wake Aerodynamics

For a given position on a blade producing a circulation change, $\Delta\Gamma$, by conservation of circulation, a helical vortex filament of strength $-\Delta\Gamma$ is shed into the flow. In order to represent 3D vortex filaments in our axisymmetric reference frames, we will also make the approximation that they can be smeared into equivalent axisymmetric vortex sheets in the meridional (m) and tangential (θ) directions. The smeared axisymmetric vortex sheets then have circulation to length ratios (densities) of γ_m and γ_θ in their respective directions. Because we are modeling the wake internal to the duct, we cannot guarantee a cylindrical wake, and therefore cannot simply model the wake with straight vortex cylinders. We will still use the concept of a wake cylinder, however to help us model discrete sections of the wake; so we continue with a description of how we smear a helical vortex filament into a cylindrical sheet.

We begin with a shed vortex sheet, the geometry of which we approximate by a left-handed helix with non-dimensional apparent pitch, h_B (see fig. 6a):

$$h_B = \frac{2\pi}{B}\frac{\mathrm{d}m}{-\mathrm{d}\theta}. \tag{15}$$

We assume that vortex filaments are shed parallel to the relative inflow velocity, $\boldsymbol{W}$. To dimensionalize the lengths for a given smeared cylindrical surface, we multiply by the cylinder radius, r, to obtain the dimensional length:

$$h_B r \approx \frac{2\pi r}{B}\left(\frac{W_m}{-W_\theta}\right). \tag{16}$$

In order to obtain a cylindrically smeared shed circulation density, γ_θ (as seen in fig. 6b), we take the shed vortex filament strength at a given radial station and divide by the dimensionalized apparent pitch. In addition, we need to apply an additional negative to ensure the vortices resulting from the left-handed helix are oriented correctly in our right-handed system. Thus

$$\gamma_\theta = -\frac{-\Delta\Gamma}{h_B r} = -\Delta\Gamma\frac{B}{2\pi r}\left(\frac{W_\theta}{W_m}\right). \tag{17}$$

Unfortunately, eq. (17) is only generally applicable if we assume that the Ωr component of W_θ is constant in the entire wake, which we do not. For our ducted case, in which the wake radius changes, we only know Ωr right at the

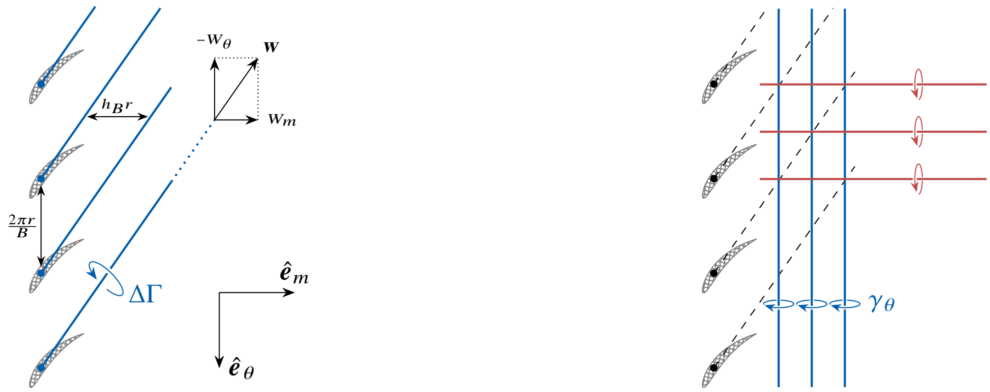

(a) Wake screw geometry. **(b) Axisymmetric smeared cylinder.**

Fig. 6 2D vortex sheets are generated from ratios of circulation to lengths between vortex sheets.

rotor lifting-line, but not generally in the remainder of the wake. We therefore want a more general expression for γ_θ based on requiring the wake to be force-free, or in other words, we demand static pressure continuity across the vortex sheets. The relatively lengthy derivation for a suitable expression is covered well in the DFDC theory document [30], so we will forego repeating it here for brevity. By assuming static pressure continuity across wake sheets; inviscid, low Mach conditions; and a calorically perfect gas, we can use eq. (17) along with various thermodynamic relations to arrive at the following expression for γ_θ that is applicable to our ducted wake case:

$$\gamma_\theta = -\frac{1}{2C_{m_{\text{avg}}}}\left(-\left(\frac{1}{2\pi r}\right)^2\left(\widetilde{\Gamma}_2^2 - \widetilde{\Gamma}_1^2\right) + 2\left(\widetilde{h}_2 - \widetilde{h}_1\right)\right), \tag{18}$$

where $C_{m_{\text{avg}}}$ is the average absolute meridional velocity on a wake panel node:

$$\boldsymbol{C}_m = C_z\hat{z} + C_r\hat{r}, \tag{19}$$

the 1 and 2 subscripts indicate blade elements on either side of the shed vortex sheet, and $\widetilde{h}$ is the accumulation of changes in enthalpy across upstream rotors:

$$\widetilde{h} = \sum_{i=1}^{N} \Delta h_{\text{disk}_m}; \tag{20}$$

where the jump relation Δh_{disk} is defined as

$$\Delta h_{\text{disk}} = \Omega\frac{B\Gamma}{2\pi}. \tag{21}$$

2. Precomputation of Wake Geometry

The approach DFDC takes to modeling the wake geometry is perhaps one of the most efficiency inducing aspects of the code, and we take the same approach in DuctAPE. We generate the axisymmetric wake lines as a "grid" defined by the solution of an elliptic partial differential system, using the solid bodies as the boundaries of the grid. By solving the appropriate governing partial differential equations, we can generate a grid that is approximately aligned with the actual streamlines for the system. Thompson et al. provide further insights into the benefits of this approach [38]. By defining the wake geometry to lie on an elliptic grid, we can discretize the axisymmetric wake lines into axisymmetric vortex panels and apply the circulation density (vortex strength distribution) from eq. (18) along the discretized wake panels. The derivation for the non-linear system solved for the wake grid is covered well in the DFDC theory, and we omit it here for brevity.

E. Solve Approach

1. State Initialization

To initialize the states, we first solve the panel method given a uniform freestream. We then run a fast blade element momentum theory (BEMT) solver, namely CCBlade [39] using the uniform freestream and body induced velocities. From the BEMT solution, we set the rotor source panel strengths and compute the blade element circulation. We then initialize the wake strengths from the freestream and induced velocities and blade element circulations.

2. Solver Method

The solve algorithm we use in this work is shown in algorithm 1. This method is similar to the approach taken in DFDC, but we have modified and reorganized most of the steps in order to allow for efficient automatic differentiation through the solve. The most notable difference in architecture is that the DFDC implementation updates states inside the residual before estimating other states in a vaguely Gauss-Sidel manner. We have moved all of the state updates outside of the residual enabling the use of implicit automatic differentiation methods. Doing so leads to slightly longer convergence times for individual analyses, but ends up being faster overall in a gradient-based optimization setting by avoiding having to pass derivatives through every iteration of the solve.

F. Viscous Drag Model

1. Duct

As an approximation of the viscous drag on the duct we assume a fully turbulent boundary layer and use Head's well-known entrainment method [40]. To solve the ODE functions of Head's method, we employ a straightforward second-order Runge-Kutta method. We determine initial conditions by starting with the momentum thickness value from the Schlichting empirical fit for a turbulent flat plate as well as the boundary layer shape factor for a turbulent flat

Algorithm 1 Solution Method

Initialize body, rotor, and wake strengths
while unconverged **and** iterator < iteration limit **do**
· Solve the linear system for the body vortex strengths ▷ using eq. (1).
· Calculate new estimates for the blade element circulation ▷ using eq. (8).
· Calculate new estimates for the wake vortex strengths ▷ using eq. (18).
· Calculate new estimates for the rotor source strengths ▷ using eq. (10).
· Calculate relaxation factors for each each state variable.
· Update states according to relaxation factors.
· Check for convergence.
end while
Post-process Solution

plate. We then obtain an estimated drag coefficient from the solved boundary layer using the Squire-Young formula [41], which relates the momentum thickness, shape factor, and edge velocity at the surface trailing edge (note the TE subscripts) to the drag coefficient:

$$c_d = \frac{2\delta_{2_{TE}}}{c}\left(\frac{U_{e_{TE}}}{U_\infty}\right)^{\frac{5+H_{12_{TE}}}{2}}. \tag{22}$$

To obtain an estimate of the total, dimensional drag force on the duct, we apply the definition of drag coefficient to get the dimensional drag per unit length:

$$D' = \frac{1}{2}\rho V_\infty^2 c\left(\hat{c}_{d_{\text{upper}}} + \hat{c}_{d_{\text{lower}}}\right); \tag{23}$$

where c is the chord length of the duct. We then integrate the drag per unit length about the circumference of the duct, using the duct exit diameter as the characteristic length.

2. Center Body

For the center body, we utilize a method similar to those used to approximate fuselage drag. We take the drag coefficient for the center body (C_D) to be

$$\begin{aligned} C_D &= \frac{D}{0.5\rho V_\infty^2 S_{\text{ref}}}, \\ &= C_f f_{\text{form}} \frac{S_{\text{wet}}}{S_{\text{ref}}}; \end{aligned} \tag{24}$$

where ρ_∞ is the freestream density, V_∞ is the freestream velocity, C_f is the flat-plate skin-friction coefficient, f_{form} is a form factor correction, S_{wet} is the wetted area of the center body, and S_{ref} is a reference area. Setting the two expressions equal, we solve for drag as

$$D = 0.5\rho_\infty V_\infty^2 C_f f_{\text{form}} S_{\text{wet}}. \tag{25}$$

We use the form factor expression from Shevell [42] based on fineness-ratio (l/d):

$$f_{\text{form}} = 1 + \frac{2.8}{(l/d)^{1.5}} + \frac{3.8}{(l/d)^3}. \tag{26}$$

We take C_f to be the skin friction coefficient from Schlichting for a flat plate of length equal to the length of the center body and the velocity at the body trailing edge.

III. Verification and Validation

In this section we present verification and validation of DuctAPE, showing outputs compared to DFDC, as well as characterizing some of the limitations of DuctAPE through a comparison with experimental data.

