



DUCTED HEAT EXCHANGER THERMAL AND AERODYNAMIC SHAPE OPTIMIZATION

Eytan J. Adler, Andrew H. R. Lamkin, and Joaquim R. R. A. Martins

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

Abstract

Interest in aircraft electrification and hydrogen fuel cells is driving demand for efficient waste heat management systems. Ultimately, most of the heat must be rejected to the freestream air. Ducted heat exchangers, also called ducted radiators, are the most common and effective way to do this. Engineers manually design ducted heat exchangers by adjusting the duct's shape and heat exchanger's configuration to reduce drag and transfer sufficient heat. This manual approach misses potential performance improvements because engineers cannot simultaneously consider all of the complex interactions between the detailed duct shape, heat exchanger design, and operating conditions. To find these potential gains, we apply gradient-based optimization to a three-dimensional ducted heat exchanger computational fluid dynamics model. The optimizer determines the duct shape, heat exchanger size, heat exchanger channel geometry, and coolant flow rate that minimize the ducted heat exchanger's power requirements while rejecting enough heat. Gradient-based optimization enables the use of nearly 100 shape design variables, creating a large design space and allowing fine tuning of the optimized design. When applied to an arbitrary, poorly-performing baseline, our method produces a nuanced and sophisticated ducted heat exchanger design with five times less cruise drag. Employing this method in the design of electric and fuel cell aircraft thermal management could uncover performance not achievable with manual design practices.

Keywords: Multidisciplinary design optimization, heat transfer, aerodynamics

Nomenclature

A	area	η_{pump}	coolant pump efficiency
D	drag	ρ	density
L	lift	$(\cdot)_{\text{HX}}$	heat exchanger model quantity
$\dot{m}$	mass flow rate	$(\cdot)_{\text{CFD}}$	CFD model quantity
P	power	$(\cdot)_{\text{coolant}}$	coolant quantity
p	pressure	$(\cdot)_{\text{ecrz}}$	economy cruise
q_c	compressible dynamic pressure	$(\cdot)_{\text{hscrz}}$	high-speed cruise
r	residual	$(\cdot)_{\text{in}}$	quantity at heat exchanger face
T	temperature	$(\cdot)_{\text{out}}$	quantity at heat exchanger exit
v	true airspeed	$(\cdot)_{\text{rto}}$	rolling takeoff
W	weight	$(\cdot)_t$	stagnation quantity
η_{motor}	electric motor and inverter efficiency	$(\cdot)_{\infty}$	freestream quantity
η_{prop}	propulsive efficiency		

1. Introduction

Hydrogen fuel cell and electric aircraft have a fundamental thermal management challenge: fuel cells and electrical components accumulate waste heat within their structure [1, 2]. This is unlike conventional jet engines, which exhaust most of their waste heat. Further complicating the problem is the fact that fuel cell and electrical component waste heat is at low temperatures, usually less than 100 °C. This provides a smaller temperature difference between the hot components and ambient air compared to conventional combustion engines, making efficient thermal management challenging.

Heat exchangers are the primary way that aircraft transfer heat from liquid coolant to the surrounding air. If placed directly in the freestream, a heat exchanger's heat transfer is proportional to aircraft speed, while the power needed to overcome the heat exchanger's drag scales with the cube of the speed [3]. Embedding the heat exchanger in a duct (Figure 1) can overcome this unfavorable scaling [4, Sec. IX.B]. A well-designed duct inlet slows the air, reducing the pressure loss across the heat exchanger. The nozzle downstream of the heat exchanger can convert some of the waste heat into thrust [5]. The thrust does not fully offset the drag of the duct and heat exchanger at low freestream Mach numbers [6], but it may at higher speeds with more available ram pressure.




Figure 1 – The duct slows the air, reducing the pressure drop across heat exchanger.

Designers need a heat exchanger model, a duct model, and a method to couple the two. Modern heat exchanger design relies on empirical analysis methods [7–9]. These methods assume a specific type of heat exchanger and provide scaling correlations to determine the heat exchanger's performance for modified heat exchanger geometry or flow conditions. Computational fluid dynamics (CFD) analyses of heat exchangers enable detailed design of fin and channel shapes [10–13]. However, this methodology is usually limited to simulating a single channel. Meshing and analyzing an entire heat exchanger requires orders of magnitude more computation, making it impractical to incorporate this approach directly into a complete ducted heat exchanger model.

Duct designers commonly use quasi-one-dimensional compressible flow models and CFD. Quasi-one-dimensional methods can estimate forces, flow conditions, and necessary cross-sectional areas. The popular Numerical Propulsion System Simulation (NPSS) [14] and pyCycle [15] codes use quasi-one-dimensional methods to model ducts in engine design. Jasa et al. [16] use a quasi-one-dimensional third stream ducted heat exchanger to augment waste heat rejection for a supersonic turbofan configuration. Brelje et al. [17] implement a quasi-one-dimensional duct for aircraft thermal management system design. Quasi-one-dimensional models can use empirical correction factors to approximate the performance effects of different shapes [18], but the models do not directly capture the detailed shape. Shaping inlets to maximize pressure recovery and minimize drag is a challenging task and is critical for achieving the best performance [19]. To avoid separation on the exterior or interior of the inlet at a range of design conditions, a designer must determine an appropriate angle and radius for the inlet lip. CFD allows designers to directly model the shape and investigate these design tradeoffs. CFD analysis of ducts is commonplace, and there are many examples of CFD-based duct shape optimization ranging from simple flow through configurations to complex turbofans [20–26].

Simulating a ducted heat exchanger requires the coupling of heat exchanger and duct models. For

conceptual design, empirical heat exchanger models have been integrated into quasi-one-dimensional duct models [17, 27, 28]. For more detailed shape design, Drela [29] couples a simplified heat exchanger model to MSES, a viscous/inviscid two-dimensional airfoil code [30]. He prescribes a heat rejection rate, which dictates the air’s enthalpy change across heat exchanger mesh cells. The model computes the pressure drop across the heat exchanger as a specified fraction of dynamic pressure at the heat exchanger’s face. His model assumes that the velocity at the heat exchanger’s exit is normal to the heat exchanger’s face. Drela also introduces terms to the integral boundary layer model in the heat exchanger cells to capture the heat exchanger’s boundary layer thinning effect. Another way to model a heat exchanger within a duct CFD analysis is with porous media [31]. This approach introduces viscous and inertial source terms to the momentum governing equation to represent the pressure drop and resistance across the heat exchanger. These terms can be direction dependent to simulate the heat exchanger’s flow-straightening effect. A separate source term in the energy governing equation models the heat transfer to the fluid. Patrao et al. [32] use a porous media approach to simulate a heat exchanger embedded in a duct for intercooling and recuperation within a hydrogen turbofan core. The momentum and heat transfer source terms are derived from empirical pressure drop and heat transfer correlations. A genetic optimizer tunes duct shape parameters of the two-dimensional, axisymmetric model. The optimization problem does not include heat exchanger sizing and channel design variables; nor does it include heat transfer as an objective or constraint. Hendricks et al. [33] also perform CFD-based coupled duct and heat exchanger analysis, but not optimization.

CFD-based aerodynamic shape and thermal optimization of a ducted heat exchanger can automatically determine the duct shape and heat exchanger properties that optimally balance heat transfer, overall drag, inlet pressure recovery, coolant pumping power, and other important metrics. Shape optimization with CFD and the adjoint method enables aerodynamic design involving hundreds or more shape design variables and has been applied to many airfoil and wing design problems [34], but not to coupled duct and heat exchanger design. Part of the challenge is making it computationally tractable. Most applications of numerical optimization to engineering problems use gradient-free optimizers to simplify the problem setup. Gradient-free optimizers can only handle up to about 10 design variables before they become prohibitively inefficient [35, Sec. 7.1]. If efficient gradient computation is available, gradient-based optimizers can optimize thousands of design variables. This allows enough design freedom to precisely parameterize three-dimensional shapes and handle many practical constraints. However, efficiently computing gradients requires substantial implementation effort, particularly for complicated and coupled numerical models.

