SIZING METHODOLOGIES FOR AIRCRAFT WITH MULTIPLE ENERGY SYSTEMS

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF
AERONAUTICS AND ASTRONAUTICS
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

J. Michael Vegh

June 2018

This dissertation is online at: http://purl.stanford.edu/gj158ch1759

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Juan Alonso, Primary Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Brian Cantwell

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Ilan Kroo

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Michael Saunders

Approved for the Stanford University Committee on Graduate Studies.

Patricia J. Gumport, Vice Provost for Graduate Education

This signature page was generated electronically upon submission of this dissertation in electronic format. An original signed hard copy of the signature page is on file in University Archives.

Abstract

Electric aircraft have been a topic of significant interest, especially in recent years. A wide variety of different aircraft configurations that utilize electric propulsion at short ranges have been proposed and are currently under development. However, physical limits in battery technology have significantly hampered their entry-into-service. A proposed solution has been to use hybrid-electric or hybrid battery approaches that allow the user to leverage multiple energy systems with contrasting specific energy and specific power properties to allow for lighter, more feasible aircraft along with the opening of new regions in the design space. Many of the designs, particularly those that use distributed electric propulsion possess a complex trade-space that needs global optimization methods to properly evaluate.

This dissertation develops a methodology for basic sizing of aircraft that enables the designer to employ global optimization methods at the earliest stages of the design process with substantial improvements in code speed, and, depending on the method used, robustness. Furthermore, it reveals an inherent weakness in developing and sizing electric and hybrid-electric aircraft that only use a single value for specific energy and specific power; the results here show representative cases where modeling off-design performance rather than simply taking a nominal value can reduce aircraft weight by as much as factor of two. This also illustrates that for electric aircraft, making intelligent design decisions in the earliest stages of the design process is much more crucial than in traditional designs where specific energy is not nearly as limiting. Additionally, it explores and diagnoses several unusual issues that may arise in the context of solving aircraft design problems that possess multiple energy systems, while offering algorithmic solutions that improve both code speed and robustness. More specifically, these designs show an unsettling characteristics in that, at the basic level, solutions to the aircraft sizing problem may be nonunique. Now, it should be noted that different methods are more susceptible to finding extraneous solutions which will be discussed in greater detail. Finally, it develops frameworks utilizing machine learning regression techniques that learn the trade space of sizing solutions, meaning that what may be computationally expensive calculations can be initiated from a nearly converged starting point.

Acknowledgments

Firstly, I would like to thank the SMART Scholarship office and the Aeroflight Dynamics Division for funding my graduate research and offering me employment upon completion of this degree. Secondly, I would like to thank my advisor, Juan Alonso, for all of the support and insight he has offered over the preceding years, as well as offering the freedom to work on a variety of interesting projects. Thirdly, I would like to thank my dissertation committee. Ilan Kroo, for the interest he has taken in my work, greatly strengthening the results presented here. Brian Cantwell, who has shaped my graduate career both in the classes I have taken from him and the warmth, encouragement, and knowledge he has given through all of our interactions. And Michael Saunders who has been an extremely kind and insightful presence in all of our conversations.

Fourthly, I would like to thank my fellow members of the Aerospace Design lab, both past and present for the valuable insight and collaboration they have offered. Particularly Trent Lukaczyk, Emilio Botero, Timothy MacDonald, and Andrew Wendorff, as well as the rest of the SUAVE team for their work developing and maintaining the code alongside myself. Additionally, I would like to thank Mike Colonno for writing the initial mission solver on which much of our work has built on. I would also like to thank the various mentors I had in my undergraduate career, especially Carl Knowlen. Working for him inspired my passion for research and strongly influenced my decision to pursue a PhD.

Additionally, I would like to thank my friends who have stood by me through the years, including Ben and Andy Bothwell (who was oddly insistent on being included in this section of the report), and who have in many ways been like a second family to me. I would also like to thank Kyle Howe, who has been a constant friend and source of encouragement throughout both my undergraduate and graduate careers. Furthermore, I would like to thank my many other friends who have been a tremendous blessing to myself and others, and for whom there is unfortunately, not nearly enough space to mention by name here.

Finally, I would like to thank my family. My brother who has stood by me for my entire life as well as my sister as someone I could always count on. I would like to particularly thank my parents both for the sacrifices they have made as well as for the example of selflessness, integrity, and work ethic that they have set for me. This work is dedicated to them.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Motivation

Commercial air travel demand has been rising dramatically in recent years. Alongside this rising demand has come wildly fluctuating fuel prices because of a combination of technological advances, changing government policies, as well as market shifts. For instance, fuel prices in 2012 doubled from 2003 values, at that point having risen to account for approximately 50% of airline operating costs for wide-body airplanes and to 30% for regional jets [12]. Furthermore, operation is expected to increase by 1.4 to 3 times by 2025 from 2004 levels [54]. A general rule of thumb is that worldwide aircraft travel doubles roughly every 15 years [2]. This increased demand will result in enormous additional greenhouse gas emissions in the absence of substantial reductions in overall fuel burn [1].

Traditional aircraft engines emit CO2 and water vapor along with smaller amounts of NOx, SOx, and soot. One unique challenge in aviation is that these emissions are released at higher altitudes (atmospheric composition changes with altitude). Therefore, they will have a different, often more pronounced impact than emissions released at sea level. NOx in particular is known to react with ozone, leading to increasing ozone at lower altitudes and decreasing ozone layer thickness at higher altitudes [76]. Ozone is peculiar in that, although it blocks ultraviolet radiation, it is itself a greenhouse gas which contributes to global warming. This increase in ozone concentration tends to be more prevalent at northern latitudes. SOx and water vapor are known to deplete ozone, so the release of these emissions may in some cases mitigate each other [76]. Water vapor is also a greenhouse gas, although it possesses less of a direct warming effect than CO2 or NOx. Contrails associated with the release of heated water vapor are known to form cirrus clouds, which also contribute to global climate change. These contrails tend to be stronger when aircraft are flown at higher altitudes, due to the more extreme temperature difference between engine exhaust and the ambient air. Such tradeoffs illustrate the complex, sometimes counterintuitive requirements associated with climate impact that are especially important to be mindful of when designing and

planning for new aircraft configurations.

Another major associated issue is noise. For conventional aircraft, propulsion noise is known to dominate over airframe noise during takeoff [32]. On the other hand, increasing engine bypass ratios over the last several decades have led to significant reductions in noise during takeoff and landing [6]. Electric propulsion systems are significantly quieter than current aircraft propulsion systems such as piston engines and turbomachinery. Therefore, widespread adoption could lead to significantly fewer active noise constraints for aircraft operation [69]. Increasing government regulations limiting these emissions are a significant driver in aircraft engine and materials improvements and their associated development costs. Additionally, in optimizing traditional tube-and-wing aircraft for reduced greenhouse gas emissions and noise requirements, there may be an associated penalty in operating cost [6]. In other words, an aircraft that is designed to reduce harmful greenhouse gas emissions and noise may actually burn more fuel than aircraft designed using more traditional methods.

Successful, widespread adoption of all-electric aircraft would eliminate direct greenhouse emissions from vehicle operation, making them as "clean" as the electrical grid used to charge the battery and eliminating much of the adverse climate impact from operating at higher altitudes. A major advantage from a environmental standpoint is that water vapor is not produced in the exhaust; the release of water vapor at higher altitudes is a significant driver in the environmental impact of commercial aviation. Furthermore, electric energy often costs a fraction of the equivalent amount of fossil fuel energy for a number of reasons, including increased conversion efficiency [69]. Hybrid-electric aircraft can also potentially offer substantially reduced fuel burn as well as more flexibility in aircraft operation. This can come by loading the vehicle with different combinations of fuel and batteries depending on the mission requirements [38]. Therefore there are powerful economic as well as environmental incentives for electrically-powered aircraft. On the other hand, the integration of electric propulsion systems into ground-based transportation poses fewer technical challenges than in aviation, as the weight penalty due to reduced specific energy is not nearly as restrictive there [44].

1.2 Battery Overview

Now, there has been considerable interest in the design of electrically-driven aircraft, using fuel cells and especially, batteries [90, 8, 44]. One of the primary advantages of batteries is a perceived lack of infrastructure required vis. a vis. fuel cells, in particular, liquid hydrogen fuel cells. However, present-day battery technology does not allow for the development of commercial-scale all-electric aircraft at significant ranges due to limitations in the specific energy, and in some cases specific power of lithium-ion batteries [88, 98].

One proposed solution to this energy density issue is lithium-air batteries. Lithium-air batteries

have projected pack-level specific energies ranging from 1,000-2,000 W-h/kg, and specific powers from .4-.67 kW/kg [51, 92]. One unique aspect of a lithium-air battery that must be taken into account in the design process is that, as the battery discharges, oxygen molecules collect on the cathode, which in turn, causes the battery to gain mass [17]. For larger batteries (such as those which may potentially be used for commercial aviation), this mass gain may be considerable. The chemical reaction using a non-aqueous electrolyte is of the form

$$2\mathrm{Li} + \mathrm{O_2} \Rightarrow \mathrm{Li_2O_2}. \tag{1.1}$$

Because the battery discharges based on the reaction of lithium with oxygen, any aircraft designed with these batteries would grow heavier throughout its mission (unless the battery uses stored onboard oxygen). Fortunately, this mass gain can be calculated using a relatively simple relation:

$$\dot{m} = \frac{MW_{\mathrm{O_2}}}{V_{0_{\mathrm{Li_2O_2}}} \cdot F} \cdot P. \tag{1.2}$$

Lithium-air batteries may also be designed using an aqueous electrolyte solution. In this case, they are also called lithium-water batteries. The chemical reaction is shown below.

$$2\mathrm{Li} + \frac{1}{2}\mathrm{O_2} + \mathrm{H_2O} \Rightarrow 2\mathrm{LiOH} \tag{1.3}$$

Some experts expect lithium-air batteries to be commercially available by 2030, although recent studies highlight significant design issues that may call this date into question [92, 59]. One reason for this is that the oxygen that reacts with the cathode lithium needs to be pure; even small amounts of moisture and carbon dioxide from the air could react with the lithium in the battery, rendering it useless [17]. Different strategies to deal with this issue include keeping pure compressed oxygen onboard the aircraft as well as including a filtering system to remove impurities. Additionally, lithium is highly reactive and needs to be handled with caution to avoid potential explosive incidents [68]. A key challenge in designing these batteries is the need to include excess lithium in the anode to compensate for losses during cycling.

Another type of metal-air battery that may be of interest in the design of electric aircraft uses an aluminum-air chemistry. Aluminum-air batteries boast impressive specific energy characteristics ($\sim$1,300 W-h/kg) and a high technology readiness level but poor specific power ($\sim$.2kW/kg). The given specific energy and specific power are representative of batteries produced in 2002 [102]. A prototype electric car powered by a combination of aluminum-air and lithium-ion batteries has been developed and tested by a company in Israel called Phinergy; commercial production was estimated to begin in 2017 [47]. The prototype car is pictured below [100].

Figure 1.1: Phinergy Prototype Aluminum-Air Car

Now, in addition to the Phinergy's automobile concept from Figure 1.1, there is a small nine passenger, 650 nautical mile commuter aircraft under development by Eviation, a company, interestingly enough, that is currently chaired by one of Phinergy's founders. A computer generated model can be seen in Figure 1.2 [81].

Figure 1.2: Eviation Alice Commuter Aircraft

It utilizes a similar energy system configuration to both Phinergy's automobile concept as well as one of the aircraft featured in Chapters 6, 8, and 9; to be more specific, a lithium-ion battery is used for general operation with an aluminum-air battery utilized as a range extender to enable significantly greater flight range. The aircraft is notable in that it possesses a battery weight fraction

of 65% with an all composite structure to compensate for specific energy limitations.