A. Verification Against DFDC

As we have established, the methodology behind DuctAPE is based heavily on DFDC. Therefore, we present a set of comparisons between DuctAPE and DFDC. We compare with an example available in the DFDC source code using a single ducted rotor across a range of operating conditions, specifically across a range of advance ratios including a hover condition.

The geometry used is shown in fig. 7. The various geometry details, as well as DFDC run files and DuctAPE analysis scripts are available in the companion repository[†] to this work. We note here that DuctAPE also differs from

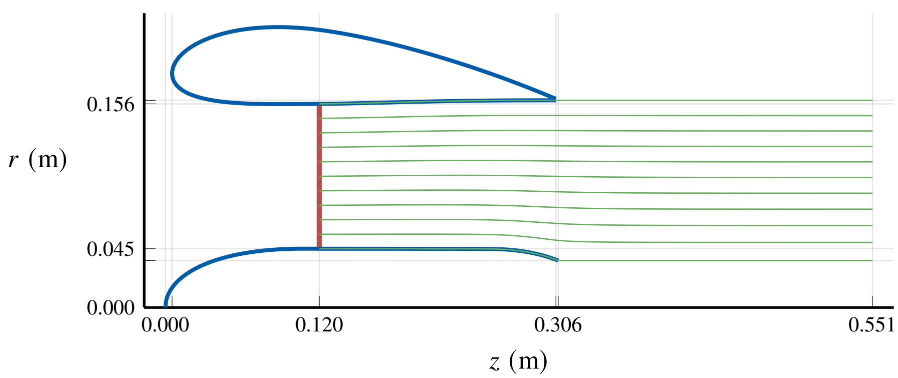

Fig. 7 Single rotor verification case geometry generated in DuctAPE. Duct and center body geometry in blue, rotor lifting line location in red, and approximate wake streamlines in green, where markers indicate panel egdes.

DFDC in the geometry re-paneling approach. The DuctAPE geometry re-paneling approach aligns the duct, center body, and wake panels aft of the rotor and distributes them linearly. We align the panels so that there is a consistent number of panels between discrete locations (such as rotor positions and body trailing edges) in the geometry, thereby avoiding discontinuities that would be incompatible with a gradient-based optimizer. For example, the number of center

[†]https://github.com/byuflowlab/ductape-2025-companion-repository

body and duct panels ahead of and behind the rotor need to stay constant if the rotor position is selected as a design variable in an optimization. Without the number of panels ahead of and behind the rotor staying constant, there would be discontinuities as the rotor passed over panels along the solid bodies.

Scanning table 1, we see that the differences between DFDC and DuctAPE are generally less than 1/2% for major output values for both a hover and a cruise case. Figure 8 shows comparisons of total thrust and power coefficients (fig. 8a) and total efficiency (fig. 8b), across the range of advance ratios, showing excellent matching across the entire range.

Table 1 Comparison of solver outputs for a cruise ($J = 1.0$) and hover ($J = 0.0$) case. Errors relative to DFDC.

Values at J=1.0	DFDC	DuctAPE	% Error
Rotor Thrust (N)	70.0	70.21	0.3
Body Thrust (N)	6.99	6.98	-0.1
Torque (N·m)	5.5	5.52	0.32
Rotor Efficiency	0.63	0.63	0.1
Total Efficiency	0.69	0.69	0.06

Values at J=0.0	DFDC	DuctAPE	% Error
Rotor Thrust (N)	91.8	91.83	0.03
Body Thrust (N)	106.45	107.02	0.53
Torque (N·m)	6.58	6.58	0.04

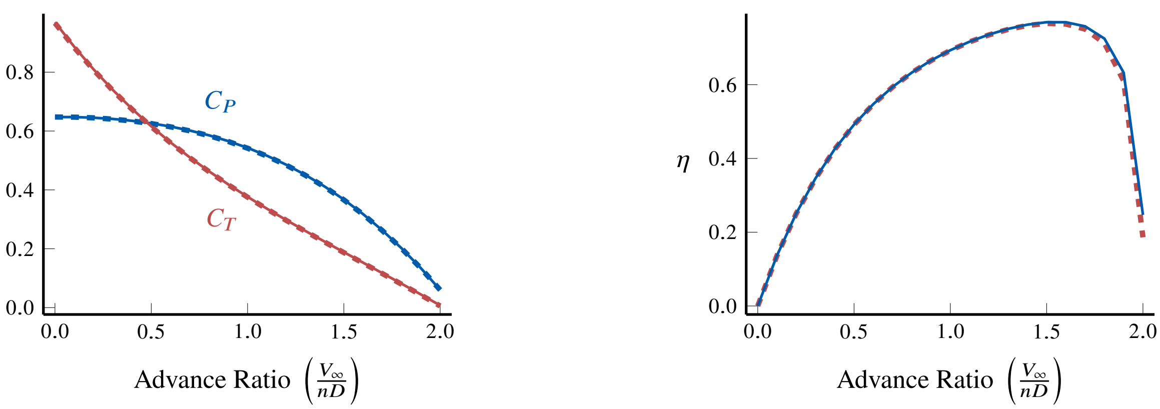

(a) Power and thrust comparison.

(b) Efficiency comparison.

Fig. 8 Comparison of power and thrust coefficients and efficiency for DFDC (dashed) and the DuctAPE implementation (solid) across a range of advance ratios.

B. Validation with Experimental Data

For validation, we compare DuctAPE outputs with data from a series of experiments performed by Hamilton Standard in the late 1960s [43]. The geometry for the experiments is shown in fig. 9. Coordinates for the duct and center body, as well as the location of the rotor are provided in the Hamilton Standard report [43]. Details on the axial location of the center body leading edge are missing, so we have somewhat arbitrarily chosen the geometry shown here based on photographs in the report. We determined, however, that the location of the center body leading edge has negligible impact on the results of the DuctAPE analysis for this case. We discretized the geometry and wake to allow both for numerical stability (avoiding too many panels), while also being sufficiently refined to keep relative changes in thrust and power coefficients below 1%. As with the verification case above, the specific geometry details and run scripts for this validation case—as well as digitized versions of the data tabulated in the Hamilton Standard report—are available in the companion repository to this work.

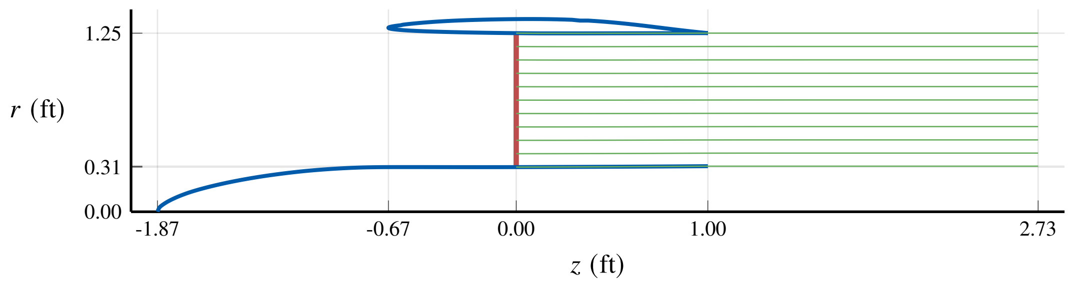

Fig. 9 High-speed validation case geometry generated in DuctAPE. Duct and center body geometry in blue, rotor lifting line location in red, and approximate wake streamlines in green, where markers indicate panel edges.

Rotor blade geometry information is also provided in the Hamilton Standard report for each of the cases, but some details are lacking in the definition of the blade section airfoil geometry. Based on the provided thickness and ideal lift coefficient distributions, we determined sets of NACA 16-series airfoils for which we ran XFOIL analyses. We then applied rotational (Du-Selig[44] and Eggers[45]) corrections to the resulting lift and drag polars which we also smoothed using B-Spline regressions to smooth out non-physical artifacts in the polars. Since the exact airfoil is unknown, we took minor liberties in applying additional pitch (within 2 degrees) to the rotors in an attempt to better match the rotor power.

We begin with rotor power and thrust coefficients compared in fig. 10a. We use faded markers for cases in which the rotor tip speed exceeded a critical Mach number of 0.7, and may therefore experience transonic effects not captured by XFOIL. We also include first-order uncertainty approximations based on measurement uncertainty and variable definitions provided in the Hamilton Standard report [43]. We see generally good matching of DuctAPE outputs compared with the experimental rotor data. Comparing with the average experimental values, we have an average absolute error of 3.7% for power coefficient and 1.5% error for thrust coefficient. Figure 10b shows comparisons of the total power and thrust coefficients (including duct forces). We note that general good matching remains. Specifically, the power coefficients are identical and the total thrust coefficient has an average absolute error of 2.1%. This indicates that our viscous drag model approximates the duct drag sufficiently well; which in this case balances almost completely with the duct-induced thrust. Thus very little difference is seen in the thrust between the rotor and total thrust.

These results are encouraging as we see DuctAPE able to capture the aerodynamics of ducted rotor systems well, even when stretching the underlying low Mach assumption of the methods. Though as noted, the low Mach assumption can only be stretched so far, as the accuracy is shown here to be dependent on the blade element polars provided to DuctAPE. If blade sections see transonic effects, it may still be possible for DuctAPE to model those cases accurately, provided that the polars are generated with a tool suitable for capturing transonic aerodynamics.