This work demonstrates the first application of gradient-based optimization to the design of ducted heat exchangers, including the bidirectional coupling of the duct and heat exchanger physics and design. We use Reynolds-averaged Navier–Stokes (RANS) CFD to simulate the fluid flow inside and outside the duct. An empirical heat exchanger model predicts the pressure drop and heat transfer, using averaged inlet flow conditions from the CFD at the heat exchanger’s face. At the heat exchanger’s exit, we match flow conditions between the heat exchanger and CFD models. The optimizer adjusts the duct shape, heat exchanger geometry, and coolant pumping rate to minimize the power needed to overcome the ducted heat exchanger’s drag, weight, and pump power.

2. Methods

We separate the ducted heat exchanger into two independently-modeled regions shown in Figure 2: 1) the fluid domain in and around the duct, and 2) the heat exchanger. The fluid domain is modeled with RANS CFD, discussed in Section 2.1. The temperature and pressure changes through the heat exchanger are computed using empirical correlations, described in Section 2.2.

The fluid state depends on the heat transfer and pressure drop through the heat exchanger, and the heat exchanger’s heat transfer and pressure drop depend on the fluid state. This two-way coupling is partially handled by MPhys¹ and partially by using the optimizer as a solver. MPhys is a tool that

¹<https://github.com/OpenMDAO/mphys>




Figure 2 – The ducted heat exchanger is split into the fluid and heat exchanger domains. The Reynolds-averaged Navier–Stokes equations model the fluid, while the heat exchanger performance is predicted by empirical relationships.

takes advantage of OpenMDAO² [36] to perform high-fidelity multidisciplinary analyses and efficiently compute derivatives of the coupled system’s outputs with respect to design variables. The optimizer takes advantage of these derivatives to determine the search direction in the design space. MPhys has been used to couple heat transfer and fluid analysis models [37]. It has also enabled mixed-fidelity thermodynamic coupling [23, 25, 38], similar to this work. Section 2.3 describes the coupling method in detail.

2.1 Fluid domain

ADflow³ [39] computes the fluid flow in and around the duct. It uses the second-order finite volume method to solve the compressible RANS equations for structured multiblock and overset meshes. To converge to steady state, ADflow uses multigrid, approximate Newton–Krylov, and Newton–Krylov methods [40]. We use the approximate Newton–Krylov solver to converge the total residual 12 orders of magnitude relative to the freestream residual. In this work, we use the Spalart–Allmaras turbulence model [41]. He et al. [22] compare Spalart–Allmaras, $k-\varepsilon$, and $k-\omega$ SST turbulence models to experimental data [42] for an internal flow problem with heat transfer and find that the Spalart–Allmaras model aligns most closely with the experiment.

The optimizer requires gradients of functions of interest, such as drag, with respect to the design variables, such as the duct’s shape. To compute these quantities, ADflow solves the discrete adjoint equations. It takes advantage of algorithmic differentiation to compute many of the terms in those equations [43]. Using the adjoint method enables the computation time of these gradients to be nearly independent from the number of design variables, depending instead on the number of functions of interest [35, Sec. 6.7].

We generate a structured multiblock mesh for the duct and its internal region with Fidelity Pointwise. We assume that the duct is horizontally symmetric, so we simulate only the $+y$ (starboard) side and place a symmetry boundary along the $x-z$ plane. We extrude a volume mesh from the duct surface and then fill the duct’s interior with a butterfly topology. The resulting mesh is extruded 100 chord lengths into the farfield with pyHyp⁴ [44], a tool based on work by Chan and Steger [45]. The mesh used for optimization (Figure 3) has 1,240,704 cells and a maximum y^+ between 0.65 and 0.78 on the baseline geometry, depending on the flight condition. As the duct’s shape changes throughout the optimization, the CFD volume mesh is warped to fit the new surface. We use an inverse-distance weighting method, implemented in IDWarp⁵ [44], to warp the mesh.

We perform a mesh convergence study by solving the coupled physics (see Section 2.3) at a range of

²<https://openmdao.org/>

³<https://github.com/mdolab/adflow>

⁴<https://github.com/mdolab/pyhyp>

⁵<https://github.com/mdolab/idwarp>




Figure 3 – The structured multiblock mesh used for optimization has 1.2 million cells.

mesh refinement levels. The results at the economy cruise condition are listed in Table 1 and shown in Figure 4. We observe an order of accuracy of 1.91, close to the theoretical second-order value. The grid error convergence, computed using grid convergence indices [46], is 98% of what would be expected with a 1.91 order of accuracy. This indicates that the meshes are within the asymptotic region. Furthermore, the boundary condition pressure and temperature values change by only 0.02% between the 1.2 and 9.9 million cell meshes. For optimization purposes, it is more important for the functionals to capture the correct trends than to very closely represent their true values without discretization error. Coppeans et al. [47] show that, despite slight differences in the optimized shape, airfoils optimized with coarser meshes achieve nearly the same objective function value when reanalyzed on a fine mesh as optimizations directly with the fine mesh. Lyu et al. [48] demonstrate similar behavior for transonic wing optimizations, finding that even meshes with an estimated 6% error in drag represent the design space sufficiently well. However, the validity of this conclusion depends on the key flow features in the solution and the optimization problem formulation.

Table 1 – Mesh convergence study on the baseline geometry at the economy cruise condition. We use the second finest mesh for optimization.

Cells	Drag (N)	Heat transfer (kW)	Cooling power (kW)	p_{in} (kPa)	$p_{t,in}$ (kPa)	$T_{t,in}$ (K)
9,925,632	302.3	612.6	60.07	46.70	44.61	314.09
1,240,704	309.6	612.3	61.41	46.69	44.60	314.09
180,200	336.7	611.1	66.41	46.63	44.56	314.13

2.2 Heat exchanger

The heat transfer and pressure drop estimates are based on an empirical crossflow plate–fin compact heat exchanger model with an offset strip fin geometry [17]. The model is implemented in the Open-Concept⁶ [49] aircraft design toolkit. It uses the effectiveness–NTU method for computing the heat transfer [8, 9]. The pressure drop and heat transfer predictions use experimental correlations from Manglik and Bergles [50] to estimate the Fanning friction factor and Colburn factor. For additional modeling details, see Brelje et al. [17].

⁶<https://github.com/mdolab/openconcept>




Figure 4 – The meshes achieve a 1.91 observed order of accuracy and are in the asymptotic region.

The heat exchanger is assumed to be a rectangular prism built of layers of streamwise channels through which air flows and spanwise channels through which liquid coolant is pumped. The three dimensions of the rectangular prism are derived from the lengths of lines projected to the CFD mesh using pyGeo [51], shown in Figure 5. The length of the heat exchanger in the streamwise direction is the average length of the lines projected from the heat exchanger face to exit boundaries. The other two dimensions, the height and width of the heat exchanger's face, are defined such that the heat exchanger face has the same area and height-to-width ratio as the CFD mesh's heat exchanger inlet and exit boundaries (Figure 6). The height and width measurements used to compute the height-to-width ratio are averaged between the face and exit boundaries. As the duct shape changes, these lengths are automatically updated to match the new geometry.




Figure 5 – The height, width, and depth of the heat exchanger's CFD boundaries are each measured twice and averaged. The black lines show the measured lengths. The blue surfaces are the heat exchanger face and exit boundaries.

2.3 Fluid and heat exchanger coupling

The heat exchanger's heat transfer and pressure loss depend on the fluid conditions at its face, which are determined by the CFD. In turn, the CFD fluid state depends on the outlet conditions computed by the heat exchanger model. We use a boundary condition formulation (Figure 7) to integrate the




Figure 6 – The heat exchanger height and width are computed to match the area and height-to-width ratio of the CFD mesh’s heat exchanger boundaries.