One drawback of aluminum-air batteries is difficulty in recharging the battery without a significant loss in capacity. A potential solution is the use of an AlCl3/EMImCl ionic liquid electrolyte, although this significantly lowers the practically achievable specific energy [82]. In addition, the AlCl3/EMImCl electrolyte may also be "prohibitively expensive" [61]. A proposed economic model involves replacing and recycling the aluminum plates after discharge instead, which entails an estimated cost of about $1.1/kg of aluminum anode [102]. Like lithium-air batteries, aluminum-air batteries accumulate oxygen while discharging. Aluminum-air batteries require water to operate, as seen by the aluminum-air discharge chemical reaction

$$4\mathrm{Al} + 3\mathrm{O_2} + 6\mathrm{H_2O} \Rightarrow 4\mathrm{Al(OH)_3}. \tag{1.4}$$

The mass of each constituent chemical (in this case, oxygen, water, and aluminum) can be calculated via relatively simple relations. These are shown for the mass flow of oxygen and overall required mass for water, respectively in (1.5) and (1.6).

$$\dot{m_{\mathrm{O_2}}} = \frac{MW_{\mathrm{O_2}}}{4V_{4\mathrm{AL(OH)_3}}} \cdot F \cdot P \tag{1.5}$$

$$m_{\mathrm{H_2O}} = \frac{MW_{\mathrm{H_2O}}}{2V_{4\mathrm{AL(OH)_3}}} \cdot F \cdot E \tag{1.6}$$

As an aside, the aluminum-water-oxygen reaction produces more energy than the given values for specific energy and specific power suggest; this is because undesirable byproduct accumulation in this reaction precludes the consumption of a fair portion of the aluminum reactant. Other advantages of metal-air batteries include a generally lower fabrication cost, flat discharge voltage, as well as a relative lack of temperature dependence [68]. A non metal-air approach to alleviating the specific energy issue is lithium-sulfur batteries. Lithium-sulfur batteries have a demonstrated specific energy of 500 W-h/kg, vastly exceeding that of lithium-ion. Furthermore, sulfur is cheap and abundant, which should, in theory, reduce manufacturing costs for these batteries. Lithium-sulfur batteries have already been used to power aircraft, particularly in high-altitude UAV designs, such as the Helios concept, pictured in Figure 1.3 [72].

Helios was a prototype UAV utilizing solar panels that set the altitude world record for level flight at 96,863 feet. The aircraft broke apart during flight in 2003 due to aeroelastic flutter as a result of its extremely high aspect ratio [72]. A number of practical challenges to their commercial development exist, however. For instance, because of the insulating nature of sulfur, conducting materials such as carbon must be added to the cathode to allow for ion transportation [64]. Perhaps the most significant challenge to their implementation and rechargability is volumetric expansion during discharge which occurs as LiS_2 forms [70]. Nonetheless, they are an area of extensive interest within the materials science community.

Figure 1.3: NASA Helios Aircraft

1.3 Conceptual Design

In the earliest stages of the aircraft design process, one needs to estimate the gross takeoff weight of the vehicle as the performance of the aircraft is directly determined by lift throughout a given mission. Conceptual design is inherently iterative due to the fact that many of the frameworks/systems of equations employed have no analytical solution [80]. Additional design requirements may increase the computational cost of these iterations, especially in the case of unusual aircraft configurations; traditional tube-and-wing designs rely heavily on the use of simple empirical correlations to estimate design performance, while more unusual designs with less data available require the direct use of expensive physics-based tools in an earlier stage of the design process to make design decisions [34]. This makes each aircraft iteration much more computationally expensive; the designer should therefore reduce the number of sizing iterations to as few as possible. Furthermore, for extremely unconventional cases, the designer may have little-to-no intuition as to what a "good" initial guess for the aircraft might look like which can lead to nonsensical results; a more informed starting point on the other hand might converge to a credible design. While the cost of introducing higher-fidelity analysis may be mitigated through the use of clever MDO architectures, the lower specific energy of battery systems has a cascading effect on the mass, and by extension, structural size of the aircraft [65, 88]. Properly evaluating the propulsion system is essential to the sizing process of futuristic electric aircraft, as even generous estimates of technology improvements still falls well short of what one may obtain with fossil fuels [69]. To illustrate some of the difficulties, taking the Breguet range equation

$$R = \frac{V}{g * SFC} * \frac{L}{D} * \ln(\frac{m_i}{m_f}), \tag{1.7}$$

and rearranging for fixed range

$$\frac{m_i}{m_f} = \exp(\frac{g * R * SFC}{V} * \frac{1}{\frac{L}{D}}) \tag{1.8}$$

$$= 1 + (\frac{m_i - m_f}{m_f}) \tag{1.9}$$

$$m_i - m_f = m_f * (\exp(\frac{g * R * SFC}{V} * \frac{1}{\frac{L}{D}}) - 1), \tag{1.10}$$

demonstrates that the required "fuel" mass ($m_i - m_f$) is proportional to the exponential of the effective specific fuel consumption of the aircraft (which specific energy is in turn inversely proportional to) minus a constant. Analogous equations also exist for electric aircraft but the adverse effect on total weight is more pronounced due to the fact that overall vehicle weight does not decrease during flight as it does with combustion-powered aircraft [44, 67]. To demonstrate, (1.7) can be expressed as

$$R = \eta_{propulsion} * Esp_{fuel} * \frac{L}{D} * \frac{1}{g} * \ln(\frac{1}{1 - \frac{m_{fuel}}{m}}). \tag{1.11}$$

This can be more easily compared to a fixed-mass form of the equation

$$R = \eta_{propulsion} * Esp_{bat} * \frac{L}{D} * \frac{1}{g} * \frac{m_{battery}}{m}. \tag{1.12}$$

A plot of these two curves for the same values of specific energy, velocity and $\frac{L}{D}$ are shown vs. fuel and battery mass in Figure 1.4 [44].

As Figure 1.4 shows, at higher "fuel" mass fractions, there is a significant penalty associated with flying an aircraft even without considering specific energy differences. Battery mass fractions of .45 or above are expected in the design of electric aircraft, which would yield a range penalty of 25% for the same specific energy. Furthermore, many proposed battery concepts increase in mass as they are discharged (see Section 1.2), compounding this effect. Even the most optimistic estimates for battery capacity fall well short of fossil fuel specific energy, as seen in Table 1.1, which lists some representative battery choices with jet fuel to compare.

Table 1.1: Energy Source Comparison

source	specific energy (W-h/kg) (approximate)	specific power (kW/kg)
jet fuel	12,000	
lithium-air battery	2,000	.66
lithium-sulfur battery	500	1
aluminum-air battery	1,300	.2
lithium-ion battery	200	1

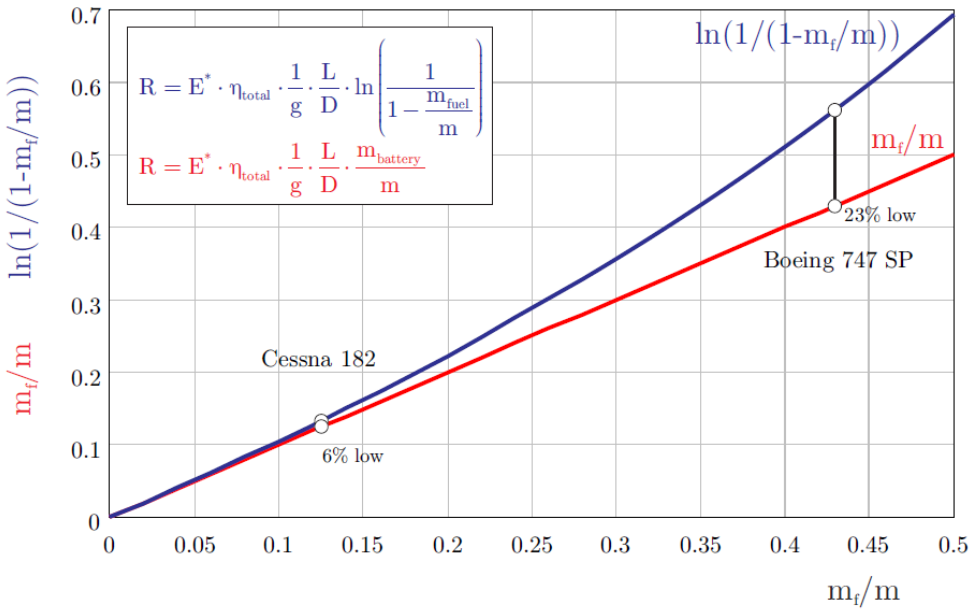

Figure 1.4: Range Equation Comparison

Table 1.1 lists the battery configurations from most optimistic (lithium-air, which is an area of active research) to least optimistic (lithium-ion, which is currently used in automobiles). Aluminum-air batteries are currently available with these power/energy characteristics, but have serious limitations such as low specific power and non-rechargability. Even the most optimistic battery candidate has 1/6th the specific energy of current propulsion systems, while current widely-available technology has a specific energy 1/60th that of jet fuel. While some of this performance gap can be made up for in increased propulsion system efficiency, configuration improvements, such as increasing L/D or reducing the structural weight are crucial to implementing these "green" aircraft with any significant amount of range. This may necessitate the use of higher fidelity methods. Furthermore, because this significantly increases the range factor, it also means that the overall aircraft weight will be more sensitive to design choices, making it more important to "nail down" the configuration early on in the design process. This means the introduction of higher fidelity analysis in the conceptual design stage could have a much bigger impact. Therefore, efficiently determining the required size of the vehicle is critical in the development of these futuristic aircraft concepts. One generalized problem

to be discussed in this thesis is as follows:

$$\min f(x) \tag{1.13}$$

$$\text{s.t. } g(x) > 0 \tag{1.14}$$

$$h(x, y) > 0 \tag{1.15}$$

$$y > 0, \tag{1.16}$$

where x is a set of design variables, f is the objective, and y is a set of terms called "sizing variables," which include parameters such as vehicle mass, battery energy, as well as propulsion system power. $h(x, y)$ is the vector of constraints associated with each y (and design parameters x). For a modern commercial aircraft, y could correspond to the gross takeoff weight of the vehicle while $h(x, y)$ could correspond to a fuel margin constraint. To give a concrete example of how this works, an iterative framework can be seen below.

$$W = C * y \tag{1.17}$$

$$W_{empty} = f(W, x) \tag{1.18}$$

$$W_{fuel} = F(x, y) \tag{1.19}$$

$$W_{out} = W_{empty} + W_{payload} + W_{fuel} \tag{1.20}$$

$$h(x, y) = \frac{W - W_{out}}{W} \tag{1.21}$$

where C is some constant to scale the problem such that it becomes numerically tractable. The weight breakdown is evaluated as a function of the design variables x and sizing variables y for a conventional gas-powered aircraft, as seen in (1.18). The mission is then flown, and the fuel weight required to fly the mission is calculated as a function of x and y, as seen in (1.19). Required weight is then calculated as seen in (1.20), for a fixed payload. Now, other valid formulations will be investigated that may be more computationally efficient in a given situation; an example is an optimizer decomposition approach, where an optimizer enforces constraints rather than solves an inner loop for each design decision. In this case, when the gradient of $h(x, y)$ is computed, such as in a optimization framework, it is more numerically consistent when normalized by a constant instead of W, i.e.

$$h(y) = \frac{W - W_{out}}{C}. \tag{1.22}$$

Differences between these formulations are elaborated upon in Chapter 3. Electric aircraft, which are more pertinent to the subject of this dissertation, would include requirements such that the battery state of charge remains positive (or has some chosen margin), and the power requirements of the mission do not exceed the maximum power output of the battery. Note that in the context

of this work, a "sizing evaluation" refers to when the vehicle's geometry and weight is defined and runs the mission based on the inputs (x, y). The sizing evaluation is chosen as a measure of computational efficiency as, in the extreme case, high fidelity methods would be run to estimate weight and aerodynamic performance for each individual aircraft. This work will investigate relative the computational cost and robustness of a number of aircraft-sizing schemes (Chapter 3) applied to four different aircraft classes (Chapter 6).