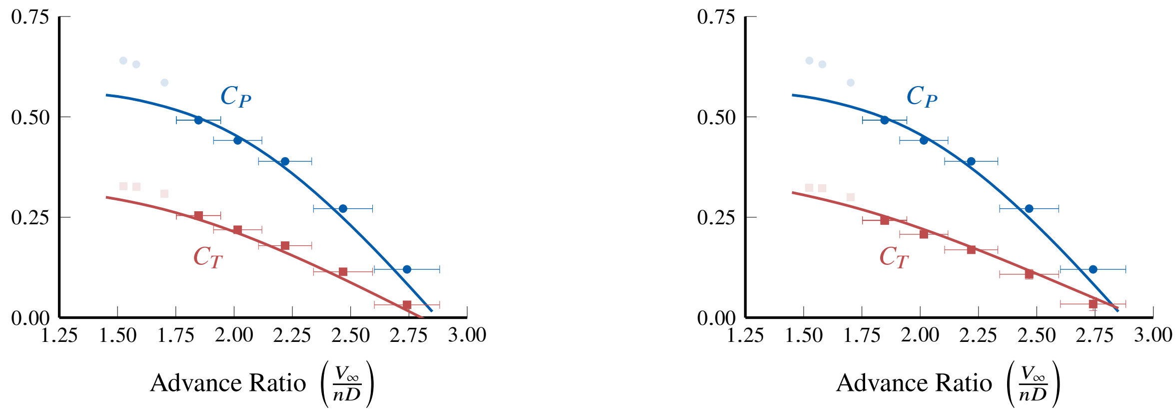

(a) Rotor power and thrust comparison. **(b) Total (rotor + duct) power and thrust comparison.**

Fig. 10 Comparison of rotor and rotor + duct power and thrust coefficients for DuctAPE (solid lines) and Hamilton Standard data (markers), where the cases with tip Mach number above 0.7 are faded.

The Hamilton-Standard experiments also included lower speed experiments with a bell-mouth inlet duct geometry. We found that DuctAPE struggled to match the experimental data in this case, notably due to large separation of the nacelle at higher speeds. This failure to model cases where viscous effects dominate highlights one of the limitations of DuctAPE—namely, the aerodynamic models are inviscid and therefore cannot model the coupled effects of separated flows, even when the integral boundary layer of section II.F is applied. Therefore DuctAPE analyses become increasingly inaccurate the more severe separation becomes. This point will become important later in section IV.

IV. Gradient-based Optimization of Electric Ducted Rotors

In order to demonstrate the suitability of DuctAPE for the gradient-based optimization of electric ducted rotors, we present three example optimizations: 1) An optimization of a propulsive configuration for cruise conditions only; 2) an optimization of a propulsive configuration for hover+cruise; and 3) an optimization of a dual-purpose (hover+generation) configuration for an AWE application. By including optimization problems at different ends of the use case spectrum, we show the current extent of DuctAPE's flexibility to be employed across a range of applications.

A. Optimization Problem Definitions

All of the optimizations in this work use the same general optimization problem, though with variations applicable to each case. The general problem definition is:

$$
\begin{aligned}
\text{Minimize} \quad & \text{expended energy;} \\
\text{with respect to:} \quad & \text{rotor geometry and rotation rate, and duct geometry;} \\
\text{subject to:} \quad & \text{thrust constraint(s), rotor tip Mach constraints, and various geometry constraints.}
\end{aligned}
\tag{27}
$$

where the objective function is nominally defined as

$$
\begin{aligned}
J &= \frac{E_{(V=0)}}{t_h + t_c} + \frac{1}{t_h + t_c}\frac{1}{V_{c_2} - V_{c_1}}\int_{V_{c_1}}^{V_{c_2}} E(V)\mathrm{d}V \\
&= \frac{P_{(V=0)}t_h}{t_h + t_c} + \frac{t_c}{t_h + t_c}\frac{1}{V_{c_2} - V_{c_1}}\int_{V_{c_1}}^{V_{c_2}} P(V)\mathrm{d}V \\
&= t_{h/t}P_{(V=0)} + \frac{1 - t_{h/t}}{V_{c_2} - V_{c_1}}\int_{V_{c_1}}^{V_{c_2}} P(V)\mathrm{d}V;
\end{aligned}
\tag{28}
$$

which is an expended energy minimization. We divide both the hover and cruise energy by the sum of time spent in hover and cruise ($t_h + t_c$) to frame the objective in terms of the ratio of time spent in hover: $t_{h/t} = t_h/(t_h + t_c)$ and we sweep this time ratio from no hover ($t_{h/t} = 0$) to $t_{h/t} = 0.1$ or $t_{h/t} = 0.3$ in the studies below. The range of cruise velocities (V_{c_1}, V_{c_2}) is 30-50m/s, and we perform a 3-point Gauss-Legendre integration to calculate the value of the integral for average cruise power. For simplicity, we assume that all cruise conditions take place between velocities of 30m/s and 50m/s, and that all operations take place either in hover or cruise with nothing modeled in between. Note that we scale the objective by a notional energy value, the computed constraint values by their respective constraint bounds, and design variables as needed in order to scale the optimization problem for better numerical performance.

The design variables for all of the optimizations include: the chord distribution along the rotor blade (c_r), the twist distribution along the rotor blade (θ_r), the rotor rotation rate (Ω) at each operating point, leading edge radial position, trailing edge radial positions for each operation phase (hover and cruise), and class shape transformation (CST) [46] parameters defining portions of the surface geometry of the duct. Note that the specific duct geometry parameters are detailed in section IV.B below.

We also set several constraints. We set required thrust value constraints for the various cases which are discussed in more detail below. We set limits on the blade tip Mach number to be 0.6 for all optimizations. We set the minimum allowable duct cross-sectional area to be 0.0045m^2 (7in.^{2}) to maintain sufficient duct internal volume to place any control or sensory electronics that might be necessary for operation. Though without knowing exactly what might need to go inside the duct, we selected this area somewhat arbitrarily to generally avoid impractically thin cross-sections. We also require the twist distribution to monotonically and smoothly decrease from blade hub to tip which helps the optimizer maintain reasonable rotor geometries. In addition, we set constraints on the duct nacelle parameters which ensure non-zero trailing edge thickness as well as avoiding surface crossover.

1. EDF Optimization Problem Variations

For both EDF optimizations we use the same airfoil polar parameterization employed in section III.A for the rotor blade elements. In general, the thrust (T) constraints in this work are given relative to the weight (mg) for which a single propulsor is responsible. In hover, each propulsor is directly responsible to carry a fraction of the vehicle weight (dividing up the vehicle weight between each propulsor). In cruise, we assume an average vehicle lift to drag ratio of 10. We therefore require an average thrust (determined using the same Gauss-Legendre quadrature procedure as used in the objective function) of 1/10th of the lift force. For the individual propulsor optimized in this work, this is means we constrain the thrust to be at least 10% of the weight each propulsor is responsible for carrying in hover conditions. We chose a lift to drag ratio of 10 based on findings in previous work indicating this to be a reasonable assumption for novel electric vehicles [47]. Mathematically, we define the thrust constraints for an individual propulsor to be:

$$T_{(V=0)} \geq mg; \tag{29}$$

$$\frac{1}{V_{c_2} - V_{c_1}} \int_{V_{c_1}}^{V_{c_2}} T(V) \mathrm{d}V \geq 0.1mg. \tag{30}$$

For the cruise only optimization, we only apply eq. (30) as there is no hover condition to constrain. We also sweep the required thrust constraint from roughly 1N to 30N. For the hover+cruise optimization, we keep the required thrust constant, setting the mass allotted to a single propulsor to lift to be $m = 15$kg and sweep the ratio of time in hover from $t_{h/t} = 0$ to $t_{h/t} = 0.1$.

2. AWE Optimization Problem Variations

In the AWE optimization, we perform an optimization of a dual-purpose ducted rotor to be used for an AWE kite. For context, a fixed-wing AWE kite using drag-based generation methods takes off vertically from the ground and slowly climbs to its operational altitude. Once at altitude, the kite transitions from vertical flight to an elliptical or figure-eight type flight path across the wind, which is analogous to the cruise condition in the previous optimization but is instead a power generation state. Therefore the rotor is required to act both as a propeller for hover and vertical climb as well as a wind turbine in cross-wind flight conditions. (Note we will call the cross-wind flight conditions generative conditions for the remainder of this work.) We selected the generative condition velocity range (30-50m/s) based on previous work [48] which took publicly available videos of the Makani M600 prototype [49–51] into account. Despite the difference in operation, the mathematics of the optimization problem in this AWE case are nearly identical to eq. (27) except that in the generative condition the rotor is ideally producing drag, or in other words, negative thrust. As such we set the thrust constraint in the generative condition to be greater than -20% of the hover thrust constraint. This thrust constraint is therefore a constraint on the maximum allowable drag (D) of the ducted rotor in a power generation configuration:

$$
\begin{aligned}
\frac{1}{V_{c_2} - V_{c_1}} \int_{V_{c_1}}^{V_{c_2}} T(V) \mathrm{d}V &\geq -0.2mg, \\
\frac{1}{V_{c_2} - V_{c_1}} \int_{V_{c_1}}^{V_{c_2}} D(V) \mathrm{d}V &\leq 0.2mg.
\end{aligned} \tag{31}
$$

The value of 20% is somewhat arbitrary, as the kinematics and dynamics of AWE kite flight are quite complex; but we use a drag constraint to approximate requirements for maintaining sufficient momentum across the flight path to stay aloft.

One additional difference in this optimization is that we take the airfoil polar parameterization from the EDF optimizations and slightly modify it to match a symmetric airfoil geometry which would be more suitable for a dual-purpose rotor than the airfoil polar used in the previous optimizations. As with the EDF case including hover, we keep the mass value from which the thrust constraints are determined constant at $m = 15$kg, and in this case we sweep the ratio of time in hover from $t_{h/t} = 0$ to $t_{h/t} = 0.3$.

B. Geometry and Geometric Design Variables

For the overall configuration, we use a rotor with a diameter of 0.254m (10in.). We set the duct chord length to be equal to the rotor diameter as a reasonable duct aspect ratio for these exploration studies. We also chose to center the rotor axially within the duct as a reasonable preliminary design. We leave the further development of DuctAPE allowing variable duct length and rotor position in gradient-based optimizations to future work. For simplicity, we set the center body chord length to be equal to the duct and align the leading and trailing edges. We also set a center body (and therefore rotor hub) radius of 25% of the rotor tip radius. Since it would have little effect beyond the design variables already present in the optimizations, we choose to maintain a constant center body for all optimizations.

Rotor geometry design variables include chord and twist values along the blade. Specifically, we use six design variables spaced equally along the blade for both chord and twist. We pre-select 13 positions along the blade for analysis and linearly interpolate between the design variables to obtain the values for analysis. An example of the chord and twist interpolation is shown in fig. 11. Note that we keep the airfoil sections constant for all optimizations, using variations (as described above) of the airfoil polar used in section III.A for the whole rotor blade in all optimizations.