CFD and empirical heat exchanger model. The CFD domain does not include the heat exchanger region. Instead, subsonic outflow and inflow boundary conditions are applied in the CFD domain at the heat exchanger’s face and exit, respectively. The heat exchanger model cannot consider spatially varying properties across its face, so we average the states to define a single pressure, temperature, and so on to pass from the CFD to the heat exchanger model. At the heat exchanger exit, the fluid states from the one-dimensional heat exchanger model must match the spatially-varying CFD model. We use a subsonic inlet boundary condition at the heat exchanger exit, which provides a relatively uniform fluid state. Heat exchangers tend to have relatively uniform outflow conditions with a velocity direction normal to the heat exchanger exit [29]. Our coupling method approximates this behavior.




Figure 7 – The physics coupling the CFD and heat exchanger models, shown with an eXtended Design Structure Matrix [52]. The optimizer drives the residuals to zero by solving for the boundary condition states— p at the heat exchanger face outflow boundary and p_t and T_t at the heat exchanger exit inflow boundary. Mass-averaged quantities at the heat exchanger face boundary computed by the CFD are passed to the heat exchanger. Residuals match the pressure and temperature at the heat exchanger exit between the two models.

The transfer of information between the CFD and heat exchanger model should conserve mass, momentum, and energy. To conserve mass, we feed the mass flow rate from the CFD at the heat exchanger face to the heat exchanger model. The heat exchanger model maintains constant mass flow rate. Thus, conserving mass through the whole system requires that the two CFD heat exchanger boundaries have the same mass flow rate. We form a residual that is the difference between the mass flow rate computed by the CFD at the outlet and inlet boundaries:

$$r_{\text{mass}} = \dot{m}_{\text{in}}^{\text{CFD}} - \dot{m}_{\text{out}}^{\text{CFD}} \quad (1)$$

Momentum and energy are conserved at the face and exit interfaces by matching two intensive flow

states at each. At the face, total pressure, total temperature, and static pressure are fed to the heat exchanger model (only two are needed but passing three simplifies the heat exchanger model's calculations). At the exit, we form residual equations between the heat exchanger and CFD models for total pressure and static temperature to ensure that two intensive states match:

$$r_{\text{pressure}} = p_{t, \text{out}}^{\text{HX}} - p_{t, \text{out}}^{\text{CFD}} \quad (2)$$

$$r_{\text{temperature}} = T_{\text{out}}^{\text{HX}} - T_{\text{out}}^{\text{CFD}} \quad (3)$$

We use mass averaging to convert spatially-varying quantities to a single value at the heat exchanger CFD boundaries. Mass and area averaging of intensive quantities both return similar values under most practical conditions. However, an optimizer will exploit any shortcoming of the model. We test both mass-averaged and area-averaged quantities. With area-averaged quantities, the optimizer shapes the inlet to reverse flow at the heat exchanger face's boundary. We suspect that the optimizer drives toward reversed flow conditions for two reasons. The first is physical: slower flow into the heat exchanger reduces the pressure drop across the heat exchanger and the skin friction on the inlet walls. The second reason is that the optimizer is exploiting numerical shortcomings in the model. This is likely a combination of the heat exchanger model ignoring any spatial variation in the incoming flow and ADflow allowing inflow from a subsonic outflow boundary condition. Area averaging may exacerbate these problems; reversed flow at the boundary reduces the mass flow without affecting the pressure and temperature fed to the heat exchanger model. The reduced mass flow lowers the dynamic pressure according to the heat exchanger model, which decreases the heat exchanger pressure drop and lessens drag. Using mass-averaged quantities helps discourage, though does not fully eliminate, this behavior because flow reversal also reduces the mass-averaged pressure and temperature (Figure 8).




Figure 8 – A simplified averaging example with uniform flow properties except that half is flowing in, while the other half is flowing out. Mass averaging accounts for the different flow directions, returning an average pressure of 0 kPa. Area averaging does not consider direction and returns the fluid's pressure of 100 kPa.

Because we are using intensive states, we must also match the areas of each model at the interfaces. The heat exchanger model assumes constant cross sectional area and inherits that area from the CFD mesh as described in Section 2.4. Thus, all that is left is to enforce that the heat exchanger inlet and exit boundaries in the CFD mesh have equal area:

$$r_{\text{area}} = A_{\text{in}}^{\text{CFD}} - A_{\text{out}}^{\text{CFD}} \quad (4)$$

The unknown states needed to specify the subsonic outlet (heat exchanger face) and inlet (heat exchanger exit) CFD boundary conditions are static pressure at the subsonic outlet and total pressure and temperature at the subsonic inlet. This results in three states (outlet static pressure, inlet total temperature, and inlet total pressure) and three residuals (mass flow, static temperature, and total pressure) for every flight condition included in the optimization. Rather than solving this coupled problem at every optimization iteration, we use the optimizer as a solver by passing the states as design variables and residuals as equality constraints that must equal zero. A single area residual constraint is added that can be met with the duct's geometric variables.

Even with mass averaged quantities at the heat exchanger CFD boundaries, the optimizer still tends to reverse the flow at the heat exchanger face. We prevent this by enforcing a minimum normal velocity at the heat exchanger inlet boundary. To do this, we define an array of patches at the heat exchanger face boundary (Figure 9). Each patch computes the mass-averaged velocity in the x direction (normal to the boundary) within its region. A Kreisselmeier–Steinhauser (KS) aggregation [53] then computes a smooth minimum of the mass-averaged x velocities from all patches. The aggregated minimum velocity becomes a constraint in the optimization problem.




Figure 9 – The mass-averaged x velocity is independently computed for each patch. A smooth minimum of all computed x velocities becomes a constraint in the optimization problem to provide a buffer from flow reversal.

In addition to preventing numerical exploitation of the model, the minimum velocity constraint serves the physical purpose of providing a buffer from flow reversal. If we set the lower x velocity limit to 0 m/s, the optimizer can design the duct to be right on the edge of flow reversal at the simulated flow conditions. However, the ducted heat exchanger must also perform well in off-design flight conditions that are not included in the optimization problem. Furthermore, flow separation is a notoriously challenging problem to simulate—the RANS results will not perfectly reflect reality. Setting the minimum allowed x velocity to some positive value provides a flow reversal buffer that maintains off-design performance and helps protect against numerical uncertainty. In Section 3.2, we test a range of minimum x velocity values to investigate the constraint’s impact on the ducted heat exchanger’s design.

2.4 Geometry

The baseline duct geometry is a NACA 9408 lofted along a superellipse with an exponent of 5. The baseline duct is 1.5 m long and 0.6 m tall by 0.4 m wide at its narrowest interior section. The initial heat exchanger is 0.1 m long and located at the streamwise center of the duct.

To deform the duct shape, we use the free-form deformation (FFD) method [54], implemented in pyGeo⁷ [51]. The optimizer moves the FFD volume’s B-spline control points to shape the duct. We use a two-level FFD hierarchy. The parent FFD (Figure 10a) controls the variable exit flap and the inlet, nozzle, and heat exchanger lengths. The exit flap design variable scales the height and width of the parent FFD’s last streamwise face. It also scales the height and width of the second-to-last face by half the variable value. The three layers in between the heat exchanger exit and second-to-last face prevent the exit flap deformations from warping the heat exchanger exit boundary. The reason for the three layers is a result of the FFD using cubic splines; the splines’ basis functions have no influence past three control points. The child FFD (Figure 10b) controls the shape of the

⁷<https://github.com/mdolab/pygeo>

duct. The control points in the two outermost layers can move in the radial direction along the y - z plane. The control points in the three innermost layers in the child FFD are held constant; they are added to help smoothly propagate interior duct shape changes along the heat exchanger face and exit boundaries. The child's FFD control points are embedded within the parent FFD; changes to the parent FFD variables are applied to the child FFD before the child FFD's shape variables are modified. This FFD setup keeps the heat exchanger inlet and exit boundaries nearly flat while maintaining the quality of the volume mesh. Keeping the boundaries flat is important because otherwise the optimizer can numerically exploit the boundary condition formulation by warping the boundary faces. Both FFDs and their control point deformations are mirrored across the symmetry plane. This preserves continuity across the symmetry plane to avoid sharp corners.