1.4 Optimization

1.4.1 Gradient-Based Methods

Mathematical optimization, in general, is a methodology used to decrease some objective function (such as fuel burn or operating cost), usually subject to some constraints under a certain set of rules which should (if all goes well) lead to at least a local minimum, assuming the problem is properly bounded [10]. Before beginning an optimization case, the user needs to make a decision as to which rules to apply. These algorithms are inherently iterative, requiring an initial guess, affecting the speed of the optimization process and possibly the location of computed optima [71]. There is often a trade off between speed and robustness that one must take into account when choosing an optimization algorithm. Gradient-based optimization, which uses local derivative information to determine a search direction is a popular choice, especially when computational resources are limited. A major shortcoming is a tendency to converge to local, rather than global minima for nonconvex problems. The first recorded use of a gradient-based method for optimization can be traced back to Cauchy in 1847 [20, 10]. Gradient (or steepest) descent is a popular initial choice in introducing the concept of optimization because the gradient points in the direction of maximum decrease in the objective. As the name implies, the gradient itself is chosen as the search direction. It works based on some initial x_i, updating it as

$$x_{i+1} = x_i - \alpha * \nabla f(x), \tag{1.23}$$

where α is a step length. Often times, this algorithm incorporates a line search to determine α in order to ensure sufficient decrease in the objective. $\nabla f(x)$ can be calculated in a variety of ways; for complex, iterative systems, finite differencing to approximate the gradient is a general, robust, albeit computationally slow choice. Gradient descent, while intuitive, can be computationally slow on even simplistic problems, and has a tendency to exhibit oscillatory behavior as it approaches a solution. This is because, for an exact line search, each search direction is orthogonal to the previous one, slowing convergence. Other methods, such as conjugate gradient and Newton methods use additional information to determine a step direction. Newton methods for example use second order

information to determine a step direction, i.e.

$$x_{i+1} = x_i - \alpha * H^{-1} \nabla f(x), \tag{1.24}$$

where H is the Hessian. Note that Newton's method approximates $f(x)$ as a quadratic, and then uses the exact solution of that quadratic function to update the design variable vector x. For a quadratic function $f(x)$, this means that the method will converge in one major iteration. Now, the aforementioned methods are for unconstrained problems; for constrained problems, one of the more popular approaches is to define a Lagrangian to minimize, i.e. there is an objective $f(x)$ to be minimized, subject to the set of constraints

$$g(x) > 0 \tag{1.25}$$

$$\hat{g(x)} = 0. \tag{1.26}$$

With this set of constraints $g(x)$,the Lagrangian can be defined as

$$L(x, \lambda, \hat{\lambda}, s) = f(x) - \hat{\lambda^T} \hat{g(x)} - \lambda^T (g(x) - s^2), \tag{1.27}$$

where s is a set of slack variables chosen to transform the set of inequality constraints $g(x)$ into a set of equality constraints and λ and $\hat{\lambda}$ are sets of Lagrange multipliers. Slack variables are only active when s is greater than zero; an s value less than zero signifies an infeasible point, while an s = 0 signifies an active constraint at a point x. The Lagrangian has the property that, at a critical point, the gradient of L with respect to x, λ, $\hat{\lambda}$, and s is 0. This is part of what are known as the first order KKT conditions, which are necessary (but not sufficient) to show optimality [71]. Additional first order KKT conditions include that $\hat{g(x)}$ and $g(x) - s^2$ must also equal zero to ensure that the constraints are satisfied, along with the requirement that the Lagrange multipliers must all be greater than or equal to zero. The second order condition requires that the Hessian of the Lagrangian be positive definite to ensure a feasible result.

Now, a class of optimizer that will be extensively used in this dissertation is a Sequential Quadratic Program (SQP). SQP methods determine the active set and are based on an application of Newton's method to the Lagrangian [71]. The model the problem as a quadratic at a point x, and solve the problem to find the solution to that quadratic, yielding a new point. To be more specific, start with a matrix $A(x)$ defined as the Jacobian Matrix of the gradients of all of the equality constraints as well as transformed inequality constraints, collectively defined as $g(x)$.

$$A(x) = \left[\nabla g(x)\right] \tag{1.28}$$

Begin with defining a quadratic program

$$\min_p f(x) + \nabla f(x)^T p + \frac{1}{2} p^T \nabla^2_{xx} L(x, \lambda) p \tag{1.29}$$

$$s.t.\ A(x)p + g(x) = 0. \tag{1.30}$$

Now, the set of first-order KKT conditions can be defined as

$$F(x, \lambda) = \begin{bmatrix} \nabla f(x) - A(x)^T \lambda \\ g(x) \end{bmatrix} = 0. \tag{1.31}$$

Finding the Jacobian of this matrix with respect to x and λ in turn forms a set of nonlinear equations that can be set to zero and solved using Newton's method (also known as Newton-Raphson).

$$F(x, \lambda)' = \begin{bmatrix} \nabla^2_{xx} L(x, \lambda) & A(x)^T \\ A(x) & 0 \end{bmatrix} = 0 \tag{1.32}$$

One can then set these terms equal to each other, forming a system of equations to determine a step p, i.e.

$$\begin{bmatrix} \nabla^2_{xx} L(x, \lambda) & A(x)^T \\ A(x) & 0 \end{bmatrix} \begin{bmatrix} p_x \\ p_\lambda \end{bmatrix} = \begin{bmatrix} -\nabla f(x) + A(x)^T \lambda \\ -g(x) \end{bmatrix}. \tag{1.33}$$

Now, assuming the LHS matrix in (1.33) is nonsingular, which happens when both $A(x)$ is full rank as well as when the matrix $\nabla^2_{xx} L(x, \lambda)$ is positive definite along the constraint lines [71], the step direction can be easily determined. This is equivalent to solving the problem shown in (1.29) and (1.30) [71].

1.4.2 Global Optimization

In some cases, gradient-based optimization may not be sufficient to properly evaluate the design space and make informed decisions. This can occur when, for instance, an abundance of local minima exist, meaning that different initial starting points for a gradient-based optimizer would result in wildly different optima with substantially differing performance gains. In these cases, one must find other intelligent ways to search the design space. Often, these are based on stochastic, population-based methods, which seed a bounded design space with an initial population, compare them, and based on their relative fitness somehow exchange information among each of the candidates [101]. One relatively intuitive example of a population-based global optimizer is a Particle Swarm algorithm, which, was originally developed as part of an attempt to simulate flocks of birds [55]. Now, a Particle Swarm Optimizer (PSO) works by seeding some "space" with a number of candidate designs, distributed throughout some bounds assigned to the design variables, often with some randomly assigned "velocity." Based on the fitness of each candidate "particle" along with its

current velocity, the trajectory of each particle is altered based on its position relative to both the best known point as well as the "fittest" current particle, causing the set of particles to sweep the design space. An illustration of this algorithm in action for an unconstrained 2D problem can be seen in Figures 1.5–1.8.

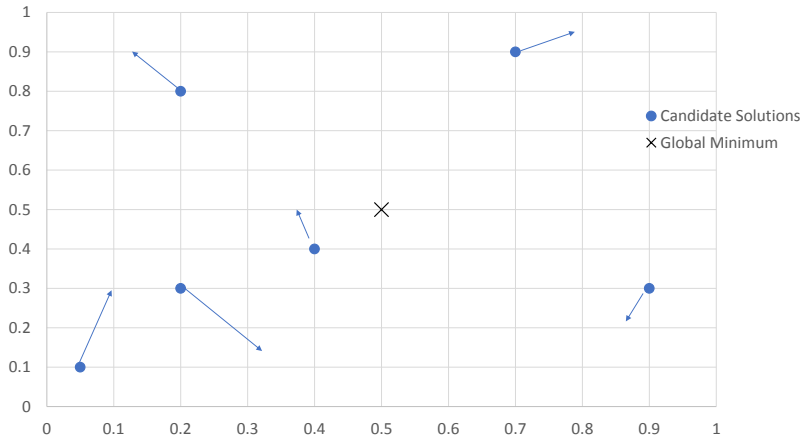

Figure 1.5: Particle Swarm Optimization (Initial Seed)

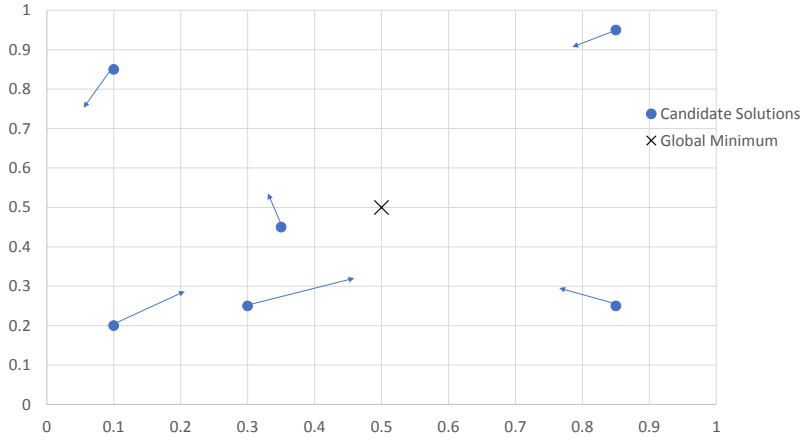

Figure 1.6: Particle Swarm Optimization (Iteration 1)

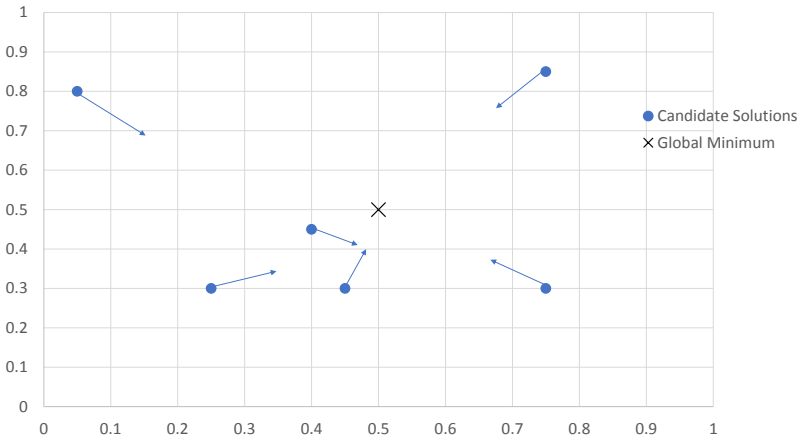

Figure 1.7: Particle Swarm Optimization (Iteration 2)

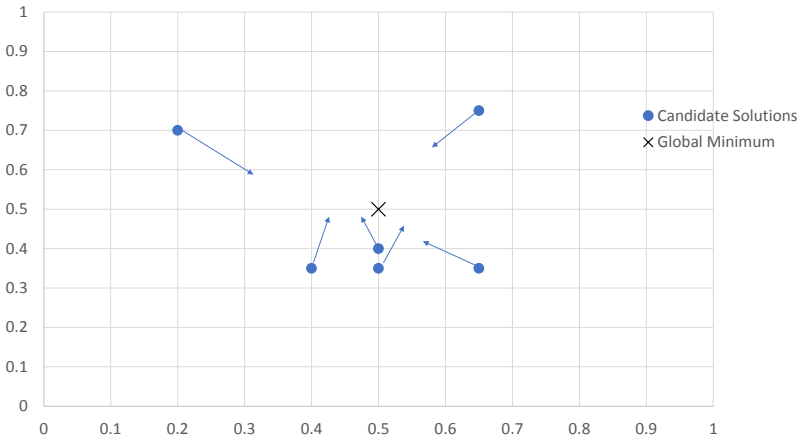

Figure 1.8: Particle Swarm Optimization (Iteration 3)

The various candidate designs start out distributed randomly around the design space, but as the algorithm proceeds, they begin to cluster around the global minimum. As a heuristic algorithm there is no guarantee of convergence, but in practice, it is less prone to becoming stuck in local minima than quasi-Newton methods, albeit at a substantially increased computational cost. On the other hand, this algorithm has the benefit that no gradients are calculated, removing the need for computationally expensive finite difference steps. In general, this is not nearly enough to achieve speed parity

with quasi-Newton methods. Now, a PSO algorithm may also be altered to handle constrained optimization problems. The Augmented Lagrangian Particle Swarm Algorithm (ALPSO) is one such example of this [50]. it uses Lagrange multipliers to model the constraints, i.e. the constrained optimization problem is formulated similar to (1.27). Now, most global optimizers do not use a disciplined handling of constraints, such as that shown in (1.25)-(1.33), vastly reducing their speed and robustness for systems with complex constraint tradeoffs compared to Quasi-Newton methods [50, 85]. Oftentimes, the stochastically generated designs all violate at least some of the imposed constraints; a variety of techniques have been proposed to generate feasible designs such as penalty functions as well as "repair" algorithms [101]. Repair algorithms use heuristic techniques to make these candidate designs feasible after the initial sampling. An example of a repair algorithm is a so-called hybrid algorithm, which uses local-search techniques to move each sample towards a feasible point [85]. These repair techniques, while tending to be more robust than other constraint-handling options for global algorithms, also tend to be somewhat slower [40]. Another proposed global search strategy includes "multi-start" approaches, where local optimization searches are performed using a set of randomly generated points, where the optimization algorithm returns the "best" point obtained [43].