We parameterize the duct geometry with a combination of B-Spline control points and CST [46] coefficients as shown in fig. 12. We first define two quadratic B-Splines that meet at the center of the duct (where the rotor is placed). The first of these quadratic B-Splines is the green spline with square control points shown in fig. 12, and the second is the red spline with diamond control points. The radial position of the first control point of the spline which is situated ahead of the rotor (the green one with square control points) is a design variable that controls inlet area (see the tick labeled I_{1_r} in fig. 12). Similarly, the radial position of the last control point of the spline which is situated behind the rotor (the red one with diamond control points) is a design variable that controls the exit area (see the tick labeled O_{3_r} in

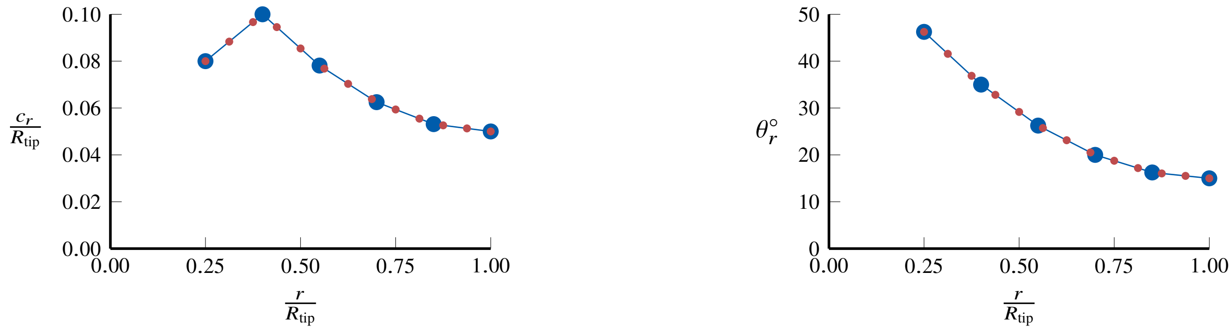

(a) Example chord (c) parameterization.

(b) Example twist (θ) parameterization.

Fig. 11 Visual explanation of the rotor geometry parameterization employed in this work. Blue circles indicate design variables, and smaller red circles indicate interpolated points for analysis.

fig. 12). In the optimizations that include both hover and cruise conditions, we allow the optimizer to select different exit areas for each, but require the inlet areas to be the same for all conditions. Rather than rotating an entire section of the duct geometry to approximate a variable exit, we allow the geometry to change smoothly in order to simplify the parameterization and avoid discontinuities incompatible with gradient-based optimization. To define the specific casing geometry, we want a round leading edge and a sharp trailing edge. To get a round leading edge, we define an additional cubic B-Spline (see the other red spline with circle control points in fig. 12) that has three of the same control points as the front quadratic B-Spline (see the overlapping green square and red circle control points in fig. 12) The cubic spline includes one additional control point (the red circle out in front in fig. 12) placed along the normal vector (represented by a dotted line) relative to the end of the quadratic spline. Using a cubic spline here gives us a round leading edge, rather than the corner seen between the green and blue lines in fig. 12. Placing the point along the normal vector enforces tangency between the casing and nacelle surfaces at the leading edge. The distance of the control point along the normal controls the leading edge radius of the casing and is fixed for the optimizations in this work. The quadratic spline aft of the rotor (red with diamond control points) already results in a sharp trailing edge, so we use that for the back half of the casing geometry We set the remaining control points of all of the B-Splines in such a way that the casing surfaces ahead of and behind the rotor smoothly connect at their junction (note all of the points along the dotted line at $r = 0$ in fig. 12). To define the nacelle geometry, we use a standard CST parameterization with five coefficients, the first of which is set and the remainder being design variables. We place the resulting points normal to the quadratic B-Splines, however, rather than normal to the axis. This allows the optimizer far more flexibility in positioning the inlet and outlet radii than a standard airfoil parameterization would. Note that fig. 12 is an exaggerated example of the duct geometry to clearly shown show the parameterization.

As noted in section III.B, there is a clear cutoff for when DuctAPE loses accuracy, specifically when boundary layer separation occurs on the duct. In addition, our drag model becomes difficult to use for cases such as hover, where there

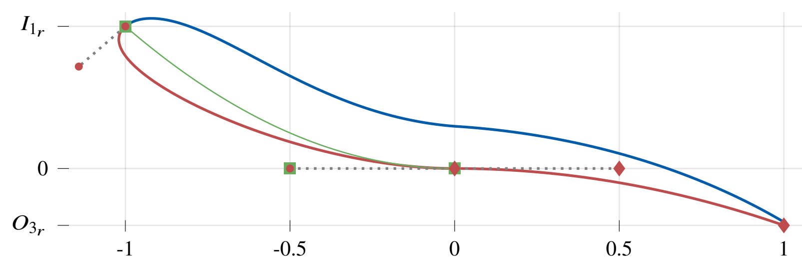

Fig. 12 Visual explanation of the duct geometry parameterization employed in this work: Nacelle and casing curves as well as front quadratic B-Spline. Gray dotted lines indicate control points placed to enforce continuity. Note that the nacelle geometry is positioned relative to the front quadratic B-spline for the front half, and relative to the casing outlet surface for the back half.

are complicated viscous coupling effects well outside the capabilities of DuctAPE to model. For this reason, we limit the design space available to the optimizer while attempting to maintain as much flexibility as possible. Doing so prevents the optimizer from venturing into regions of the design space dominated by viscous effects that DuctAPE is incapable of modeling. It is left to higher fidelity tools to more fully optimize the geometry to be nearer separation conditions if desired.

In preliminary exploration, we found there to be several potential design variables to which separation is very sensitive. These sensitivities led to numerical difficulties as the optimizer moved into regions of the design space where the underlying assumptions of DuctAPE's methodology began to break down. For the studies published in this work, we do not apply our drag model to the hover cases. Instead, we implement the optimizations using the following techniques in order to inherently avoid overly large adverse pressure gradients in the hover cases until very near the trailing edge, thereby helping the optimizer avoid numerical difficulties. The first technique we apply is to set the parameters with the most control of the duct leading edge radius to be constant. Specifically, we fix the first CST parameter and the relative position of the second cubic B-Spline control point. Setting these two parameters constant effectively keeps the leading edge radius constant and avoids lip separation. The next technique addresses the outlet geometry in two steps: first, we set a constraint on the maximum exit ratio for the hover case in a preliminary optimization to keep the exit area reasonable when no drag model is applied. We then take the result of the preliminary optimization and use a root finder with our drag model to determine the hover exit ratio that results in a boundary layer shape factor value of $H = 2.75$ at the trailing edge which is roughly 90% of the separation criteria ($H = 3.0$). We then reset the exit ratio upper bound constraint to the root finder solution value and re-optimized the remaining design variables. We consider these techniques to be relatively conservative as there are potentially lip suction effects as well as boundary layer re-energization aft of the rotor that could delay separation but are not captured by our drag model.

For the optimizations in this work, we use Sparse Nonlinear OPTimizer (SNOPT), a sequential quadratic programming algorithm for large-scale, constrained optimization [52]. For the most part we keep the various settings in SNOPT as their default values. One exception was setting the allows non-linear constraint violation magnitude during line searches to be 1.0 rather than the default 10.0 in an effort to help the optimizer avoid numerical difficulties when in nearly separated regions of the design space. To provide SNOPT with a Jacobian of the objective and constraint functions, we use ForwardDiff.jl, a Julia package for forward mode automatic differentiation [53].

C. Optimization Results

1. EDF Cruise-only Optimizations

Figure 13 shows the optimal duct geometries as the mass defining the required thrust is swept from 3kg to 27kg. As the thrust values increase (moving from the geometries in blue to those in red in fig. 13) the optimal geometries exhibit increasing inlet areas. This opening up of the inlet as the thrust maintains the stagnation point at the leading edge of the duct, thereby keeping the drag due to the inlet down by avoiding high adverse pressure gradients at the duct lip. Additionally, the optimal CST parameters reduce until the duct area constraint becomes active, leading to a slight reduction in drag due to the nacelle geometry for all the optimal designs.

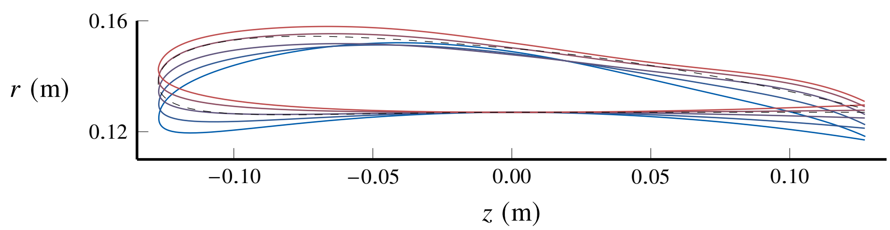

Fig. 13 EDF cruise-only initial (dashed black line) vs optimal geometries. Optimal values are plotted for several thrust constraints associated with a range of masses from m = 3kg in blue fading to m = 27kg in red.

Looking just at the inlet, one might expect that it is generally better to reduce duct drag to alleviate some of the burden on the rotor to produce the required thrust. If we shift our focus to the trailing edge, however, we also notice the exit area opening up with increasing thrust constraint. In this case, we actually see greater duct drag for the lower thrust cases. The difference is small relative to the rotor thrust contribution, but in fig. 14 we see a breakdown of the optimal rotor and duct contributions to thrust. Despite the additional duct drag, the optimizer has found it advantageous to reduce the exit area in order to minimize the energy expended.