Figure 10 – A two-level geometry parameterization maintains desired properties of the mesh, such as flat heat exchanger boundaries, while allowing large changes to the duct's size and shape.

At times, the optimizer found strange ways to exploit the parameterization, adding bumps and wiggles to further decrease the objective. To prevent this behavior, we add a handful of constraints formed from linear combinations of shape variables. One linear constraint enforces that the difference between the inner and outer shape control points along the symmetry plane at the second-to-last top streamwise section is no greater than the average of the same quantity in the FFD slices along the corner and side of the duct. A similar constraint is imposed for the corresponding control points on the bottom of the duct. We also add linear constraints to force every pair of trailing edge control points to move together to maintain the trailing edge thickness. Two more linear constraints dictate that the two pairs of control points along the symmetry plane at the inner top and bottom of the duct before and after the heat exchanger move an equal and opposite amount in the vertical direction. These prevent the optimizer from using the parameterization to shift the entire duct vertically, which it managed to do with a less restrictive parameterization. Finally, we impose a lower bound of -0.15 on the trailing edge control points at the two corners of the superellipse.

Lastly, we add thickness and radius constraints to achieve a practical duct design (Figure 11). For structural, manufacturability, and component packaging reasons, the duct must have some thickness between the interior and exterior surfaces. The thickness constraints account for these practical considerations. Another practical consideration is the leading edge radius of the inlet lip. The radius must be sufficiently large to avoid flow separation at a range of possible operating conditions. The full range of operating conditions is not considered in the optimization problem, so the leading edge radius constraints protect against degraded off-design performance.




Figure 11 – Thickness and leading edge radius constraints on the baseline geometry.

2.5 Optimization problem

The optimization problem (Table 2) minimizes the power demanded by the ducted heat exchanger subject to heat transfer, geometric, and coupling residual constraints. The optimizer can vary the duct geometry, heat exchanger hot (coolant) and cold (air) side channel geometry, and a variable exit flap position at each flight condition. We perform a three-point multipoint optimization that roughly represents a fuel cell propeller aircraft problem. Table 3 lists the three flight conditions. The takeoff point provides a demanding heat transfer condition. Two cruise points ensure that drag is low at a range of cruise flight conditions. The objective is a weighted sum of the total cooling power, P_{total} , at each of the flight conditions. The chosen weightings incentivize low cooling power in cruise (two-thirds of the objective) and add a penalty for high cooling power in takeoff (one-third of the objective).

We define the cooling power as the sum of the power required to overcome drag, power required to overcome weight, and electricity to drive the coolant pump. All terms must represent power from the same source to be able to sum them. For example, it would be unfair to sum thrust power and electrical power because one unit of thrust power demands 1.3 units of electrical power to overcome motor and propeller losses. We put all powers in terms of electrical power, which could be provided by a fuel cell or battery. The three terms are defined as

$$P_{\text{drag}} = \frac{Dv}{\eta_{\text{prop}}\eta_{\text{motor}}} \quad (5)$$

$$P_{\text{weight}} = \frac{Wv}{\frac{L}{D}\eta_{\text{prop}}\eta_{\text{motor}}} \quad (6)$$

$$P_{\text{pump}} = \frac{\dot{m}_{\text{coolant}}\Delta p_{\text{coolant}}}{\rho_{\text{coolant}}\eta_{\text{pump}}} \quad (7)$$

$$P_{\text{total}} = P_{\text{drag}} + P_{\text{weight}} + P_{\text{pump}} \quad (8)$$

Drag is computed by integrating the force and momentum on the duct surfaces, including the heat exchanger face and exit. Weight is the sum of the heat exchanger and duct weights. The heat exchanger weight is computed by OpenConcept as the volume of the plates and fins times the density of aluminum. The duct weight is assumed to scale with its surface area—we use an areal density of 3 kg/m^2 . All relevant parameters are listed in Table 4.

In the aircraft design process, drag and weight changes affect the aircraft sizing. These effects change the amount of heat that must be rejected and could increase the cooling power, depending on how it is bookkept. While subtle changes are not included in our objective, we expect them to scale

Table 2 – The optimizer minimizes the average cooling power of two cruise points and a takeoff point.

	Variable	Quantity
minimize	$1/3 (P_{\text{total, rto}} + P_{\text{total, ecrz}} + P_{\text{total, hscrz}})$ (Equation 8)	
by varying	duct shape	96
	inlet length	1
	heat exchanger length	1
	nozzle length	1
	variable exit flap	2
	heat exchanger air side channel height	1
	heat exchanger air side channel width	1
	heat exchanger coolant side channel height	1
	heat exchanger coolant side channel width	1
	coolant mass flow rate	3
	heat exchanger inlet static pressure	3
	heat exchanger outlet total pressure	3
	heat exchanger outlet total temperature	3
	117 design variables	
subject to	heat transfer $\geq$ heat transfer at flight condition (Table 3)	3
	mass flow residual = 0 (Equation 1)	3
	outlet total pressure residual = 0 (Equation 2)	3
	outlet temperature residual = 0 (Equation 3)	3
	area residual = 0 (Equation 4)	1
	duct thickness $\geq$ baseline thickness	1,000
	duct leading edge radius $\geq$ baseline leading edge radius	20
	minimum x velocity at heat exchanger face $\geq$ 15 m/s	3
	bump linear constraints (Section 2.4)	2
	trailing edge linear constraints (Section 2.4)	6
	vertical shift linear constraints (Section 2.4)	2
	1,046 constraints	

Table 3 – The flight conditions include a takeoff with a demanding heat transfer constraint and two cruise points.

	Takeoff	Economy cruise	High-speed cruise
Mach	0.25	0.45	0.5
Altitude	1,000 ft	23,000 ft	25,000 ft
Heat transfer	600 kW	400 kW	450 kW
Angle of attack	4°	1°	0°

Table 4 – Parameters are chosen to roughly represent current technology.

Parameter	Value	Notes
L/D	Lift-to-drag ratio	17
η_{prop}	Propulsive efficiency	80%
η_{motor}	Motor efficiency	95%
η_{pump}	Coolant pump efficiency	75%
ρ_{coolant}	Coolant density	1,020 kg/m ³ Water glycol mixture
T_{coolant}	Coolant inlet temperature	90°C
	Duct areal density	3 kg/m ²

with our objective. In other words, the ducted heat exchanger design that minimizes our objective is similar to the one that would be best if this iterative design were considered.

Constraints ensure that the ducted heat exchanger rejects sufficient heat at the specified operating conditions. We also use the optimizer to converge the coupling residuals between the heat exchanger

and CFD models, discussed in Section 2.3. The coupling residuals ensure that the two models are consistent, conserving mass, momentum, and energy. Thickness constraints force the optimizer to return designs that are thick enough to be manufacturable and contain the necessary structure and components. Leading edge radius constraints help produce inlet designs that perform adequately at flight conditions that are not included in the optimization problem. Finally, we constrain the minimum velocity into the heat exchanger boundary, as described in Section 2.3.

We use SNOPT [55] through the pyOptSparse interface⁸ [56] to perform the optimization. SNOPT is a gradient-based optimizer, enabling it to efficiently handle a large number of design variables [35, Sec. 1.4]. This allows us to give the optimizer significant design freedom to precisely shape and size the duct and heat exchanger. We converge all optimizations to an optimality metric of 10^{-5} and feasibility metric of 10^{-7} , as defined by SNOPT. The optimizations run on NASA's Pleiades cluster using Intel Cascade Lake Xeon Gold 6248, Skylake Xeon Gold 6148, and Broadwell Xeon E5-2680v4 nodes.