Other global search strategies make use of surrogate representations of the design space. One pseudo-global example of a surrogate-based optimization method is Trust Region Model Management (TRMM) [83]. TRMM separates the design space between low and high fidelity evaluations, using corrections to ensure that the objective, constraints, and their gradients match at the point of interest. It optimizes on the corrected low fidelity problem within a constrained portion of the design space (the trust region), then samples the high fidelity model at the corrected low fidelity optimum. Here, the high fidelity evaluations could be candidate aircraft from an design code, while the low-fidelity evaluates could be a surrogate representation of them. One "nice" characteristic of Trust Region Model Management is that it is known to converge to either a stationary point or local optimum from a given initial guess [3]. Now, in the context of this dissertation, due to the low specific energy and specific power of the propulsion systems outlined here, large infeasible regions of the design space exist. Surrogate-based optimization is known to struggle with problems that possess this characteristic, so it will not be investigated extensively here [29]. Many other global search strategies exist, and the algorithms discussed here are by no means exhaustive.

1.5 Literature Survey

Significant work has been published in the conceptual design of electric and hybrid-electric aircraft. Datta and Johnson, for instance, investigated the design of fuel cell and battery powered helicopters, taking into account discharge losses in the batteries, focusing specifically on how the design of the propulsion system affects system-level performance for a fixed airframe, while detailing areas of

potential improvement such as the drive system [26]. Hepperle qualitatively compared a variety of energy systems, including fuel cells and conventional turbomachinery and outlined some modified range equations for constant mass aircraft, showing some results for constant efficiency and an assumed L/D [44]. Boeing created an extensive report on potential N+3 future passenger aircraft called Subsonic Ultra Green Aircraft Research (SUGAR), wherein Phase I highlighted the SUGAR Volt, which incorporated electric propulsion capabilities [14]. Phase II looked at N+4 technologies, with more analysis into the electrification of passenger aircraft [15].

Figure 1.9: Boeing Sugar Volt

Stckl, van Toor, and Lobentanzer performed trade studies on the range and total weight of lithium-air batery powered aircraft for varying number of passengers, comparing them to a range of turboprop aircraft [92]. More recently, Duffy, Wakayama, and Hupp explored the sizing of three different all-electric aircraft configurations (a helicopter, stopped rotor, as well as tilt rotor design), comparing their overall operating cost for 100 nautical mile VTOL missions vs. a conventional internal combustion engine helicopter design assuming a battery specific energy of 400 W-h/kg [28]. Perhaps most relevant to the work here, Johnson, Silva, and Solis performed fundamental sizing and trade studies of some of these electric and hybrid-electric aerial taxi aircraft, taking into account a model that estimates the tradeoffs between specific energy and discharge rate [53]. The vast majority of these proposed designs require a significant improvement in battery specific energy over present-day lithium-ion batteries to allow for meaningful range/convergence of their sizing algorithms.

One aircraft design that is currently under development is the Kitty Hawk Flyer. Kitty Hawk is a startup company funded by Google's Larry Page; their unusual aircraft concept can essentially be described as a multicopter with a chair on top. The aircraft is FAR 103 compliant, and a prototype

has been built and is in testing as of this writing. A picture of the aircraft can be seen in Figure 1.10 [94].

Figure 1.10: Kitty Hawk Flyer

One topic of special interest is distributed electric propulsion, which is a somewhat radical departure from traditional concepts. Distributed electric propulsion works via the use of propellers positioned along the length of the wing, powered by electric motors; it is enabled because electric systems tend to be scale invariant, whereas combustion systems are more efficient at larger scales. Potential benefits include increased dynamic pressure over the wing, reducing the required wing size to a more cruise-optimal shape. Such designs allow for redundancy, in that several motors/propellers can fail while allowing the aircraft to continue functioning. Finally, in typical helicopters, the complicated rotor system with cyclic, collective, and a swashplate may potentially be made lighter and simpler by being replaced with vectored thrust and smaller electric components [93]. Proponents of these configurations argue that integration of electric technology in this way can allow all-electric design, in some instances, to outperform conventional aircraft without radical battery technology improvement assumptions [69]. Green, Schiltgen, and Gibson developed a conceptual design tool that specifically focuses on components and characteristics of hybrid turboelectric concepts that make use of distributed propulsion, but argue that in order to truly capture the benefits of such a concept, one needs to evaluate propulsion-airframe interaction effects [39]. A number of electric aircraft are currently in development, which would fall under the category of personal air vehicles. These include the Joby S2, which uses propellers distributed along the wing for vertical takeoff and landing while stowing them during cruise flight (to reduce drag) [91]. The wing-tip mounted propellers are used during forward flight. An artist's rendering of the aircraft can be seen in Figure 1.11.

Figure 1.11: Joby S2

Other, similar aircraft in development include the Lilium Jet as well as Zee.Aero's personal air vehicle concept. Lilium is a company based out of Germany; their 5-seater tilt-wing concept uses 36 in-wing ducted fans which can rotate to allow for vertical takeoff and landing [93]. The aircraft itself possesses a canard, with ducted fans along both the canard and rear-wing. It boasts a top speed of 300 km/hour with a range of 300 km. A prototype 2-seater design known as the Lilium Eagle has been built and flown, and is in the midst of testing as of the time of the this writing. An artist's rendering of Lilium's design can be seen in Figure 1.12.

The aforementioned distributed electric propulsion concepts are part of a broader push towards electric personal air transportation (or eVTOLs), which reached a significant milestone as part of the Uber Elevate summit. A white paper associated with this summit posits that significant reductions in traffic congestion can be achieved via on-demand air transportation due to a convergence of a variety of developing technologies including artificial intelligence, battery technology, as well as material science innovations [46].

Hybrid-electric regional transports are another aircraft configuration of interest. For instance, Pornet et. all investigated the design of a hybrid-electric transport jet, where a long-range transport was retrofitted with batteries and flown with a shorter range, and the batteries operated in cruise, while allowing for longer range missions that utilize the combustion engine [78]. One goal of these designs such as this is to leverage the increased efficiency of electric systems vis a vis combustion systems to reduce operating costs as well as emissions. Another important aspect of these designs is the ability to operate electrically for short ranges, allowing for more flexibility in reduced cost/emissions depending on the choice of vehicle operation [38].

Figure 1.12: Lilium 5 Seater Concept

Marwa et. all produced some interesting results that show, analytically, that under a specific set of conditions (that is, when the hybrid system is disjoint, i.e., that only one system runs at a time), the range equation is binary, i.e. for fixed parameters, if there is a minima in range, it is between the extremes where the aircraft is powered by a single fuel source [66]. In other words, under these conditions, neglecting other aspects of the propulsion system, and assuming constant efficiencies, aircraft range vs. "electrification factor" is convex. As with many of these studies, real benefits only begin to appear at battery specific energies well beyond current lithium-ion technology.

Now, as mentioned previously, battery specific energy is an enormous limiting factor in the conceptual design of these systems. Specific power may also be of concern, although many studies do not take this into account. For example, reference [4] looks into some detail into the operation of hybrid-electric aircraft at a wide range of electrification choices, but neglects to mention any potential investigation of power-limited batteries (correspondence with the author confirms that this aspect of the design was ignored). Fredericks makes mention of sizing diesel-electric aircraft using an assumed battery specific energy and specific power without taking into account tradeoffs based on battery design and discharge choices [34]. One notable exception is Isikveren et. al., who, for a simplified case derive analytical expressions for the sizing/weight buildup of an aircraft that take into account both energy and power requirements for aircraft using more than one energy system [49].

Furthermore, in the context of designing these aircraft, nearly all approaches focus on simple trade studies, with little to no discussion of the optimization process. Moore and Fredericks argue that gradient-based approaches to these concepts are insufficient to properly explore the design space of a large number of configuration choices [69]. To give a concrete example, the choices

in size, number, and especially position of the various propellers/fans create a highly multimodal design space with a wide range of local minima. Patterson and German develop a multidisciplinary analysis environment focused on distributed electric propulsion to handle propeller-wing interactions, although they do not include any optimization results themselves [74]. On the other hand, Varesi and Radan investigate the multiobjective optimization of automobiles with a single parameter, called "degree of hybridization (DoH)," using a genetic algorithm [96]. DoH is defined as the ratio of maximum electric power output to the total power output of the vehicle. Their results showed that several choices of objective (e.g. gradeability, acceleration times) were highly nonconvex vs. this parameter and exhibited multiple local maxima and minima. This phenomena is expected to arise in many of these hybrid-electric aircraft designs.

1.6 Contributions

Original contributions of this dissertation can broadly be categorized into two "boxes." The first is discovery of previously undocumented issues that have a tendency to arise in aircraft design problems of this nature, i.e. that employ multiple energy systems. Their exact causes will be explored in some detail. The second is ways to solve, or at least, mitigate these issues, which, fortunately can significantly increase code speed as well.

One major issue is the tendency of these systems to possess a multiplicity of sizing solutions depending on the algorithm used to solve the problem as well as the starting point for the iteration scheme. It was found that certain (typically faster) solution methods exhibit this tendency more than others, and that delaying the usage of these algorithms until the system is moderately more converged substantially decreases the likelihood that errant solutions are found. The second major problem is that design problems for certain classes of aircraft, particularly ones with a propulsion framework such as a series-hybrid design, are unsolvable using some commonly utilized algorithms (such as fixed point iteration). A solution here is an intelligent backtracking algorithm, which enables the problem to be robustly solved using a variety of solution algorithms.

Because the cases feature higher dimensionality with a complex interplay between the energy systems, convergence becomes significantly slower. To that end, several different solution algorithms were explored to increase code speed under a variety of conditions. Furthermore, a framework for using machine learning algorithms to "learn" the solution space and start each iteration process from a more converged result was also developed. This dissertation also reveals, that, particularly for hybrid-electric and hybrid battery designs, correct modeling of battery specific power limitations is critical in estimating required aircraft weight; using only a single nominal value for specific energy and power resulted in designs that were sometimes twice as heavy as designs that look at the full potential range of battery operation. All algorithms and frameworks were implemented within an open-source aircraft design software to allow access for the larger aerospace community.