As a sanity check, we took some of the lower required thrust optima and manually increased the outlet radius slightly while keeping the other variables constant. As we increased the outlet radius, we saw a decrease in duct drag, but also a reduction in rotor thrust, with the net effect being an overall reduction in total thrust to the point of constraint violation. Since the thrust constraint is active in all these optimizations, the optimal solution does indeed lie at a point with higher duct drag. Though it may be unintuitive to increase the duct drag for the cases requiring lower thrust, we see the benefit of optimizing the rotor and duct system together in order to find designs that are more optimal than would be found by

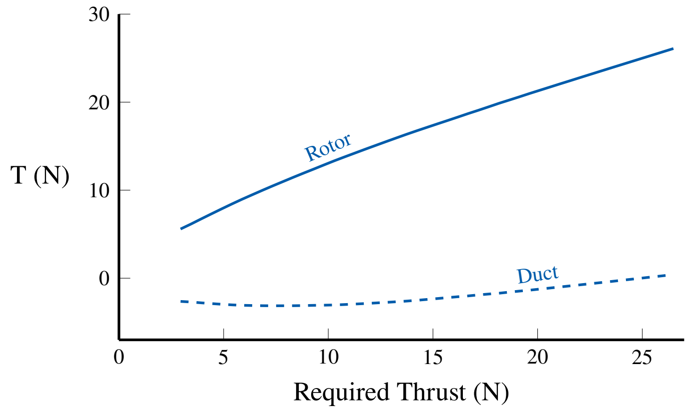

Fig. 14 Optimal rotor (solid line) and duct (dashed line) thrust components, where the total thrust constraint is active for all optimizations.

performing optimization separately. In this case, it is far more important to favor the rotor design over minimizing duct drag due to the exit area in order to minimize the objective while meeting the constraints. On the other hand, the rotor thrust is far less sensitive to other changes in duct geometry that reduce drag, such as inlet radius and nacelle shape. Being able to use gradient-based optimization techniques with DuctAPE allows us to fully leverage the various system sensitivities to arrive at a true optimum.

Figure 15 shows the optimal chord and twist distributions as well as rotation rates for the same cases plotted in fig. 13. In general, we see what we would expect for optimal solutions as the thrust constraint increases: As the thrust requirement increases, the rotor chord, twist, and rotation rate values likewise increase, though the increase in rotor geometry and rotation rate is not linear with change in thrust requirement. Looking back at fig. 14 we also see a slight non-linearity in the thrust components. With higher thrust requirements, the duct exit opens up, decreasing the duct drag and reducing the rotor's responsibility to meet the thrust requirement. As we see toward the right side of fig. 14, our studies approach conditions where the duct induces sufficient thrust to have net zero drag.

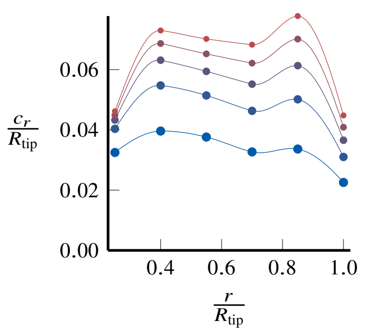

(a) EDF optimal chord distribution.

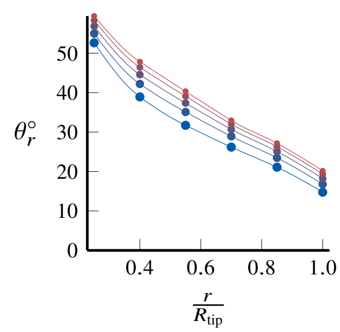

(b) EDF optimal twist distribution.

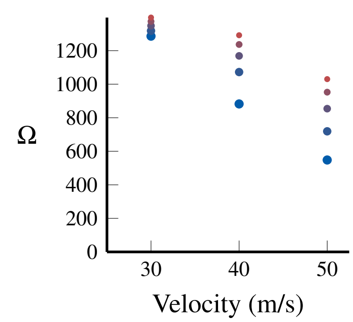

(c) EDF optimal rotation rate values.

Fig. 15 Optimal values for chord, twist, and rotation rate (from left to right) for the EDF cruise-only optimizations. Optimal values are plotted for cases ranging from m = 3kg in blue to m = 27kg in red.

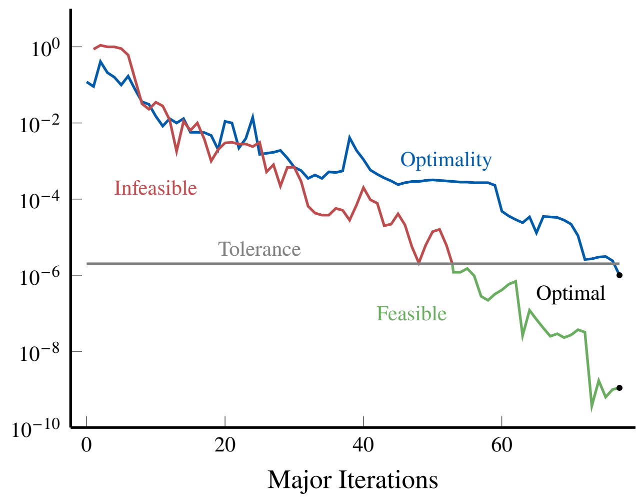

Fig. 16 Example of the progress of the optimality (in blue) and feasibility (in red and green) relative to the optimality and feasibility tolerance (in gray) as the optimization progresses. The black dots indicate the iteration at which both the optimality and feasibility conditions are met and the optimization terminates successfully.

Figure 16 shows an example (specifically the $m = 15$kg case) of the optimality and feasibility values reported by SNOPT at each major iteration as the optimization progresses. The optimal solution is shown to be well within the feasibility tolerance and also meets the optimality tolerance. Not only is the optimizer finding solutions improved from initial inputs, but also the solutions it finds are both feasible and mathematically optimal. With this, we conclude that DuctAPE is operating as desired and functioning in a gradient-based optimization setting.

2. EDF Hover+Cruise

We now take into account both cruise and hover conditions. In this case we keep the required thrust constant at roughly 15N and sweep the ratio of time in hover from $t_{h/t} = 0$ (the same problem as cruise-only in the previous section) to $t_{h/t} = 0.1$. Figure 17 shows the final duct geometries for the hover and cruise conditions at the ends of the hover ratio range. We see that the optima consist of the nearly the same duct geometry across the entire span of hover time ratios. This makes sense as we are not varying the required thrust as we did before. Note that the variations in nacelle geometry are small, so unlike fig. 13 we have plotted only the initial design along with the optima for the ends of the range of hover ratios: $t_{h/t} = 0.001$ and $t_{h/t} = 0.1$.

Comparing with fig. 13 we see the chosen optimum duct geometry in cruise is nearly identical to the cruise-only case. Then comparing the hover geometries in fig. 17a to the cruise geometries in fig. 17b we see some opening up of the exit area for the hover case. Typically, geometry for static conditions benefits from a bell-mouth inlet shape to enable a more open exit area by reducing the degradation of boundary layer health closer to the duct lip. As the inlet radii are required to be the same for both hover and cruise cases, it appears that it is more advantageous to prioritize the cruise condition inlet geometry to reduce drag in cruise than potentially gain more thrust in hover by moving toward a more

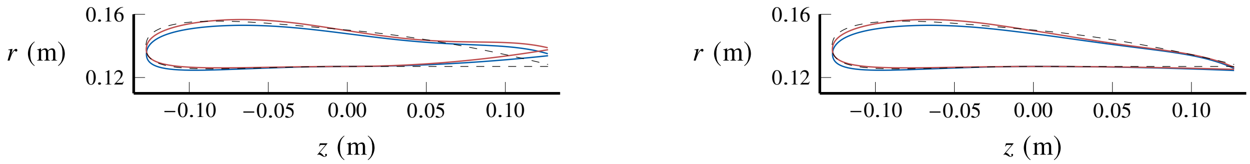

(a) Hover. **(b) Cruise.**

Fig. 17 EDF hover and cruise condition initial (dashed lines) vs optimal geometries. Optimal values are plotted for $t_{h/t} = 0.001$ in blue and $t_{h/t} = 0.1$ in red.

bell-mouth shaped inlet.

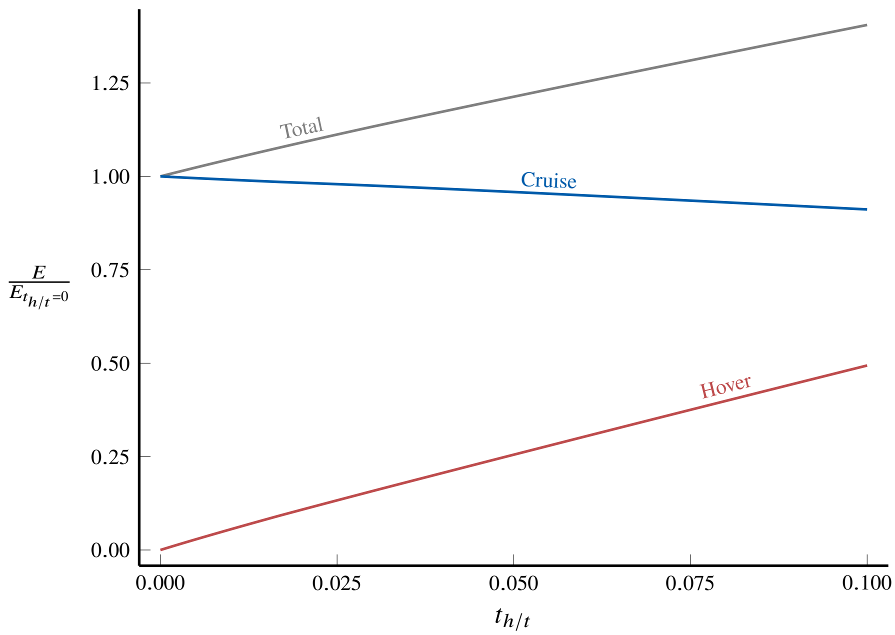

Fig. 18 Optimal total and cruise and hover components of energy relative to the $t_{h/t} = 0$ case.