3. Results

The optimizer transforms the arbitrary and poorly-designed baseline, which cannot meet the takeoff heat transfer constraint, into an efficient and refined ducted heat exchanger—slashing cruise drag by five times. To better understand the optimum design, we perform additional analyses and optimizations to study the influence of the heat transfer in Section 3.1. We also examine the effect of the minimum x velocity constraint on the optimum design in Section 3.2.

Figure 12 shows the optimized ducted heat exchanger analyzed at the high-speed cruise condition. The most striking change from the baseline is that the optimizer has turned the superelliptical duct into a shape that is much closer to a circle. This change reduces the duct surface area for a given heat exchanger frontal area. A lower surface area results in a lighter and lower drag duct, both of which lessen the cooling power objective. It is not perfectly circular due to a combination of including nonzero angles of attack and limitations of the geometric parameterization. Figure 13 plots SNOPT's optimality and feasibility metrics throughout the optimization, which took 1.6 days on 364 processors.




Figure 12 – The high-speed cruise flight condition with pressure coefficient contours on the duct and temperature contours on the symmetry plane. The optimizer rounds the corners of the superelliptical baseline design to produce a far more circular duct shape.

⁸<https://github.com/mdolab/pyoptsparse>




Figure 13 – SNOPT converges tightly to a minimum in the design space that meets all constraints. The tightly-converged feasibility ensures that the residuals are satisfied, giving a physical result.

Figure 14 breaks down the drag, weight, and pump power contributions to the cooling power objective at each flight condition. The takeoff flight condition has the highest cooling power. This should not be too surprising, since the takeoff point has the most demanding heat transfer constraint. Takeoff also makes up only one third of the objective, so the optimizer is incentivized to give up takeoff performance if it enables improved cruise performance.

The demanding takeoff heat transfer constraint drives the design toward a ducted heat exchanger that can transfer more heat than is needed at the cruise conditions. Two design characteristics help reduce this oversizing: a higher coolant mass flow rate (twice the pumping power) and a more open flap position at takeoff than at cruise. The higher coolant mass flow rate (listed in Table 5) enables more heat transfer within the same heat exchanger frontal area. A more open flap position (Figure 15) increases the mass flow rate of air through the duct. A greater air mass flow rate enables more heat transfer, but adds drag because the higher dynamic pressure at the heat exchanger's face increases the pressure drop across it. Both of these features sacrifice cooling power at takeoff through increased pump power and drag. However, they help make the ducted heat exchanger more appropriately sized for the two cruise conditions that carry more weight in the objective.




Figure 14 – A breakdown of the contributions to total cooling power at each of the three flight conditions. The optimizer produces a design that sacrifices pump power and drag at takeoff to reduce cruise drag.




Figure 15 – The optimized design opens the flap to increase the nozzle area by nearly 50% at the takeoff flight condition. The larger exit increases the mass flow rate of air at takeoff, limiting how oversized the ducted heat exchanger is in cruise.

The pressure bars in Figure 16 illustrate how the ducted heat exchanger design manipulates the dynamic pressure along its length at the three flight conditions. The height of the bar represents the freestream compressible dynamic pressure—the difference between the freestream total pressure and freestream static pressure. Approaching the duct inlet, the dynamic pressure (light blue) is converted to static pressure (dark navy) by slowing the incoming flow. There is a total pressure loss across the heat exchanger, noted by the red section downstream of the heat exchanger. The total pressure loss is a function of the dynamic pressure at the heat exchanger's face. Total pressure losses cause a drag penalty and represent energy that cannot be recovered in the duct's nozzle.




Figure 16 – A graphical representation of the ducted heat exchanger's manipulation of dynamic pressure along its length. The total bar height is the freestream compressible dynamic pressure. The light blue region is the dynamic pressure at a given streamwise position, while the dark navy region is the difference between the static pressure and the freestream static pressure. Total pressure loss is denoted by the red region.

At the cruise conditions (Figure 16b and Figure 16c), most of the compression upstream of the heat exchanger happens before even entering the inlet. This phenomenon is consistent with the inlet area being more than twice the nozzle area in cruise, as shown in Figure 15, because the dynamic pressure at the nozzle exit is much greater than at the front of the inlet. This means a streamtube capturing the same mass flow will be larger at the inlet lip than the nozzle exit. By the time the flow

reaches the heat exchanger's face, 98% of the freestream dynamic pressure is converted to static pressure. The flow slows to a Mach number of about 0.06 and achieves a static pressure 1.15–1.18 times higher than freestream. The dramatic reduction in dynamic pressure of the incoming flow reduces the total pressure loss across the heat exchanger to only 11% of the freestream compressible dynamic pressure.

The takeoff condition (Figure 16a) exhibits the same trends as the cruise conditions. However, the more open flap position results in only 95% of the dynamic pressure being converted to static pressure. This, combined with nearly twice the air mass flow rate as the cruise conditions, results in a total pressure loss that is almost double that of the cruise conditions relative to the freestream compressible dynamic pressure.

Figure 17 depicts Mach number contours along the symmetry plane at the three flight conditions. The plots visualize a similar effect as the static pressure bar in Figure 16: the streamtube ingested into the duct, shown by the streamlines, expands much less for the takeoff condition (top) than the two cruise conditions. They also illustrate how, by controlling the exit flap, the optimizer returns a design where the Mach number at the heat exchanger face is nearly constant across the simulated flight conditions, despite the dramatic changes in freestream conditions.




Figure 17 – Mach contours of the optimized design at takeoff, high-speed cruise, and economy cruise, respectively. The Mach number at the heat exchanger's face is relatively constant across all flight conditions. This is a result of the optimized design ingesting a larger streamtube and compressing the air less at the takeoff condition. This allows it to meet its heat transfer constraint without sacrificing too much cruise performance.

An interesting feature shown in Figure 17 is how much the stagnation point moves between flight conditions. At the takeoff flight condition (top), the stagnation point, indicated by the darkest blue region at the inlet lip of near-zero Mach number, is close to the leading edge of the lower lip. At the two cruise conditions, the stagnation point is considerably beyond the leading edge on the duct's interior surface. The streamlines near the stagnation point that follow the duct's exterior surface must sharply curve around the leading edge, causing a low pressure region along the inlet lip that may reduce drag. This example demonstrates the impressive potential of optimization to automate inlet design. The optimizer returns a short and efficient inlet design that manages to avoid interior and exterior flow separation at a representative range of flight conditions.

Table 5 – Sizing the ducted heat exchanger to meet the takeoff heat transfer causes it to be oversized for cruise. The optimizer finds that transferring more than the minimum required amount of heat in cruise enables it to reduce drag by recovering some of the thermal energy. It maintains roughly consistent heat transfer across all flight points.

	Takeoff	Economy cruise	High-speed cruise
Heat transfer	600 kW	575 kW	616 kW
Coolant mass flow	13.3 kg/s	9.5 kg/s	10.5 kg/s
Total weight		47.5 kg	
Air side channel width		1.83 mm	
Air side channel height		8.28 mm	
Coolant side channel width		5.23 mm	
Coolant side channel height		0.74 mm	

3.1 Impact of heat transfer

In theory, the thermal energy transferred from the coolant to the air can be recovered to offset some of the drag by expanding the heated air in the nozzle [3, 6]. The magnitude of this effect depends on the amount of ram pressure, which depends on freestream dynamic pressure. It also depends on the amount of heat transferred to the air, which can be varied with air mass flow rate, coolant temperature, and coolant mass flow rate. To quantify this effect, we perform analyses and optimizations with a modified temperature boundary condition residual that matches the total temperature of the heat exchanger face and exit boundaries. By comparing the result of these simulations to the ones where the temperatures account for the heat transfer, we can better understand the influence of the energy recovery on the drag and design choices.