1.7 Dissertation Overview

Chapter 2 discusses many of the physical benefits of using multiple energy systems along with some of the basic workings and synergies that can occur when these systems are employed in tandem. Chapter 3 introduces the frameworks employed in this dissertation (and commonly, elsewhere) that are used to solve fundamental aircraft sizing and design problems, in the form of a broadly applicable mathematical system. Chapter 4 describes some of the basic choices in terms of numerical methods that are used to converge systems throughout this dissertation and discusses some of their strengths and weaknesses. Chapter 5 analyzes and progressively simplifies several classes of analytical "toy" aircraft design problems in order to diagnose some of the pathological issues that sometimes arise in these systems, which are shown for more realistic aircraft designs in later chapters. In addition, it compares a number of methods outlined in this dissertation in their relative performance towards a representative global optimization problem. Chapter 6 details four classes of aircraft to be analyzed in this report, linking their precise formulation to the systems outlined in Chapter 3. Chapter 7 explores the application of machine learning to an "in the loop" optimization problem using the methodologies outlined in previous sections of this paper and discusses the strengths and weaknesses of variety of candidate regression algorithms. Chapter 8 explores the application of these methods to two major classes of aircraft for 2-dimensional sweeps of an aircraft's design space, illustrating the pathological issues that can arise as well as solutions to these issues. It also compares the computational benefits of some of the methodologies outlined here towards problems of this class. Finally, Chapter 9 applies the various methodologies outlined here to four separate aircraft classes in gradient-based optimization frameworks along with a discussion of the results.

Chapter 2

Design Considerations

2.1 Specific Energy vs. Specific Power

In order to understand some of the important tradeoffs one must make when designing more electric, or all electric aircraft, the designer should take into account both the power and energy requirements of the energy system. This is often plotted in terms of what is called a "Ragone plot." A Ragone plot characterizes the relationship between energy density and power density of a device [22]. To illustrate, for a battery, a large part of the effect can be seen in the relationship between voltage and power, where

$$P = IV. \tag{2.1}$$

To see this more explicitly, a plot of voltage vs current density and power density for an aluminum-air battery is shown in Figure 2.1 [102].

The actual ability of a battery to discharge is measured in capacity (in amp-hours). Because one can see a roughly linearly decreasing trend of voltage with increasing power as well as current, and because energy is just the integrated power vs. time

$$E = \int_0^t P dt \tag{2.2}$$

$$E = \int_0^t IV dt. \tag{2.3}$$

Now, battery capacity is typically defined as

$$C = I * t. \tag{2.4}$$

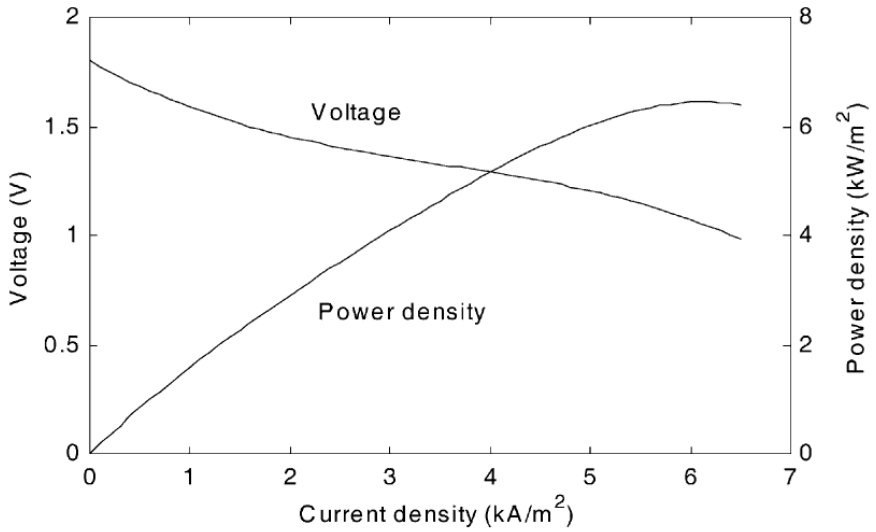

Figure 2.1: Aluminum-Air Cell Discharge Curve

Therefore, for a given battery, there is an inverse relationship between energy and power merely based on how it is discharged (assuming roughly constant voltage). Now, in actuality, this relationship is somewhat more complicated, as the voltage of a battery drops based on its depth of discharge and the current drawn; however, qualitatively, this relationship holds [26]. One way of taking this into account is captured by Peukart's Law, which models the change in capacity from varying discharge rates, and was originally developed in the late 1800s [73]

$$C = I^k * t, \tag{2.5}$$

where k is the Peukart coefficient (an empirical value determined from testing) and t is the time to discharge the battery in hours. In addition, one may design a specific battery chemistry to target a particular range of specific energy vs. specific power. Most electric aircraft conceptual design papers assume a constant specific energy and specific power (if specific power is modeled at all). A number of different frameworks exist for modeling this tradeoff, ranging from purely empirical curves to more complex, physics-based models which attempt to evaluate the open-circuit voltage vs. both current level as well as battery state of charge. As mentioned above, as batteries discharge, the open circuit voltage is known to drop off as seen in Figure 2.2, which depicts an aluminum-air battery voltage vs. discharge time for some experimental data [21].

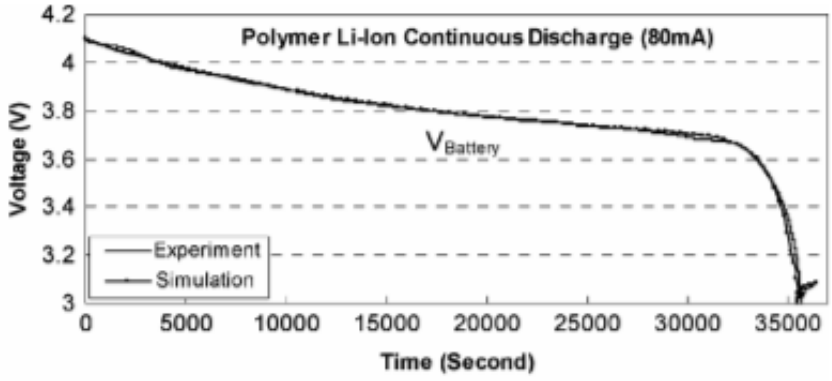

Figure 2.2: Battery Discharge Profile

This effect results in increased discharge losses in the battery at constant power, as, with the voltage drop, current must be increased to maintain power output. Many battery models attempt to account for these losses [21, 26]. However, due to the asymptotic drop-off in voltage as shown in Figure 2.2, discharge losses blow up as the battery state of charge approaches 0. This can make these models unsuitable for incorporation into an optimization process, as these discharge losses can result in unphysical discontinuities which make the problems highly nonsmooth, especially when the state of charge becomes negative. A Ragone plot for the battery constructed using this curve can be seen in Figure 2.3.

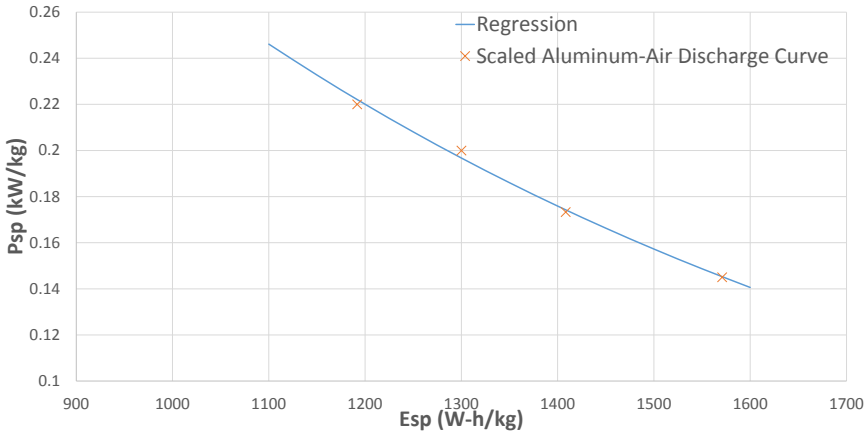

Figure 2.3: Aluminum Air Ragone Plot

Batteries are often parametrized in terms of maximum C rate, which is the rate the battery is

discharged in 1/hours. In other words, 1C means the battery is discharged in one hour, while .5C means the battery is discharged in two hours. To illustrate, the Ragone plot shown in Figure 2.3 has a range of .26C at 1,100 W-h/kg, to .15C at 1,300 W-h/kg.

2.2 Why Multiple Energy Systems?

In some cases, an aircraft design may not close when using a single energy system. Physically, this may be a result of poor specific power, poor specific energy, or both. Looking forward, several batteries are in development that possess extremely high specific energy but poor specific power, including the aluminum-air battery discussed in the previous section [102, 19]. For example, the Ragone plot shown in Figure 2.3 is plotted alongside the Ragone plot for what one might call a near-term lithium battery with corresponding C-rates superimposed.

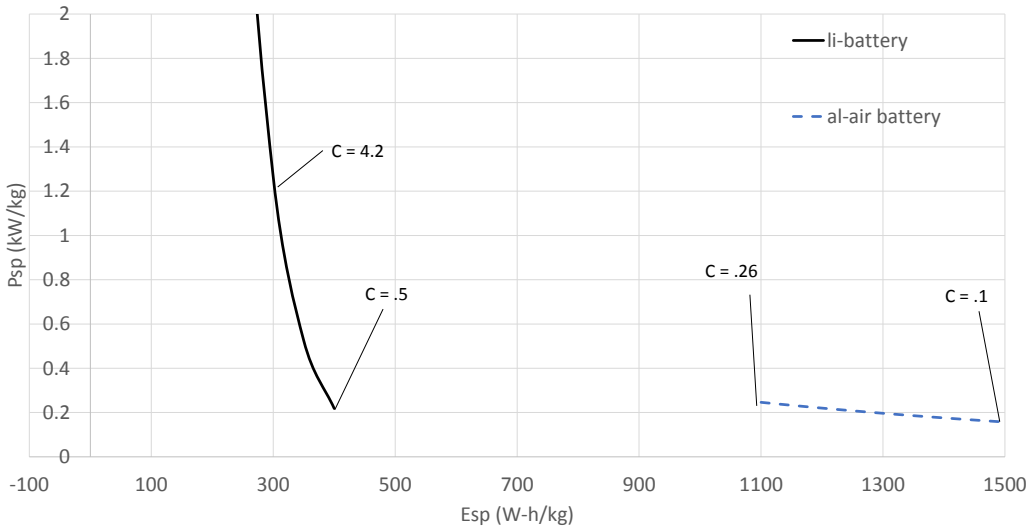

Figure 2.4: Lithium-Sulfur and Aluminum-Air Ragone Plots

As Figure 2.4 shows, when designing an aircraft that combines multiple energy systems multiple orders of magnitude may exist in terms of both specific energy as well as specific power. This presents problems from physical as well as numerical standpoints, making the aircraft design process potentially much more difficult. Furthermore, any battery chemistry in development will fall well short of the specific energy of fossil fuels ($Esp \approx 12{,}000\ \frac{W-h}{kg}$) as a result of theoretical energy density limits. In this case, one may be able to open up new configurations by tailoring the composition

of the energy system to the needs of the mission. To give a concrete example, in rotorcraft, the power required to hover can be 2-3 times the power required to cruise. A hypothetical, qualitative profile of vehicle power requirements vs. velocity can be seen in Figure 2.5, alongside an example of a mission profile in which an aircraft could benefit from the use of multiple energy systems.

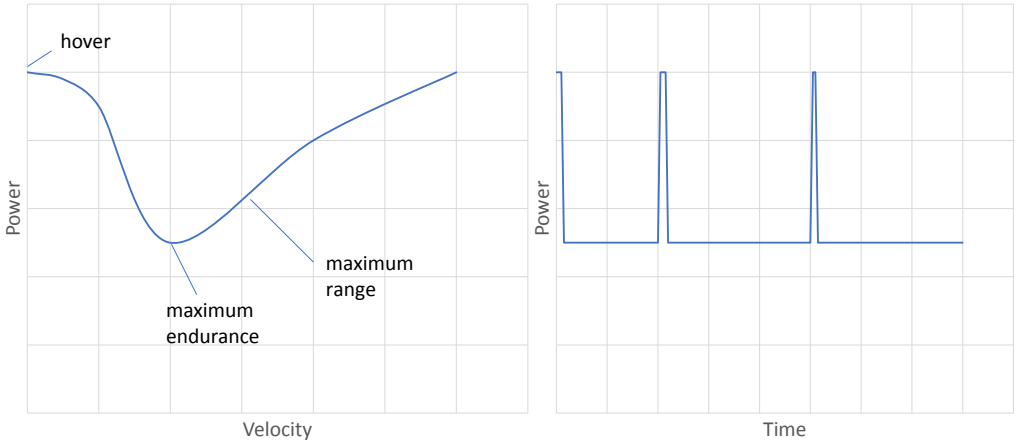

Figure 2.5: "Ideal" Hybrid Power Profile

Many designs are only required to hover for takeoff and landing for short periods of less than five minutes, so the profile shown in Figure 2.5, while somewhat extreme-looking may be representative of some aircraft. One could choose a battery system such that a high specific power battery provides the extra power needed in hover, while the high specific energy battery operates only in cruise. Hence, overall system weight can be dramatically reduced by not carrying around a large, power-limited energy system.