Figure 18 shows the relationship between the optimal energy expenditure, non-dimensionalized by the zero-hover energy, relative to ratio of time spent in hover. We can clearly see why it is advantageous to favor the cruise geometry: the relative contribution of the hover state is much less than the cruise state, especially for lower ratios of time in hover. Predictably, as the hover time ratio increases so does the relative contribution of hover to energy expense. Looking back to fig. 17, we see a slight opening up of the inlet as the hover time ratio increases. Perhaps if hover dominated (which based on linearly extrapolating the cruise and hover curves, the point at which hover begins to dominate energy expenditure is $t_{h/t} = 0.172$), we would see a greater tendency toward a bell-mouth inlet geometry. We might expect the crossover point to happen sooner for larger thrust requirements; and for sufficiently large thrust requirements, the hover condition would likely begin to dominate the optimal solution. DuctAPE has been shown to work well in the current case; but if hover conditions were to dominate the energy expenditures, the limitations of DuctAPE due to the inviscid assumption discussed in section III.B might require higher fidelity tools to be employed instead.

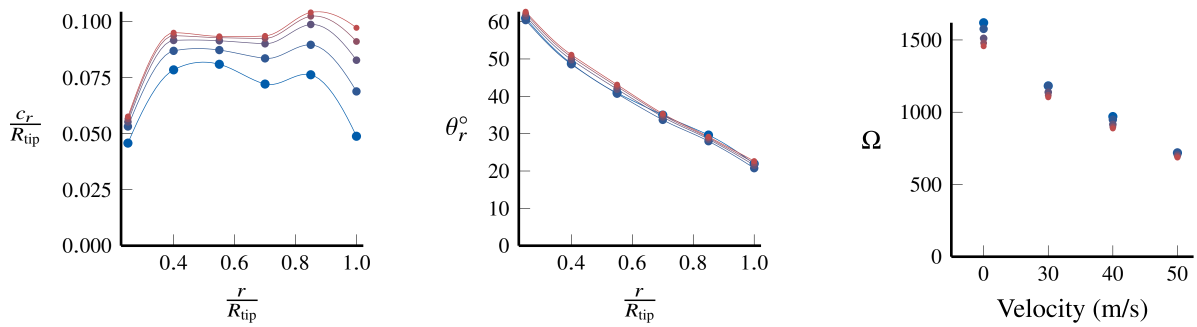

(a) EDF hover+cruise optimal chord distribution. **(b) EDF hover+cruise optimal twist distribution.** **(c) EDF hover+cruise optimal rotation rate values.**

Fig. 19 Optimal values for chord, twist, and rotation rate (from left to right) for the EDF hover+cruise optimizations. Initial points in gray, with triangles indicating design variable bounds. Optimal values are plotted from $t_{h/t} = 0.0$ in blue to $t_{h/t} = 0.1$ in red.

Analogous to fig. 15, fig. 19 shows the optimal chord and twist distributions as well as rotation rates for the same cases plotted in fig. 17. Similarly to fig. 17, fig. 19 shows very little change in the rotor geometry. Based on the lack of variation in the optimal rotor geometries, it isn't surprising to see a linear relationship with less time in hover requiring less total energy expenditure in fig. 18. The most significant changes are in the tip chord values, with slightly larger chords being chosen for longer times in hover. In general, the twist distribution and rotation rates stayed the same across the entire range of times in hover. The lack of variation across time in hover perhaps makes sense when considering the relatively low range of values considered. If we were to explore a wider range of relative hover time, we might expect to see the trends shown here continue: larger chord and twist rate values with more time in hover, and slightly decreased cruise rotation rates with increased time in hover.

3. AWE Optimizations

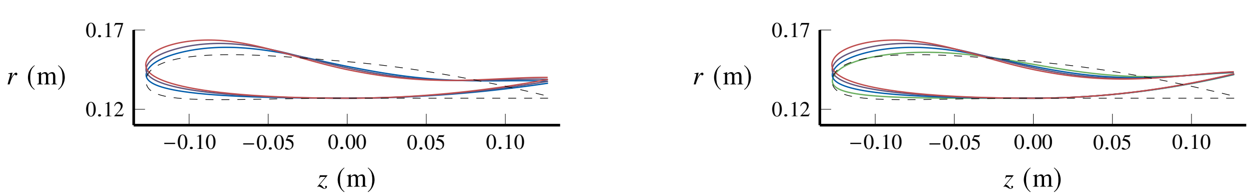

(a) Hover. **(b) Generative.**

Fig. 20 AWE hover and generative condition initial (gray) vs optimal geometries.Optimal values are plotted for $t_{h/t} = 0.0$ in green and from $t_{h/t} = 0.1$ in blue to $t_{h/t} = 0.3$ in red.

Figure 20 shows the initial and final duct geometries for the hover and generative flight conditions for the AWE case. There is some interesting contrast between the EDF and AWE cases: We compare figs. 17 and 20, though we note that the EDF case geometry does not vary from the initial designs much, and we use and plot the same initial designs for the AWE case. Therefore many of the trends we discuss now are easily seen in fig. 20 alone.

Compared to the EDF case, the optimal inlet geometry favors a more open design with greater variation in the inlet area and less variation in the hover exit area (though the larger variation in inlet can likely be attributed to the larger range of hover time ratios considered). The trends in inlet area are the same, however, with smaller inlet areas being associated with less time in hover, the zero-hover case in green being the smallest. The more open exit geometry of the AWE generative condition actually reduces the value of the shape factor on the upper side trailing edge (which directly affects the drag as calculated in eq. (22)). Therefore there is more allowance for a more open inlet in the AWE generative condition than in the EDF cruise condition where a more closed exit is required to minimize energy expenditure. In the AWE case, the duct cross-sectional area constraint is the same as used in the EDF case, but the optimum designs shift the area allocation forward more than in the EDF case. Likely there is interplay between the more open inlet and exit areas and their effect on the nacelle drag allowing the optimizer more flexibility than in the EDF cruise case where a more uniform area distribution was required to drive the duct drag down. Another interesting contrast with the EDF case is that in the AWE case, the exit areas in hover are actually slightly smaller than those for generative conditions. It appears that the inlet area for the AWE case is still driven by the "cruise" condition, though in this case the hover design also benefits. For ducted AWE rotor design then, we could potentially get away with a single, non-variable duct geometry without reducing overall energy capture by too much.

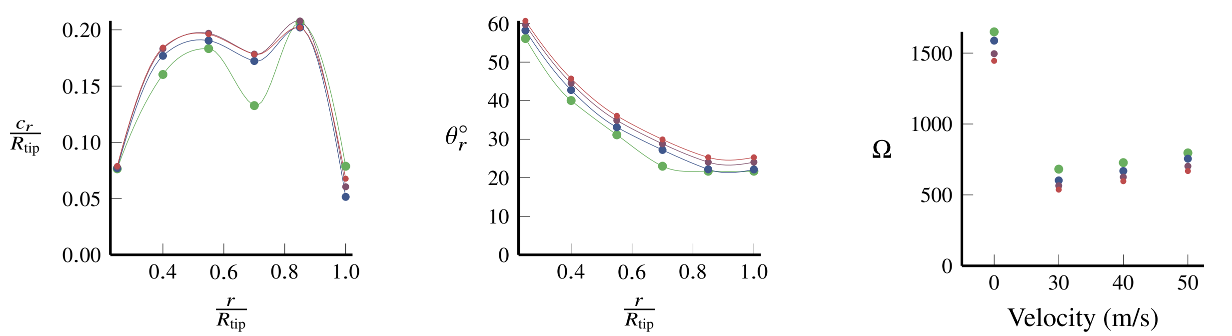

(a) AWE optimal chord distribution. **(b) AWE optimal twist distribution.** **(c) AWE optimal rotation rate values.**

Fig. 21 Optimal values for chord, twist, and rotation rate (from left to right) for the AWE optimizations. Initial points in gray, with triangles indicating design variable bounds. Optimal values are plotted for $t_{h/t}$ = 0.0 in green and from $t_{h/t}$ = 0.1 in blue to $t_{h/t}$ = 0.3 in red.

Figure 21 shows the rotor geometry and rotation rates for the AWE case optima across the range of hover time ratios explored. For the zero-hover case, there is a discernible difference in the chord and twist distributions which are highlighted in green in fig. 21. We see that as the hover ratio increases, the geometries quickly arrive at the nearly unchanging geometries for the remainder of the hover time ratios. This is in contrast to the EDF case, where the $t_{h/t} = 0$ condition is much nearer the non-zero hover ratio cases. This contrast should be expected since in the AWE case, the hover and generative operating points are performing opposing tasks, whereas in the EDF case, both hover and cruise are propulsive operations.

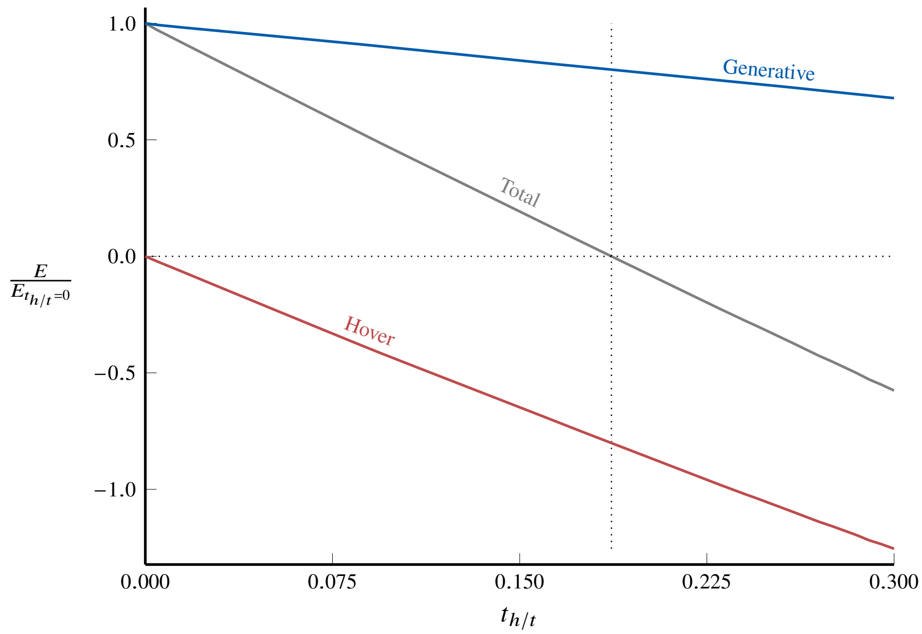

Fig. 22 Optimal total and generative and hover components of energy captured/expended.