Figure 18 – The optimized ducted heat exchanger design uses the transferred heat to offset 70% of the cruise drag.

We start by performing an analysis on the optimized design with the modified temperature residual to remove the heat transfer's effect on drag. The optimizer acts as a solver for the analysis, determining the static pressure, total pressure, and total temperature boundary condition states at each flight condition that drive the mass flow, total pressure, and modified temperature residuals to zero. Figure 18 compares the drag of the optimized design at each flight condition with and without considering heat transfer. The optimized design manages to offset 70% of the drag in cruise by recovering the thermal energy. Another theoretical trend this figure reveals is that the drag reduction from energy recovery increases with increasing flight speed, in part thanks to the increase in available ram pressure.

Next, we run an optimization with the modified temperature residual to investigate how the optimum design changes when the heat transfer is neglected. The result has active heat transfer constraints at all flight conditions because the optimizer has no incentive to increase the coolant mass flow rate beyond what is needed to meet the required heat transfer. This is unlike the optimization considering heat transfer, which increases coolant mass flow rate to transfer more heat than is required—the drag reduction from the additional heat outweighs the pump power increase in the objective. Figure 19 compares the cross section shape of the duct designed without considering heat transfer to the design optimized while considering heat transfer. The most noticeable change is the smaller cross sectional area. This is consistent with trying to exactly meet all heat transfer constraints, rather than accepting an oversized duct for cruise to improve takeoff performance and making up for it by transferring extra heat in cruise.




Figure 19 – The optimized ducted heat exchangers with and without considering heat transfer at the economy cruise point. The design optimized while neglecting heat transfer is dramatically smaller and has active heat transfer constraints at all flight conditions, unlike the design optimized while considering heat transfer.

These results have consequences for the thermal management system architecture design because there may be points in the flight envelope where the ducted heat exchanger wants more heat than the propulsion system is producing. In practice, the amount of heat transferred must match the amount dictated by the aircraft's thermal management system. For the chosen flight conditions, the ducted heat exchanger's optimal heat transfer throughout the flight is roughly set by the peak required heat transfer (600 kW in this case). Thus, it is likely in the designer's interest to limit peak required heat transfer loads to appropriately size the ducted heat exchanger for a broader range of flight conditions. Though the effect may be small, this could also nudge the aircraft design toward a flight profile where the duct is more appropriately sized throughout the flight. For example, flying faster in cruise would produce more heat and cruising at higher altitudes would slightly decrease the amount of heat the ducted heat exchanger can transfer.

3.2 Effect of minimum x velocity constraint

This section investigates the effect of the minimum aggregated x velocity constraint described in Section 2.3. To study its impact, we converge optimizations for constraint lower bounds from 12.5 m/s to 20 m/s. In all optimizations, the constraint is active at the takeoff and economy cruise conditions, but not at the high-speed cruise point. Based on the results, this constraint effectively provides a buffer from flow separation in the inlet. Tightening the constraint has little effect on the cooling power objective in cruise, but it does cause notable cooling power increases at takeoff.

Figure 20 plots how the cooling power objective at each flight conditions varies as the constraint changes. The cross sections on the top are at the economy cruise flap position. The figure shows how, as the minimum allowed x velocity increases, the takeoff performance is most strongly impacted. Part of this effect could be due to the optimizer returning a design with a longer and steeper inlet, which helps reduce the dynamic pressure at the heat exchanger face during takeoff. Cruise performance is largely unaffected, though it does begin to worsen at a high enough constraint value. At constraint values less than 12.5 m/s, the CFD solver is unable to converge intermediate designs due to inlet separation. Thus, a constraint value of 15 m/s provides a buffer from this separation and improves optimization robustness without dramatically impacting cruise cooling power.




Figure 20 – As the minimum x velocity constraint is increased, the takeoff cooling power is most strongly impacted, but cruise performance eventually begins to worsen. The duct cross sections are depicted on the top at the economy cruise flap position.

4. Conclusion

Electric and fuel-cell aircraft produce large quantities of low-grade waste heat that must be transferred to freestream air. Ducted heat exchangers are an effective way to transfer the heat, but they can impose a substantial performance penalty if they are not well designed. The design of ducted heat exchangers is a challenging and tightly coupled-problem that involves trading off drag and heat transfer at different points in the flight envelope through judicious duct and heat exchanger design.

This work demonstrates the first application of CFD-based optimization to the coupled duct shape and heat exchanger design and sizing problem. We use a geometric parameterization that gives the optimizer a large design space for the duct shape. The parameterization can deform the shape from rectangular to circular cross sections, modify the duct airfoil shapes, change cross sectional area along the duct, and independently vary the length of the inlet, heat exchanger, and nozzle. The optimization problem also incorporates heat exchanger size and channel geometry design variables. The optimizer minimizes the average cooling power penalty at three flight conditions: a heat-limited takeoff point and two cruise points.

The optimization process transforms the badly-designed baseline into a design that tactfully meets the takeoff heat transfer requirement while reducing the drag at the two cruise points by nearly five times. At takeoff, the optimizer uses a high coolant mass flow rate and open nozzle flap to transfer more heat in a smaller duct. This approach limits how oversized the duct is for cruise. In cruise, the optimized design has higher coolant mass flow rates than are needed to meet the heat transfer constraint. By transferring more heat than required, the design reduces cruise drag by 70% compared to the same design without heat transfer. The design also reduces the duct's drag and weight by performing most of the inlet compression upstream of the inlet lip. To do this, the design controls the air mass flow with the nozzle exit flap and shapes the inlet appropriately.

This work demonstrates the versatility of multidisciplinary optimization for designing ducted heat exchangers. Optimization not only automates the design process, but enables engineers to efficiently explore the design space, understand tradeoffs, and discover nuances in the coupled design that unlock performance improvements for future electrified aircraft.

Acknowledgements

The first author is supported by the Department of Defense through the National Defense Science and Engineering Graduate (NDSEG) Fellowship Program and by the University of Michigan Rackham Predoctoral Fellowship Program. He is also supported in part by the Michigan Institute for Computational Discovery and Engineering (MICDE) Graduate Fellowship program. Computational resources supporting this work were provided by the NASA High-End Computing (HEC) Program through the NASA Advanced Supercomputing (NAS) Division at Ames Research Center.

5. Contact Author Email Address

The first author can be contacted at eytana@umich.edu.

6. Copyright Statement

The authors confirm that they hold copyright on all of the original material included in this paper. The authors confirm that they give permission for the publication and distribution of this paper as part of the ICAS proceedings or as individual off-prints from the proceedings.

References

- [1] Adler, E. J., and Martins, J. R. R. A., "Hydrogen-Powered Aircraft: Fundamental Concepts, Key Technologies, and Environmental Impacts," *Progress in Aerospace Sciences*, Vol. 141, 2023, p. 100922. <https://doi.org/10.1016/j.paerosci.2023.100922>.
- [2] Brelje, B., and Martins, J. R. R. A., "Electric, Hybrid, and Turboelectric Fixed-Wing Aircraft: A Review of Concepts, Models, and Design Approaches," *Progress in Aerospace Sciences*, Vol. 104, 2019, pp. 1–19. <https://doi.org/10.1016/j.paerosci.2018.06.004>.
- [3] Meredith, F. W., "Cooling of aircraft engines with special reference to ethylene glycol radiators enclosed in ducts," Aeronautical research committee reports and memoranda no. 1683, Air Ministry, 1935.
- [4] Hoerner, S. F., *Fluid-Dynamic Drag*, Hoerner Fluid Dynamics, Bakersfield, CA, 1965. URL https://www.google.ca/books/edition/Fluid_dynamic_Drag/6K12uAAACAAJ.
- [5] Silverstein, A., "Experiments on the Recovery of Waste Heat in Cooling Ducts," Special Report NACA-SR-111, NACA, May 1939. URL <https://ntrs.nasa.gov/citations/20090015245>.
- [6] Rauscher, M., and Phillips, W. H., "Propulsive Effects of Radiator and Exhaust Ducting," *Journal of the Aeronautical Sciences*, Vol. 8, No. 4, 1941, pp. 167–174. <https://doi.org/10.2514/8.10671>.
- [7] Çengel, Y. A., and Ghajar, A. J., *Heat and Mass Transfer*, 5th ed., McGraw Hill, 2015.