Now, in designing a vehicle with multiple propulsion systems, one must first determine some way to parametrize the relationship between the two systems. When looking at hybrid-electric cars as well as aircraft, a variety of terms in addition to parametrization choices exist, such as what is variously called a "hybridization factor," "electrification factor", "degree of hybridization (DoH)," or simply "electrification" and is shown in (2.6) [18, 5].

$$DoH = \frac{P_{elec}}{P_{elec} + P_{comb}} \tag{2.6}$$

Confusingly, the term "hybridization" factor is also used to refer to the ratio of electric to internal

combustion engine power as in (2.7).

$$DoH = \frac{P_{elec}}{P_{comb}} \tag{2.7}$$

Thus, it appears that terminology has not yet been completely standardized when it comes to hybrid vehicle design, even for such a fundamental parameter.

2.3 Reasons for a Hybrid-Electric Aircraft

Many of the most important reasons for hybrid-electric automobiles do not necessarily translate directly when applied to electric aircraft. For example, a start-stop system to prevent excessive idling in the combustion engine during traffic or long stoplight waits does not apply nearly as strongly to aircraft, where the idling time tends to make up a much shorter fraction of the flight time [56]. There are a number of reasons to develop hybrid-electric aircraft that, depending on technology improvements and design choices, may or may not result in reduced fuel burn. One of the most important factors in the design of hybrid-electric aircraft is specific power of the propulsion system subcomponents; the tradeoffs are discussed in Section 2.2 above.

One promising candidate system to integrate with a battery is the diesel engine. Diesel engines are known to offer substantial improvements in specific fuel consumption over turbine engines, albeit while suffering significant power-to-weight ratio penalties [25]. To illustrate, a plot of some scaling regressions for the weight of diesel engines, turboshaft engines, as well as high power-to-weight electric motors can be seen in Figure 2.6 [99].

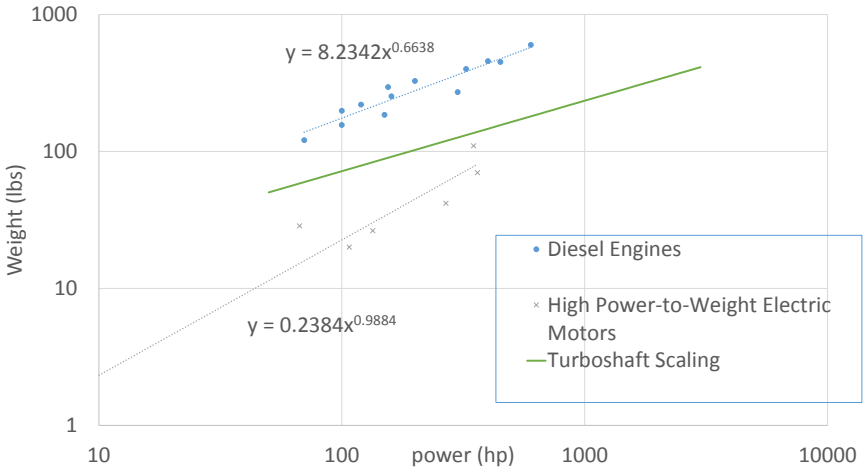

Figure 2.6: Power System Weight Regressions

For this reason, a number of new vehicle concepts are in development which implement a diesel-electric propulsion system that uses batteries and motors to account for short duration power requirements (e.g. hover), while taking advantage of the high efficiency of diesel engines [34]. Additionally, a given diesel engine (if turbocharged) tends to lapse at a higher altitude than a comparable turboshaft, offering more operational flexibility [25]. As a result, aircraft designed using diesel engines, combined with an electric system, will likely not suffer the same off-design operation penalties from hot and high hover requirements. Taken together, these characteristics could open up new potential design concepts.

One important aspect in the design of diesel aircraft involves properly modeling the turbocharging operation. To that end, a method based on a 0D model was developed and used to regress some publicly available data, allowing the user to represent the effects of differing atmospheric effects in the performance of the engine. To illustrate, using the energy equation, the incoming temperature is raised based on a term called $Spa_{turbocharger}$, as shown in (2.8).

$$T_1 = T_0 + \frac{Spa_{turbocharger}}{cp} \tag{2.8}$$

Then, assuming an isentropic process and a perfect gas

$$\frac{p_1}{p_0} = \frac{T_1}{T_0}^{\frac{\gamma}{\gamma-1}} \tag{2.9}$$

$$\rho_1 = \frac{p_1}{RT_1}, \tag{2.10}$$

the gas is then compressed based on a fixed compression ratio r_c

$$\rho_2 = \rho_1 * r_c, \tag{2.11}$$

with isentropic compression assumed to determine the temperature and pressure. T_3 was then determined based on the lower heating value of the fuel and a fuel-to-air ratio that was iterated upon

$$T_3 = T_2 + f * \frac{Q_{lhv}}{cp}. \tag{2.12}$$

No changes in gas composition properties (such as specific heat) were modeled in this simplified analysis, and a perfect gas was assumed throughout. Furthermore, if T_3 was found to be greater than the adiabatic flame temperature, the adiabatic flame temperature was assumed, and a fallout fuel-air ratio calculated from the above relation. The exhaust stroke then returns the charge to the

original volume, and with isentropic expansion is assumed, one obtains

$$\rho_4 = \rho_3 * (1 + f) \tag{2.13}$$

$$\frac{p_4}{p_3} = \frac{\rho_4}{\rho_3}^{\gamma}. \tag{2.14}$$

Turbocharger exhaust conditions were then determined based on an energy balance.

$$T_5 = T_4 - \frac{Spa_{turbocharger}}{cp(1 + f)} \tag{2.15}$$

With the gas conditions defined for all parts of the cycle, the indicated fuel conversion efficiency, $\eta_{f,i}$ as well as indicated mean effective pressure, $imep$ (the work per cycle per chamber volume) can be determined based on the enthalpy at each of the cycle stages [45].

$$\eta_{f,i} = \frac{W_{c,i}}{m_f * Q_{lhv}} = \frac{(h_4 - h_3) - (h_2 - h_1) + \frac{p_4}{\rho_4} - \frac{p_2}{\rho_2}}{f * Q_{lhv}} = 1 - \frac{T_4 - T_1}{\gamma * (T_3 - T_2)} \tag{2.16}$$

$$imep = \frac{W_{c,i}}{m_f * V_d} = f * Q_{lhv} * \rho_1 * \frac{r_c}{r_c - 1} * \eta_{f,i} \tag{2.17}$$

With the indicated mean effective pressure defined, the power available compared to the sea level standard conditions can be estimated as

$$\frac{P_a}{P_0} = \frac{imep}{imep_0}, \tag{2.18}$$

assuming that engine operating speed is not a strong function of ambient conditions. Note that turbocharged engines normally limit their power output via a wastegate which diverts airflow away from the turbocharger (to prevent overpressurization of the engine), and thereby reduce the $Spa_{turbocharger}$ term. To simulate wastegate operation, $Spa_{turbocharger}$ was determined based on the power required to make $\frac{P_a}{P_0} = 1$ below the break altitude for the standard atmosphere (taken from engine data), and held constant above this altitude, i.e.

$$Spa_{turbocharger} \leq Spa_{turbocharger,break} \tag{2.19}$$

To illustrate, a plot of $Spa_{turbocharger}$ values required to maintain sea level standard (SLS) conditions, with a standard atmospheric break altitude of 6,000 feet for both standard as well as "hot" (defined here as $T_{standard}$+20K) can be seen in Figure 2.7.

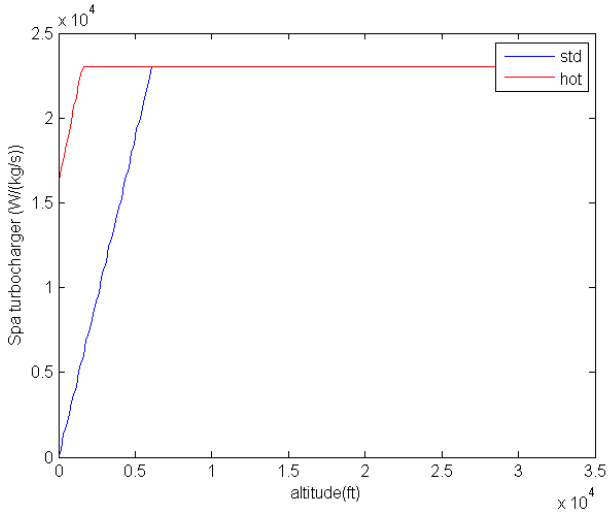

Figure 2.7: Turbocharger power required

Substituting in standard sea level atmospheric values for γ, cp, and R, along with normal operating parameters for T_3 and r_c into the cycle analysis $\frac{P_a}{P_0}$ term results in a complex function based on the freestream temperature ratio θ as well as the $Spa_{turbocharger}$ term, but is only linear with the freestream pressure ratio δ. To that end, the engine was assumed to lapse linearly with a proposed lapsing parameter $\Upsilon = \delta * \theta^{x_t}$, where the x_t term varies from engine to engine. The cycle analysis was then queried for a variety of $Spa_{turbocharger}$ values, along with a range of θ, with a power law used to determine the best fit for the coefficient x_t. A plot of this is shown in Figure 2.8, with no limit on $\frac{P_a}{P_0}$ (i.e. the wastegate is fully closed for the range of operation). Figure 2.8 shows the results of a temperature sweep of the cycle analysis model vs. the best fit Υ for a given turbocharger specific power and x_t.

Figure 2.9 illustrates that, under these assumptions, the parameter Υ causes the 0D model to collapse to a line, indicating a useful, simple choice for estimating power available in a conceptual design context. In other words, the power available is estimated at a constant sea level value below a specified Υ_{break} threshhold, at which point it then decreases linearly with increasing Υ.