Figure 22 shows the relative components of energy generation/expenditure. Note that positive numbers in fig. 22 indicate energy capture and negative values indicate energy expenditure. Looking at fig. 22 we see that the range of hover time ratios explored captures the cutoff point of the feasibility of the optimized design. At $t_{h/t} \approx 0.18$ the optimized designs are no longer able to achieve a net gain in energy, despite the generative conditions still capturing energy, the energy expenditure in hover quickly overcomes any time spent in generative conditions. Of course minimizing hover time relative to the total mission time is best, but optimization studies such as this one could be helpful in determining concept feasibility and economic viability.

Figure 23 shows thrust contributions for the EDF and AWE cases. Remember that both the EDF and AWE optimizations were all run with a constant hover thrust requirement for a 15kg load, which leads to nearly flat curves in fig. 23, rather than the non-linear curves seen back in the EDF cruise only optimizations (fig. 14). Comparing figs. 23a and 23b, we again see some similarities and differences between the EDF and AWE cases. As a similarity, the thrust components in hover are very similarly split. An interesting difference shows up in the cruise/generative case. Although both inducing drag, the EDF duct has significantly less of a drag contribution than does the AWE duct. This difference isn't entirely explained by a difference in thrust constraint either. The EDF rotor thrust is roughly 6 times larger than the duct drag, whereas the AWE duct thrust accounts for nearly half of the total allowed drag. This difference in results once again demonstrates the benefit of optimizing the duct and rotor system together rather than separately, as in both cases, hover is determined primarily from the thrust constraint. These results also further highlight the utility of DuctAPE in a range of gradient-based optimization settings, as the EDF and AWE cases both successfully optimize for their respective operating conditions.

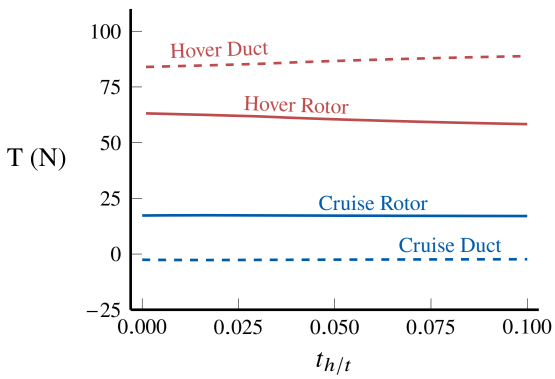

(a) Optimal thrust components for EDF optimizations.

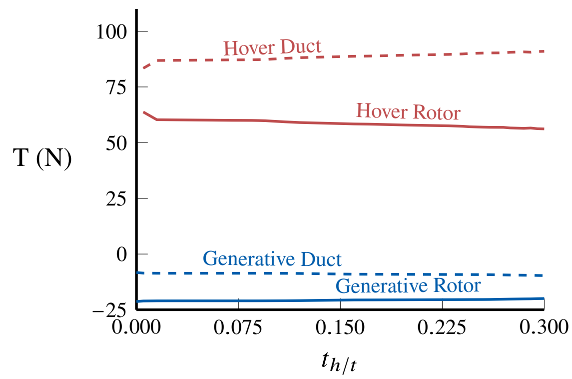

(b) Optimal thrust components for AWE optimizations.

Fig. 23 Comparison of EDF and AWE relative thrust contributions.

V. Conclusions and Future Work

In conclusion, we have developed a ducted fan analysis tool—DuctAPE—based on the underlying theory used for DFDC, including an axisymmetric panel method, blade element method, and smeared wake vortex method. We have developed DuctAPE with a variety of differences from the DFDC implementation allowing for smooth, automatic differentiation through the code to ensure it is suitable for application in gradient-based optimization settings. We have shown that despite these differences in implementation approach, DuctAPE matches DFDC within 0.5% for major output values.

We have further enhanced DuctAPE's capabilities by adding a basic integral boundary layer drag model. With the addition of this drag model, we have validated DuctAPE against experimental data, showing that DuctAPE performs well for cases without significant boundary layer separation, and showing that the drag model is sufficient for such cases.

We have shown that DuctAPE is suitable for application in gradient based optimization settings by performing aerodynamic optimizations for EDF and AWE dual-purpose rotors. These optimizations span a range of operating conditions, including both hover and cruise conditions, as well as a variety of design variables encompassing both rotor and duct geometry. Despite being limited to cases where the inviscid and incompressible assumptions of DuctAPE are generally valid, DuctAPE performs quite well in gradient-based optimization. On the other hand, for cases where coupled viscous effects are important (such as hover) care should be taken in the formulation of the optimization problem. If detailed exploration of performance trades taking boundary layer separation into account is required, then higher-fidelity tools with fully coupled viscous models should be considered.

Starting with an EDF cruise only set of optimizations, we saw that across a range of required thrust, it was generally more beneficial to prioritize rotor performance while reducing duct drag in ways to which the rotor performance was not as sensitive. Additionally considering hover along side cruise for the EDF case, we saw that prioritizing cruise performance while simply meeting hover constraints was optimal. Though with greater time spent in hover, there was some deviation toward hover-beneficial duct geometry. Finally, we explored an AWE dual-purpose rotor optimization and saw interesting contrast with the EDF case when not only propulsion, but power generation operating conditions were considered. Generally, hover performance was consistent between the two cases; but the AWE optima resulted in a significantly larger contribution of duct drag relative to rotor thrust in generative conditions.

Given the successful application to gradient-based optimization, we anticipate DuctAPE being useful in multi-disciplinary settings and will be exploring that possibility in the immediate future. Additional future work may include further improving the accuracy of DuctAPE through additional development of the viscous drag model, which would be especially improved if more complete inviscid-viscous coupling was implemented. Finally, despite sharing some of the same underlying theory of DFDC, DuctAPE lacks some of the auxiliary/additional features of DFDC that may be desirable for future work, including inverse design capabilities as well as wake shape update features.

VI. Acknowledgments

This paper is based upon work supported by NASA under award No. 80NSSC23M0218. The authors would also like to acknowledge the assistance of Ayden Bennett[‡] in uncertainty quantification of the validation data [43] used in section III.B.

References

[1] Garrow, L. A., German, B. J., and Leonard, C. E., "Urban air mobility: A comprehensive review and comparative analysis with autonomous and electric ground transportation for informing future research," *Transportation Research Part C: Emerging Technologies*, Vol. 132, 2021, p. 103377. https://doi.org/https://doi.org/10.1016/j.trc.2021.103377, URL https://www.sciencedirect.com/science/article/pii/S0968090X21003788.

[2] Casagrande Hirono, F., Torija, A. J., Grimshaw, S. D., Cousins, D., Farman, J., and Taylor, J. V., "Aerodynamic and aeroacoustic design of electric ducted fans," *Aerospace Science and Technology*, Vol. 153, 2024, p. 109411. https://doi.org/https://doi.org/10.1016/j.ast.2024.109411, URL https://www.sciencedirect.com/science/article/pii/S127096382400542X.

[3] Hosseini, S. E., and Saboohi, Z., "Ducted wind turbines: A review and assessment of different design models," *Wind*

[‡]Undergraduate Research Assistant, Department of Mechanical Engineering, Brigham Young University

Engineering, Vol. 0, No. 0, 0, p. 0309524X241282090. https://doi.org/10.1177/0309524X241282090, URL https://doi.org/10.1177/0309524X241282090.

[4] Kort, L., "Combined Device of a Ship's Propeller Enclosed by a Nozzle," , 1936. URL https://patents.google.com/patent/US2030375A/en.

[5] Lilley, G. M., and Rainbird, W. J., "A preliminary report on the design and performance of ducted windmills," Report 102, The College of Aeronautics: Cranfield, April 1956.

[6] Stipa, L., "Stipa Monoplane with Venturi Fuselage," Technical Memorandum 753, NACA, September 1934.

[7] Dyne, G., "An Experimental Verification of a Design Method for Ducted Propellers," Tech. Rep. 63, Swedish State Ship Building Experimental Tank, May 1968.

[8] Larocca, F., "Multiple Objective Optimization and Inverse Design of Axial Turbomachinery Blades," *Journal of Propulsion and Power*, Vol. 24, No. 5, 2008, pp. 1093–1099. https://doi.org/10.2514/1.33894.

[9] Gomes, R., Henriques, J., Gato, L., and Falcão, A., "Multi-point aerodynamic optimization of the rotor blade sections of an axial-flow impulse air turbine for wave energy conversion," *Energy*, Vol. 45, No. 1, 2012, pp. 570–580. https://doi.org/10.1016/j.energy.2012.07.042.

[10] Persico, G., and Rebay, S., "A penalty formulation for the throughflow modeling of turbomachinery," *Computers & Fluids*, Vol. 60, 2012, pp. 86–98. https://doi.org/10.1016/j.compfluid.2012.03.001.

[11] Pasquale, D., Persico, G., and Rebay, S., "Optimization of Turbomachinery Flow Surfaces Applying a CFD-Based Throughflow Method," *Journal of Turbomachinery*, Vol. 136, No. 3, 2014, p. 031013. https://doi.org/10.1115/1.4024694.

[12] Petrovic, M. V., Dulikravich, G. S., and Martin, T. J., "Optimization of multistage turbines using a through-flow code," *Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy*, Vol. 215, No. 5, 2001, pp. 559–569. https://doi.org/10.1243/0957650011538802.

[13] Oka, N., Furukawa, M., Yamada, K., Kawamitsu, K., Kido, K., and Oka, A., "Simultaneous Optimization of Rotor Blade and Wind-Lens for Aerodynamic Design of Wind-Lens Turbine," *Volume 3B: Oil and Gas Applications; Organic Rankine Cycle Power Systems; Supercritical CO2 Power Cycles; Wind Energy*, American Society of Mechanical Engineers, Düsseldorf, Germany, 2014, p. V03BT46A016. https://doi.org/10.1115/GT2014-25770.

[14] Khamlaj, T. A., and Rumpfkeil, M. P., "Analysis and optimization of ducted wind turbines," *Energy*, Vol. 162, 2018, pp. 1234–1252. https://doi.org/10.1016/j.energy.2018.08.106.