- [8] Incropera, F. P., DeWitt, D. P., Bergman, T. L., and Lavine, A. S., *Fundamentals of Heat and Mass Transfer*, 6th ed., John Wiley & Sons, 2007.
- [9] Kays, W. M., and London, A. L., *Compact Heat Exchangers*, 3rd ed., McGraw-Hill, New York, 1984.
- [10] Hendricks, T. J., Mcenerney, B., Drymiotis, F., Furst, B., and Shevade, A., "Design and Testing of High-Performance Mini-Channel Graphite Heat Exchangers in Thermoelectric Energy Recovery Systems," *ASME International Mechanical Engineering Congress and Exposition*, American Society of Mechanical Engineers, 2017. <https://doi.org/10.1115/IMECE2017-72411>.
- [11] Anibal, J. L., and Martins, J. R. R. A., "Adjoint-based shape optimization of a plate-fin heat exchanger using CFD," *Applied Thermal Engineering*, 2024, p. 123570. <https://doi.org/10.1016/j.applthermaleng.2024.123570>.
- [12] Anibal, J., "Aerodynamic Shape Optimization of Heat Exchangers," Ph.D. thesis, University of Michigan, Ann Arbor, MI, 2023. <https://doi.org/10.7302/7308>.
- [13] Sundén, B., "Computational Fluid Dynamics in Research and Design of Heat Exchangers," *Heat Transfer Engineering*, Vol. 28, No. 11, 2007, pp. 898–910. <https://doi.org/10.1080/01457630701421679>.
- [14] Claus, R. W., Evans, A. L., Lylte, J. K., and Nichols, L. D., "Numerical propulsion system simulation," *Computing Systems in Engineering*, Vol. 2, No. 4, 1991, pp. 357–364. [https://doi.org/10.1016/0956-0521\(91\)90003-N](https://doi.org/10.1016/0956-0521(91)90003-N).
- [15] Hendricks, E. S., and Gray, J. S., "pyCycle: A Tool for Efficient Optimization of Gas Turbine Engine Cycles," *Aerospace*, Vol. 6, No. 87, 2019. <https://doi.org/10.3390/aerospace6080087>.
- [16] Jasa, J. P., Brelje, B. J., Mader, C. A., and Martins, J. R. R. A., "Coupled Design of a Supersonic Engine and Thermal System," *World Congress of Structural and Multidisciplinary Optimization*, Beijing, China, 2019.
- [17] Brelje, B. J., Jasa, J. P., Martins, J. R. R. A., and Gray, J. S., "Development of a Conceptual-Level Thermal Management System Design Capability in OpenConcept," *NATO Research Symposium on Hybrid/Electric Aero-Propulsion Systems for Military Applications (AVT-RSY-323)*, Trondheim, NO, 2019. <https://doi.org/10.14339/STO-MP-AVT-323>.
- [18] Kellermann, H., Lüdemann, M., Pohl, M., and Hornung, M., "Design and optimization of ram air-based thermal management systems for hybrid-electric aircraft," *Aerospace*, Vol. 8, No. 1, 2020, p. 3. <https://doi.org/10.3390/aerospace8010003>.
- [19] Sóbester, A., "Tradeoffs in Jet Inlet Design: A Historical Perspective," *Journal of Aircraft*, Vol. 44, No. 3, 2007, pp. 705–717. <https://doi.org/10.2514/1.26830>.
- [20] Madabhushi, R. K., Levy, R., and Pincus, S. M., "Design of Optimum Ducts Using an Efficient 3-D Viscous Computational Flow Analysis," Conference paper GRC-E-DAA-TN55806, NASA, January 1987. URL <https://ntrs.nasa.gov/citations/19930074227>.
- [21] Chiang, C., Koo, D., and Zingg, D. W., "Aerodynamic Shape Optimization of an S-Duct Intake for a Boundary-Layer Ingesting Engine," *Journal of Aircraft*, Vol. 59, No. 3, 2022, pp. 725–741. <https://doi.org/10.2514/1.C036632>.
- [22] He, P., Martins, J. R. R. A., Mader, C. A., and Maki, K., "Aerothermal Optimization of a Ribbed U-Bend Cooling Channel Using the Adjoint Method," *International Journal of Heat and Mass Transfer*, Vol. 140, 2019, pp. 152–172. <https://doi.org/10.1016/j.ijheatmasstransfer.2019.05.075>.
- [23] Gray, J. S., Mader, C. A., Kenway, G. K. W., and Martins, J. R. R. A., "Modeling Boundary Layer Ingestion Using a Coupled Aeropropulsive Analysis," *Journal of Aircraft*, Vol. 55, No. 3, 2018, pp. 1191–1199. <https://doi.org/10.2514/1.C034601>.

- [24] Yildirim, A., Gray, J. S., Mader, C. A., and Martins, J. R. R. A., "Boundary Layer Ingestion Benefit for the STARC-ABL Concept," *Journal of Aircraft*, Vol. 59, No. 4, 2022, pp. 896–911. <https://doi.org/10.2514/1.C036103>.
- [25] Lamkin, A. H. R., Yildirim, A., Martins, J. R. R. A., and Wukie, N. A., "Advancements in Coupled Aeropropulsive Design Optimization for High-Bypass Turbofan Engines," *AIAA Aviation Forum*, San Diego, CA, 2023. <https://doi.org/10.2514/6.2023-3591>.
- [26] Abdul-Kaiyoom, M. A. S., Yildirim, A., and Martins, J. R. R. A., "RANS-Based Multipoint Aeropropulsive Design Optimization of an Over-Wing Nacelle Configuration," *AIAA Aviation Forum*, San Diego, CA, 2023. <https://doi.org/10.2514/6.2023-3588>.
- [27] Chapman, J. W., Schnulo, S. L., and Nitzsche, M. P., "Development of a Thermal Management System for Electrified Aircraft," *AIAA Scitech Forum*, American Institute of Aeronautics and Astronautics, 2020. <https://doi.org/10.2514/6.2020-0545>.
- [28] Adler, E. J., Brelje, B. J., and Martins, J. R. R. A., "Thermal Management System Optimization for a Parallel Hybrid Aircraft Considering Mission Fuel Burn," *Aerospace*, Vol. 9, No. 5, 2022. <https://doi.org/10.3390/aerospace9050243>.
- [29] Drela, M., "Aerodynamics of heat exchangers for high-altitude aircraft," *Journal of Aircraft*, Vol. 33, No. 1, 1996, pp. 176–184. <https://doi.org/10.2514/3.46919>.
- [30] Drela, M., "Newton solution of coupled viscous/inviscid multielement airfoil flows," *21st Fluid Dynamics, Plasma Dynamics and Lasers Conference*, American Institute of Aeronautics and Astronautics, 1990. <https://doi.org/10.2514/6.1990-1470>.
- [31] Musto, M., Bianco, N., Rotondo, G., Toscano, F., and Pezzella, G., "A simplified methodology to simulate a heat exchanger in an aircraft's oil cooler by means of a Porous Media model," *Applied Thermal Engineering*, Vol. 94, 2016, pp. 836–845. <https://doi.org/10.1016/j.applthermaleng.2015.10.147>.
- [32] Patrao, A. C., Jonsson, I., Xisto, C., Lundbladh, A., and Grönstedt, T., "Compact heat exchangers for hydrogen-fueled aero engine intercooling and recuperation," *Applied Thermal Engineering*, Vol. 243, 2024, p. 122538. <https://doi.org/10.1016/j.applthermaleng.2024.122538>.
- [33] Hendricks, T. J., Tarau, C., and Dyson, R. W., "Hybrid Electric Aircraft Thermal Management: Now, New Visions and Future Concepts and Formulation," *2021 20th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (iTherm)*, 2021, pp. 467–476. <https://doi.org/10.1109/iTherm51669.2021.9503205>.
- [34] Martins, J. R. R. A., "Aerodynamic Design Optimization: Challenges and Perspectives," *Computers & Fluids*, Vol. 239, 2022, p. 105391. <https://doi.org/10.1016/j.compfluid.2022.105391>.
- [35] Martins, J. R. R. A., and Ning, A., *Engineering Design Optimization*, Cambridge University Press, Cambridge, UK, 2022. <https://doi.org/10.1017/9781108980647>, URL <https://mdobook.github.io>.
- [36] Gray, J. S., Hwang, J. T., Martins, J. R. R. A., Moore, K. T., and Naylor, B. A., "OpenMDAO: An open-source framework for multidisciplinary design, analysis, and optimization," *Structural and Multidisciplinary Optimization*, Vol. 59, No. 4, 2019, pp. 1075–1104. <https://doi.org/10.1007/s00158-019-02211-z>.
- [37] Anibal, J., Mader, C. A., and Martins, J. R. R. A., "Aerodynamic shape optimization of an electric aircraft motor surface heat exchanger with conjugate heat transfer constraint," *International Journal of Heat and Mass Transfer*, Vol. 189, 2022, p. 122689. <https://doi.org/10.1016/j.ijheatmasstransfer.2022.122689>.