2.4 Challenges

When designing aircraft that incorporate more then one energy system, a number of new challenges arise. Much of this is due to increased system complexity. For instance, one needs to ensure that the combined propulsion subsystems can handle not only the energy requirements of the mission, but also the power requirements at each design condition. Additionally, the developer needs to ensure

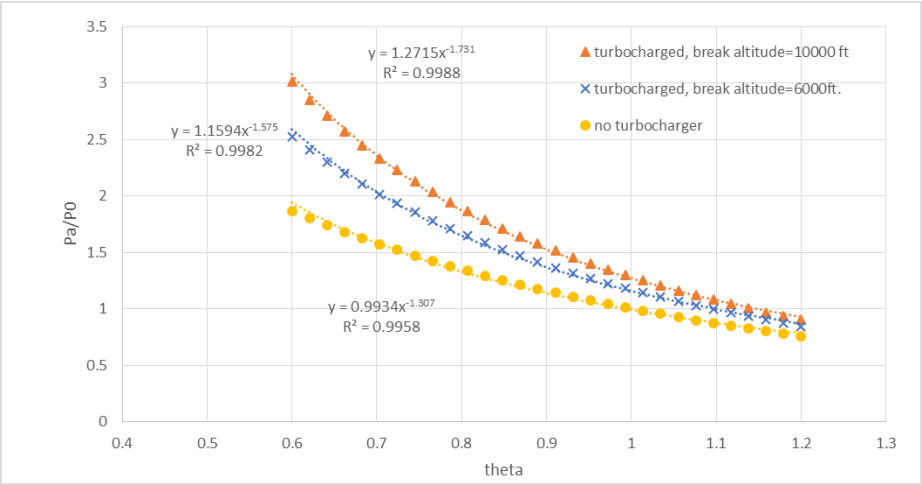

Figure 2.8: Lapsing Parameter vs. Temperature

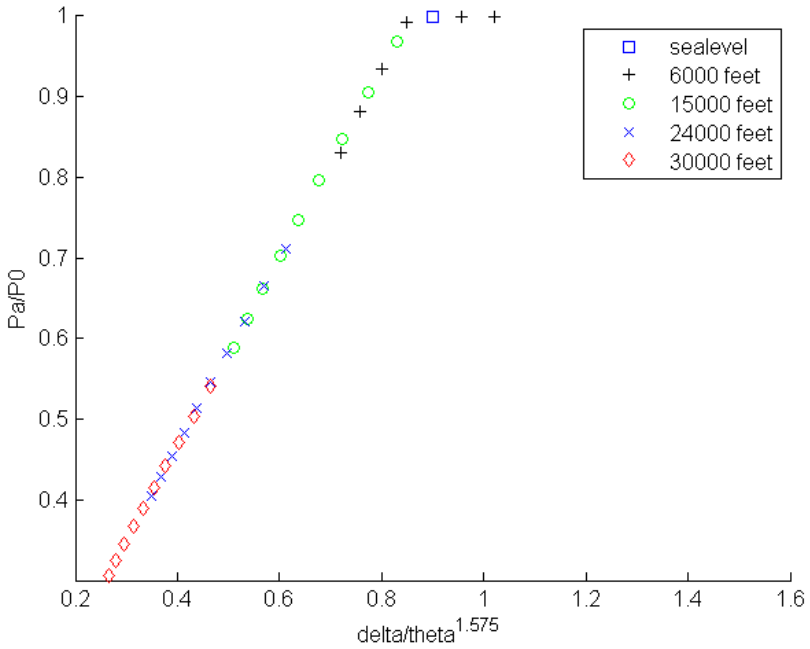

Figure 2.9: Cycle Analysis Temperature Altitude Sweep

that in the final design, each battery subsystem along with the other required subsystems such as speed-controllers, electric motors, and wiring are accounted for and able to handle the power output at all points in the mission. This results in additional overhead in setting up any design/optimization

problem, requiring a larger number of constraints that need to be carefully accounted for as well as careful consideration of the critical design requirements. When trying to analyze these configurations using legacy aircraft design tools, one must be especially careful; tolerance settings may need to be "tweaked" to ensure that the user can obtain meaningful gradient information. Furthermore, significant pre/post processing may need to be performed to evaluate these subsystems and ensure physically achievable results.

Chapter 3

Sizing Methodologies

3.1 General

This dissertation explores two broad categories of aircraft sizing procedures for four different aircraft classes. It also documents some unusual phenomena that can arise in these systems. The first method is an optimizer decomposition-based approach, which is popular in the optimization community [57]. Said algorithm can be explained as follows; the sizing variables are appended to the optimization problem's design variables with relevant constraints handled by the optimizer to ensure aircraft consistency, i.e. (1.13) is replaced with (3.1), while (1.15) is added as a constraint as shown in (3.1)-(3.5).

$$\min f(x,y) \tag{3.1}$$

$$s.t. \quad g(x) > 0 \tag{3.2}$$

$$h(y) = \frac{y - y_{out}}{C} \tag{3.3}$$

$$h(y) > 0 \tag{3.4}$$

$$y > 0 \tag{3.5}$$

where C is a vector of constants used to scale the output parameters h to roughly the same order of magnitude and y is a vector containing a set of "fundamental" sizing parameters, such as aircraft weight, battery energy, and electrical system power. This approach tends to be favored by optimization algorithm developers, as it is often more efficient to allow the optimizer to search along an infeasible region, rather than solve a large number of subproblems [57]. One disadvantage is that, in the event that the optimization fails to converge to a feasible point, the designer gains no new knowledge of the design space. In addition, it was found in a previous paper that, for hybrid-electric aircraft, formulating the problem in this manner can cause difficulties in satisfying

the sizing constraints unless the problems is well-scaled and smooth (i.e. all min/max statements are removed) [99]. This will be further elaborated upon. As the problem becomes more complex, the sizing variables may become loosely coupled, and an abundance of local may minima arise, causing difficulty even for generally reliable optimizers. Another formulation is to solve the subproblem for y as a function of x as shown in (3.6)-(3.9) and (3.3).

$$\min f(x) \tag{3.6}$$

$$s.t. \quad g(x) > 0 \tag{3.7}$$

$$h(y) = \frac{y - y_{out}}{y} \tag{3.8}$$

$$\text{solve for } y \quad \text{s.t.} \quad -\epsilon < h(y) < \epsilon \tag{3.9}$$

This "sizing-loop" method tends to be utilized by aircraft manufacturers as well as students in conceptual design. It has a number of advantages, in that, in the problems considered here, any optimal configuration would have an $h(y) = 0$. Thus, the optimizer may be less likely to become "stuck" in local minima (at least those defined by the sizing variables). Additionally, many trusted aircraft design tools in use today deploy similar methodology; as a result, the aspiring aircraft designer may adopt this methodology for reasons divorced from any mathematical properties [27, 52, 80]. Furthermore, as each optimizer iteration is properly sized, the designer can build some intuition about the sensitivity of the aircraft to input variables as the trade studies are performed or an optimization process continues. A major disadvantage is a potential loss in solution accuracy; the solution can only be converged to a certain value, practically speaking, before accumulated numerical errors from solving various subsystems (such as a mission solver or higher-fidelity aerodynamics solver) prevent further convergence. For the aluminum-air aircraft, this was found to begin occurring at a sizing tolerance of 1E-5. This means that the finite differencing steps from the optimizer should be increased to ensure that meaningful gradient information can be obtained; the users guide to SNOPT recommends a step size $\approx (tol)^{1/2}$ [37]. For more elaborate mission solvers, higher-fidelity subdiscipline analysis, or tighter solution tolerances, solving these problems may be difficult or even impossible to the tolerances required in optimization. An additional complication is that, for a given decision parameter x, solutions such that $h(y) = 0$ may be non-unique in y, which poses no small difficulty, as will be seen in Chapter 5. A basic information tree using these sizing variables can be seen in Figure 3.1. Note that this particular algorithm is called at every optimizer iteration.

Now in the event that the user, rather than running an optimization case, wants to either run trade studies for fixed design variables (such as wing loading, aspect ratio, etc.), or just analyze a single aircraft case, he or she may run a fixed point iteration solver on y such that $h(x, y) = 0$ (1.15). A visual for a single battery aircraft using a fixed point iteration or root finding method can be seen in Figure 3.2.

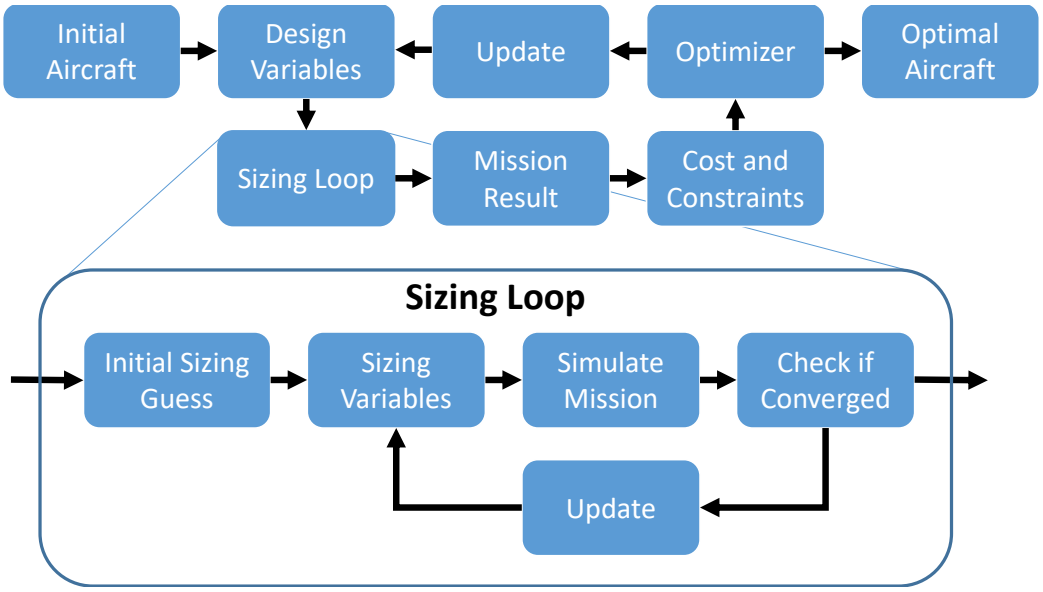

Figure 3.1: Sizing-Loop Algorithm

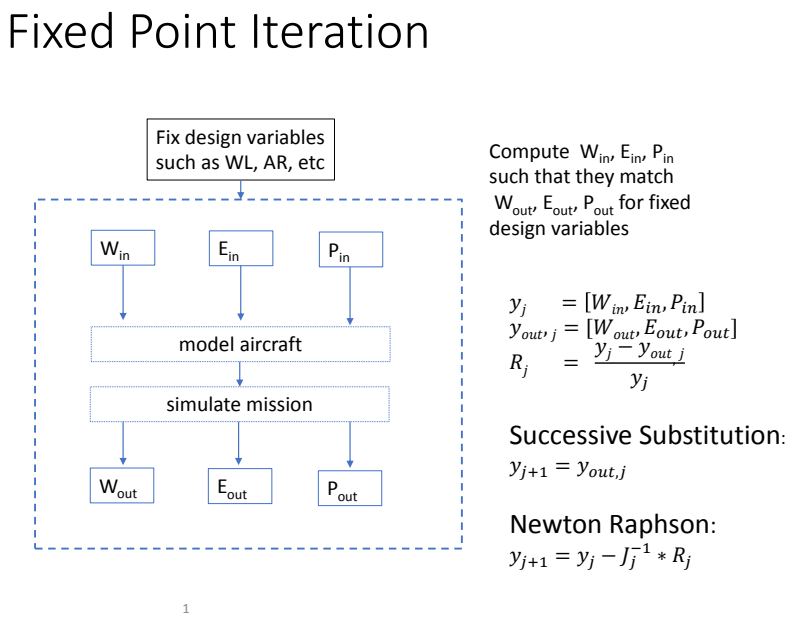

Figure 3.2: Fixed Point Iteration Diagram

3.1.1 Partial Differential Equation Analogy

Most of the results in this dissertation use some form of fixed-point iteration to solve these aircraft design problems, often with some outer loop optimizer to wrap the system. Some insight into the

behavior and limitations of this methodology can be gained by observing the result and assumptions that typically go into similar systems. Now, fixed point iteration methods are often used to solve systems of partial differential equations, such as in Computational Fluid Dynamics (CFD), which, like the problems here, can employ an outer loop optimizer for aerodynamic shape design [41]. Similarly formulated problems can also be seen in optimal control. The goal of these systems is to minimize some objective, while solving the system for a set of state variables (y), that are in turn assumed to be some unique function of the design variables (x) [42]. This work will show that certain classes of aircraft design problems, taken analogously, do not show this uniqueness characteristic, and in fact, valid designs can close for a set of design choices x with wildly varying values for y. This can occur as a result of using either separate initial guesses for y to solve the system, using different numerical methods to solve the system, or some combination of the two. Taking this analogy further, the "sizing loop" formulation, where the sizing variables are computed as a function of the design variables may also be considered as a "reduced space" approach, while the optimizer decomposition framework used here may be considered closer to a "full space" method [11]. Note that many of these cases explicitly start the optimization from a converged case of the PDE to increase code speed and ensure that it starts with a reasonable initial guess [58]. Additionally, it can be shown that, for the set of active constraints, linearly independent gradients are sufficient to satisfy the KKT conditions for a given (x,y) at a feasible point [58, 62]. This can lead to some insight into the underlying systems and may help explain some of the unusual phenomena observed in Chapter 5.