[15] Bagheri-Sadeghi, N., Helenbrook, B. T., and Visser, K. D., "Ducted Wind Turbine Optimization and Sensitivity to Rotor Position," , Nov. 2017. https://doi.org/10.5194/wes-2017-55.

[16] Wang, X., Hirsch, C., Kang, S., and Lacor, C., "Multi-objective optimization of turbomachinery using improved NSGA-II and approximation model," *Computer Methods in Applied Mechanics and Engineering*, Vol. 200, No. 9-12, 2011, pp. 883–895. https://doi.org/10.1016/j.cma.2010.11.014.

[17] Reid, L., and Moore, R. D., "Performance of Single-Stage Axial-Flow Transonic Compressor With Rotor an dStator Aspect Ratios of 1.19 and 1.23, Respectively, and With Design Pressure Ratio of 1.82," Technical Paper 1338, NASA, November 1978.

[18] Rahmatian, M. A., Nazarian Shahrbabaki, A., and Moeini, S. P., "Single-objective optimization design of convergent-divergent ducts of ducted wind turbine using RSM and GA, to increase power coefficient of a small-scale horizontal axis wind turbine," *Energy*, Vol. 269, 2023, p. 126822. https://doi.org/10.1016/j.energy.2023.126822.

[19] Walther, B., and Nadarajah, S., "Optimum Shape Design for Multirow Turbomachinery Configurations Using a Discrete Adjoint Approach and an Efficient Radial Basis Function Deformation Scheme for Complex Multiblock Grids," *Journal of Turbomachinery*, Vol. 137, No. 8, 2015, p. 081006. https://doi.org/10.1115/1.4029550.

[20] Papadimitriou, D. I., and Giannakoglou, K. C., "Compressor Blade Optimization Using a Continuous Adjoint Formulation," *Volume 6: Turbomachinery, Parts A and B*, ASMEDC, Barcelona, Spain, 2006, pp. 1309–1317. https://doi.org/10.1115/GT2006-90466.

[21] Aranake, A., and Duraisamy, K., "Aerodynamic optimization of ducted wind turbines: Aerodynamic optimization of ducted wind turbines," *Wind Energy*, Vol. 20, No. 5, 2017, pp. 877–889. https://doi.org/10.1002/we.2068.

[22] Stahlhut, C., "Aerodynamic Performance Modeling of Ducted Fans and ducted Propellers for Preliminary Design," *American Helicopter Society 70th Annual Forum*, Vol. 1, American Helicopter Society International, Inc., 2014, pp. 810–822.

[23] "Three-Dimensional Theory of Ducted Propellers," Tech. Rep. 602, D.E. Ordway, M.M. Sluyter, B.O.U. Sonnerup, August 1960.

[24] de Campos, J. F., "On the Calculation of Ducted Propeller Performance in Axisymmetric Flows," Ph.D. thesis, Technische Universiteit Delft, June 1983.

[25] Bontempo, R., and Manna, M., "Solution of the flow over a non-uniform heavily loaded ducted actuator disk," *Journal of Fluid Mechanics*, Vol. 728, 2013, pp. 163–195. https://doi.org/10.1017/jfm.2013.257.

[26] Conway, J. T., "Exact actuator disk solutions for non-uniform heavy loading and slipstream contraction," *Journal of Fluid Mechanics*, Vol. 365, 1998, pp. 235–267. https://doi.org/10.1017/S0022112098001372.

[27] Martensen, E., "Die Berechnung der Druckverteilung an dicken Gitterprofilen mit Hilfe von Fredholmschen Integralgleichungen zweiter Art," *Archive for Rational Mechanics and Analysis*, Vol. 3, 1959, pp. 235–270.

[28] Bontempo, R., Cardone, M., Manna, M., and Vorraro, G., "Ducted Propeller Flow Analysis by Means of a Generalized Actuator Disk Model," *Energy Procedia*, Vol. 45, 2014, pp. 1107–1115. https://doi.org/10.1016/j.egypro.2014.01.116.

[29] Bontempo, R., and Manna, M., "Performance analysis of open and ducted wind turbines," *Applied Energy*, Vol. 136, 2014, pp. 405–416. https://doi.org/10.1016/j.apenergy.2014.09.036.

[30] Drela, M., Youngren, H., and Sanders, S., "Ducted Fan Design Code (DFDC)," , 2005. URL https://web.mit.edu/drela/Public/web/dfdc/.

[31] Sargent, P. B., and Anemaat, W. A., "Benchmarking a Robust Panel Code for Ducted Fan VTOL Aircraft Design," *2018 Applied Aerodynamics Conference*, American Institute of Aeronautics and Astronautics, Atlanta, Georgia, 2018. https://doi.org/10.2514/6.2018-4212.

[32] Korondi, P. Z., Parussini, L., Marchi, M., and Poloni, C., "Reliability-based Design Optimization of a Ducted Propeller through Multi-fidelity Learning," *Proceedings of the 3rd International Conference on Uncertainty Quantification in Computational Sciences and Engineering (UNCECOMP 2019)*, Institute of Structural Analysis and Antiseismic Research School of Civil Engineering National Technical University of Athens (NTUA) Greece, Crete, Greece, 2019, pp. 600–619. https://doi.org/10.7712/120219.6363.18806.

[33] Saleem, A., and Kim, M.-H., "Aerodynamic performance optimization of an airfoil-based airborne wind turbine using genetic algorithm," *Energy*, Vol. 203, 2020, p. 117841. https://doi.org/10.1016/j.energy.2020.117841.

[34] Lewis, R. I., *Vortex Element Methods for Fluid Dynamic Analysis of Engineering Systems*, No. 1 in Cambridge engine technology series, Cambridge University Press, Cambridge ; New York, 1991.

[35] Katz, J., and Plotkin, A., *Low speed aerodynamics*, 2nd ed., Cambridge University Press Cambridge, UK, Cambridge, UK, 2001.

[36] Branlard, E., *Wind Turbine Aerodynamics and Vorticity-Based Methods*, Research Topics in Wind Energy, Vol. 7, Springer, 2017. https://doi.org/10.1007/978-3-319-55164-7.

[37] Ryall, D. L., and Collins, I. F., "Design and Test of a Series of Annular Aerofoils," Reports and Memoranda 3492, Ministry of Technology Aeronautical Research Council, March 1967.

[38] Thompson, J. F., Thames, F. C., and Mastin, C., "Automatic numerical generation of body-fitted curvilinear coordinate system for field containing any number of arbitrary two-dimensional bodies," *Journal of Computational Physics*, Vol. 15, No. 3, 1974, pp. 299–319. https://doi.org/10.1016/0021-9991(74)90114-4.

[39] Ning, A., "Using blade element momentum methods with gradient-based design optimization," *Structural and Multidisciplinary Optimization*, Vol. 64, No. 2, 2021, pp. 991–1014. https://doi.org/10.1007/s00158-021-02883-6.

[40] Head, M. R., "Entrainment in the Turbulent Boundary Layer," Reports and Memoranda 3152, Aeronautical Research Council, September 1958.

[41] Squire, H. B., and Young, A. D., "The Calculation of the Profile Drag of Aerofoils," Reports and Memoranda 1838, Aeronautical Research Committee, November 1937.

[42] Shevell, R. S., *Fundamentals of flight / Richard S. Shevell.*, Prentice-Hall, Englewood Cliffs, N.J, 1983.

[43] Rohrbach, C., Fortmann, W., and Croxall, A., "Variable Geometry ducted Propeller Test Program: Final Report, Volume II," Report AD835090, United Aircraft Research Laboratories, May 1968.

[44] Du, Z., and Selig, M., *A 3-D stall-delay model for horizontal axis wind turbine performance prediction*, ???? https://doi.org/10.2514/6.1998-21, URL https://arc.aiaa.org/doi/abs/10.2514/6.1998-21.

[45] Eggers, A., Chaney, K., and Digumarthi, R., "An Assessment of Approximate Modeling of Aerodynamic Loads on the UAE Rotor," 2003. https://doi.org/10.1115/WIND2003-868.

[46] Kulfan, B., "A Universal Parametric Geometry Representation Method - "CST"," *45th AIAA Aerospace Sciences Meeting and Exhibit*, American Institute of Aeronautics and Astronautics, Reno, Nevada, 2007. https://doi.org/10.2514/6.2007-62, URL https://arc.aiaa.org/doi/10.2514/6.2007-62.

[47] Critchfield, T., and Ning, A., "Mission-Focused Multidisciplinary Design Optimization of Tilt-Rotor eVTOL Propulsion System," *Journal of Aircraft (In Review)*, 2025.

[48] Mehr, J., Alvarez, E. J., and Ning, A., "Interactional Aerodynamics Analysis of a Multi-Rotor Energy Kite," *Wind Energy*, Vol. 28, No. 1, 2025, p. e2957. https://doi.org/10.1002/we.2957.

[49] Echeverri, P., Fricke, T., Homsy, G., and Tucker, N., "The Energy Kite: Selected Results From the Design, Development and Testing of Makani's Airborne Wind Turbines Part I," Tech. rep., Makani Technologies LLC, 2020.

[50] Echeverri, P., Fricke, T., Homsy, G., and Tucker, N., "The Energy Kite: Selected Results from the Design, Development and Testing of Makani's Airborne Wind Turbines Part II, Technical Artifacts," Tech. rep., Makani Technologies LLC, 2020.

[51] Echeverri, P., Fricke, T., Homsy, G., and Tucker, N., "The Energy Kite: Selected Results from the Design, Development and Testing of Makani's Airborne Wind Turbines Part III, Technical Artifacts," Tech. rep., Makani Technologies LLC, 2020.

[52] Gill, P. E., Murray, W., and Saunders, M. A., "SNOPT: An SQP Algorithm for Large-Scale Constrained Optimization," *SIAM Journal on Optimization*, Vol. 12, No. 4, 2002, pp. 979–1006. https://doi.org/10.1137/S1052623499350013.

[53] Revels, J., Lubin, M., and Papamarkou, T., "Forward-Mode Automatic Differentiation in Julia," *arXiv:1607.07892 [cs.MS]*, 2016. URL https://arxiv.org/abs/1607.07892.