- [38] Yildirim, A., Gray, J. S., Mader, C. A., and Martins, J. R. R. A., "Coupled Aeropropulsive Design Optimization of a Poded Electric Propulsor," *AIAA Aviation Forum*, 2021. <https://doi.org/10.2514/6.2021-3032>.
- [39] Mader, C. A., Kenway, G. K. W., Yildirim, A., and Martins, J. R. R. A., "ADflow: An open-source computational fluid dynamics solver for aerodynamic and multidisciplinary optimization," *Journal of Aerospace Information Systems*, Vol. 17, No. 9, 2020, pp. 508–527. <https://doi.org/10.2514/1.1010796>.
- [40] Yildirim, A., Kenway, G. K. W., Mader, C. A., and Martins, J. R. R. A., "A Jacobian-free approximate Newton–Krylov startup strategy for RANS simulations," *Journal of Computational Physics*, Vol. 397, 2019, p. 108741. <https://doi.org/10.1016/j.jcp.2019.06.018>.
- [41] Spalart, P., and Allmaras, S., "A One-Equation Turbulence Model for Aerodynamic Flows," *La Recherche Aerospaciale*, Vol. 1, 1994, pp. 5–21.
- [42] Coletti, F., Verstraete, T., Bulle, J., Van der Wielen, T., Van den Berge, N., and Arts, T., "Optimization of a U-Bend for Minimal Pressure Loss in Internal Cooling Channels–Part II: Experimental Validation," *Journal of Turbomachinery*, Vol. 135, No. 5, 2013, p. 051016.
- [43] Kenway, G. K. W., Mader, C. A., He, P., and Martins, J. R. R. A., "Effective Adjoint Approaches for Computational Fluid Dynamics," *Progress in Aerospace Sciences*, Vol. 110, 2019, p. 100542. <https://doi.org/10.1016/j.paerosci.2019.05.002>.
- [44] Secco, N., Kenway, G. K. W., He, P., Mader, C. A., and Martins, J. R. R. A., "Efficient Mesh Generation and Deformation for Aerodynamic Shape Optimization," *AIAA Journal*, Vol. 59, No. 4, 2021, pp. 1151–1168. <https://doi.org/10.2514/1.J059491>.
- [45] Chan, W. M., and Steger, J. L., "Enhancements of a three-dimensional hyperbolic grid generation scheme," *Applied Mathematics and Computation*, Vol. 51, No. 2–3, 1992, pp. 181–205. [https://doi.org/10.1016/0096-3003\(92\)90073-A](https://doi.org/10.1016/0096-3003(92)90073-A).
- [46] Roache, P. J., "Verification of Codes and Calculations," *AIAA Journal*, Vol. 36, No. 5, 1998, pp. 696–702. <https://doi.org/10.2514/2.457>.
- [47] Coppeans, A. W. C., Fidkowski, K. J., and Martins, J. R. R. A., "Comparison of Finite Volume and High-Order Discontinuous Galerkin Based Aerodynamic Shape Optimization," *AIAA SciTech Forum*, National Harbor, MD, 2023. <https://doi.org/10.2514/6.2023-1845>.
- [48] Lyu, Z., Kenway, G. K. W., and Martins, J. R. R. A., "Aerodynamic Shape Optimization Investigations of the Common Research Model Wing Benchmark," *AIAA Journal*, Vol. 53, No. 4, 2015, pp. 968–985. <https://doi.org/10.2514/1.J053318>.
- [49] Brelje, B. J., and Martins, J. R. R. A., "Development of a Conceptual Design Model for Aircraft Electric Propulsion with Efficient Gradients," *Proceedings of the AIAA/IEEE Electric Aircraft Technologies Symposium*, Cincinnati, OH, 2018. <https://doi.org/10.2514/6.2018-4979>.
- [50] Manglik, R. M., and Bergles, A. E., "Heat transfer and pressure drop correlations for the rectangular offset strip fin compact heat exchanger," *Experimental Thermal and Fluid Science*, Vol. 10, No. 2, 1995, pp. 171–180. [https://doi.org/10.1016/0894-1777\(94\)00096-Q](https://doi.org/10.1016/0894-1777(94)00096-Q).
- [51] Hajdik, H. M., Yildirim, A., Wu, N., Brelje, B. J., Seraj, S., Mangano, M., Anibal, J. L., Jonsson, E., Adler, E. J., Mader, C. A., Kenway, G. K. W., and Martins, J. R. R. A., "pyGeo: A geometry package for multidisciplinary design optimization," *Journal of Open Source Software*, Vol. 8, No. 87, 2023, p. 5319. <https://doi.org/10.21105/joss.05319>.
- [52] Lambe, A. B., and Martins, J. R. R. A., "Extensions to the Design Structure Matrix for the Description of Multidisciplinary Design, Analysis, and Optimization Processes," *Structural and Multidisciplinary Optimization*, Vol. 46, No. 2, 2012, pp. 273–284. <https://doi.org/10.1007/s00158-012-0763-y>.

- [53] Kreisselmeier, G., and Steinhauser, R., "Systematic Control Design by Optimizing a Vector Performance Index," *International Federation of Active Controls Symposium on Computer-Aided Design of Control Systems, Zurich, Switzerland*, 1979. [https://doi.org/10.1016/S1474-6670\(17\)65584-8](https://doi.org/10.1016/S1474-6670(17)65584-8).
- [54] Sederberg, T. W., and Parry, S. R., "Free-form Deformation of Solid Geometric Models," *SIGGRAPH Comput. Graph.*, Vol. 20, No. 4, 1986, pp. 151–160. <https://doi.org/10.1145/15886.15903>.
- [55] Gill, P. E., Murray, W., and Saunders, M. A., "SNOPT: An SQP Algorithm for Large-Scale Constrained Optimization," *SIAM Review*, Vol. 47, No. 1, 2005, pp. 99–131. <https://doi.org/10.1137/S0036144504446096>.
- [56] Wu, N., Kenway, G., Mader, C. A., Jasa, J., and Martins, J. R. R. A., "pyOptSparse: A Python framework for large-scale constrained nonlinear optimization of sparse systems," *Journal of Open Source Software*, Vol. 5, No. 54, 2020, p. 2564. <https://doi.org/10.21105/joss.02564>.