3.2 SUAVE Overview

All of the aircraft analysis and design/optimization studies starting in Chapter 6 will utilize SUAVE, an open-source design tool[63]. SUAVE, was created with the goal of providing an environment for both the analysis and the optimization of aerospace vehicles that would be compatible with exotic configurations. In prior work, a brief description of some of the analysis capabilities of the environment was presented [63]. SUAVE was built from the beginning with the capability to employ any number of exotic energy systems, and thus, is well-suited for studies such as this.

SUAVE's initial code was written for use with multiple information sources and extreme flexibility. Any parameter that defines the vehicle (the configuration of the vehicle, the mission, or even the analysis) can be exposed to an optimizer. Combining the available optimization variables with this analysis flexibility allows for the design of new aircraft concepts/missions. For example, a single aircraft can be optimized under multiple missions and configurations concurrently, or a family or even a fleet of aircraft can be optimized simultaneously.

A more recent paper highlighted SUAVE's flexibility in formulating optimization problems which this dissertation will heavily leverage [13]. The analysis routines are coded in such a way that arbitrary inputs should always provide smoothly varying results. For robustness, an output should

always be provided back to an optimizer even when the optimizer's attempted inputs are physically infeasible. This informs the optimization algorithm at every input and prevents a breakdown in the optimization process. Additionally, this requirement necessitates careful examination of the underlying code to ensure that analyses will always provide finite numerical results provided the inputs meet broad specifications.

Another key feature of SUAVE is the use of underlying physical principles to analyze exotic and unconventional vehicles. In these radical designs, the intuition of the classically trained designer starts to fail as the design space is shifted dramatically, and as a result handbook methods or rules of thumb become unreliable. Building intuition for this new design space requires both a tool such as SUAVE, and the capability to drive it in any way the designer deems necessary. This could potentially include outside optimization routines or uncertainty quantification tools.

3.3 Other Tools Used

Optimization cases were handled using SUAVE's pyOpt optimizer wrapper, which in turn called SNOPT (a Sequential Quadratic Programming method) [77, 37]. pyOpt is a Python based, object-oriented framework used to call and more easily communicate with a variety of optimization packages within the Python language. In addition, in Chapter 5, a particle swarm algorithm will also be used with the results compared [50]. Furthermore, to aid in computing an initial guess for the sizing variables y, a number of machine learning regression models from scikit-learn were used [75]. scikit-learn is a machine learning toolbox, also written in Python, which offers a variety of regression, classification, as well as clustering packages that can be used to handle and obtain some insight into the behavior of real life data as well as "black box" packages. Details of this approach can be seen in Chapter 7.

Chapter 4

Numerical Methods

4.1 Fixed Point Iteration Methods

A common method used to solve the inner-loop problem (see Section 3.1) is successive substitution, where y_{out} is set as a new y, and the loop iterated on until the guesses converge. This is a fairly slow, but robust method. Solving this problem is relatively simple for a single variable, but becomes more complicated as more sizing variables are added, especially when they are all loosely coupled. For example, for the aluminum-air aircraft case (see Chapter 6), the auxiliary battery often only consumes a small amount of the total energy, yet accounts for a substantial fraction of the power; thus, using total energy as a sizing variable results in highly infeasible aircraft, and needs to be handled separately. Figure 6.7 compares the residuals associated with each sizing parameter at each successive substitution iteration as an illustration of this problem. As the solution proceeded (for fixed x), there was significant oscillation in the error; furthermore, the different sizing errors were out of phase with each other, which significantly slowed the convergence rate. Furthermore, to converge all of the required parameters, some of the other associated residuals may become over-converged, which, due to the higher range factor, may affect the design enough to result in a serious reduction in gradient accuracy. One method to reduce the number of sizing iterations is to use Newton-Raphson (nr), where the algorithm in Figure 3.1 is finite-differenced in y to obtain the Jacobian with respect to the residual h, i.e.

$$J_j = \frac{D(h_j)}{D(y_j)}. \tag{4.1}$$

The new y value becomes

$$y_{j+1} = y_j - (J_j^{-1}) * h(y_j), \tag{4.2}$$

where y_j refers to the j'th iteration of y, where y is a vector. Newton-Raphson is known to exhibit quadratic convergence near the solution, which can significantly reduce the number of function

evaluations. Newton-Raphson methods used within the dissertation will use successive substitution to solve the problem until $h(y)$ is converged to within 5% (unless otherwise specified) for all values of h, then use nr to accelerate convergence. Now, the cases here take the definition of the residual from (3.8). For illustrative purposes, the residual may be non-normalized, and simply computed as

$$h(y) = y - y_{out}. \tag{4.3}$$

This case tends to be more numerically "nice" than the previous case, but comes with the downside that separate values inside the y vector may be scaled very differently across the design space, which should be addressed. Now, it should be noted that normalizing the value for h by the input value y introduces a degenerate case when a given sizing parameter may equal zero; in the context of the work here, in those cases the residual was instead normalized by the average of y and y_{out}. As an aside, this non-normalized case has the interesting property that the Jacobian becomes

$$J_j = \frac{D(y_j - y_{out,j})}{D(y_j)}. \tag{4.4}$$

Furthermore, as a comparison of the two algorithms, successive substitution and Newton-Raphson are identical when the Jacobian is set to 1 and $h(y)$ is non-normalized as in (4.3). Additionally, depending on the sizing tolerance, Newton-Raphson, due to its quadratic convergence, may cause a particular sizing iteration to become overconverged, introducing numerical noise when connected to potential outer loop optimizer.

One advantage of Newton-Raphson over successive substitution in the aircraft presented here is that it takes into account how changes in each sizing variable affect residuals associated with the other variables, resulting in less oscillatory behavior. Another method to increase code speed is to use Broyden's method with the Sherman-Morrison formula to directly update the inverse of the Jacobian [86]:

$$J_{j+1}^{-1} = J_j^{-1} + \frac{dy - J_j^{-1} dh}{dy^T J_j^{-1} dh} dy^T J_j^{-1} \tag{4.5}$$

where

$$dy = y_{j+1} - y_j \tag{4.6}$$

and

$$dh = h(y_{j+1}) - h(y_j). \tag{4.7}$$

This reduces the number of function evaluations by eliminating the need to run a finite difference scheme to form the Jacobian at each iteration which can become expensive if either the objective function itself is costly or y is high-dimensional, or both. However, this update is also less robust, and the Jacobian may need to be reinitialized to ensure convergence. A demonstration of these methods "in action" can be seen in Figure 4.1, which compares the mean squared error of h at each

sizing loop iteration for the same aluminum-air reference guess for successive substitution, Newton-Raphson, and Broyden's Methods. Gaps in the iterations correspond to finite-difference steps to build the Jacobian.

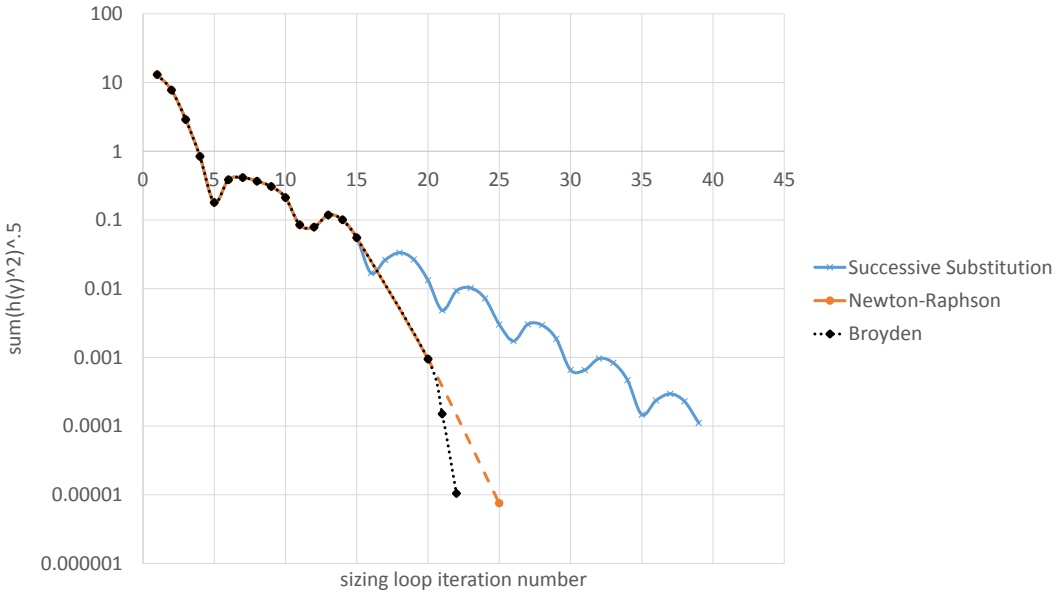

Figure 4.1: Convergence Results for higher-dimensional y

Figure 4.1 shows that, using these methods can solve the sizing loop problem considerably faster than simple successive substitution. However, the user should be careful, in that, due to their quadratic convergence, they may overshoot the solution tolerance (which here is 1E-4), which may cause significant problems when calculating gradients. In this instance, the convergence jumped from 1E-3 to 1E-5 in a single iteration. This is a less critical issue when one uses global-optimization techniques, as accurate gradients are not nearly as important there. Both successive substitution and Newton-Raphson will be compared in a 2D sweep of aircraft designs in Chapter 8, while successive substitution, Newton-Raphson, as well Broyden's method will be compared in full, gradient-based optimization cases in Chapter 9.

Figure 6.7 also shows that properly informing the initial guess for y is a crucial aspect of increasing code speed. To that end, a number of different techniques were used to reduce the number of sizing loop iterations, thereby improving the optimization efficiency. One easy way to increase code speed (which is effective for gradient-based optimizers), is to tabulate the converged solutions, and, when the input variables x are "close" to a tabulated value use the nearest converged $y(x)$ as an initial guess. However, farther away from the tabulated data, one should use other methods to determine a

"good" initial guess for the y values. To that end, a variety of regression algorithms from scikit-learn were evaluated to determine their effectiveness in finding a good guess for the initial y in the sizing loop [75]. Details can be seen in the Chapter 7 section of this report.

4.2 Backtracking

Now, as will later be seen in Chapters 8 and 9, backtracking, in certain cases, can prove crucial in ensuring robust convergence of these designs. In this case, for an optimization or root-finding problem, one updates a guess based on a step size α (≤ 1) and a search direction p.

$$x_{j+1} = x_j - \alpha * p \tag{4.8}$$

From here, one determines the value of $f(x_{j+1})$ in an optimization problem or $\sum h(x_j)^2$) in a root-finding problem, and compares their values to that at x_j. One then reduces the step size α until certain conditions are satisfied. While a variety of options exist to decide on this termination, here, α will be halved until the condition

$$\sum h(x_{j+1})^2 \leq C1 * \sum h(x_j)^2 \tag{4.9}$$

is met, where $C1 = 1$ unless otherwise specified. For optimization problems, one typically adds a condition on the gradient of the objective to ensure sufficient decrease in the augmented objective function. Additional conditions on curvature can also be added to ensure that the line search terminates near a critical point, which can aid in gradient-based optimization. Of note is that for Lipschitz continuous first partial derivatives, using backtracking for a steepest-descent method is provably convergent [7]. Additionally, one should bear in mind that, for the systems here, backtracking is performed in y (i.e. the sizing variables) when running a root-finding method on the problem.

4.3 Regression Methods

Under the scikit-learn "umbrella," a large number of regression tools were used and compared in their effectiveness in simulating a full sizing/optimization problem. The regression methods were split into linear regression algorithms, which include methods such as Gaussian Process Regression, Support Vector Regression, K Nearest Neighbors, and ensemble methods [79, 89, 24]. The ensemble methods, rather than using a specific algorithm to fit the variables, stochastically search in function space to determine what combinations of functions work well for the given dataset. Examples of these algorithms include Gradient Boosting, Bagging, Random Forest, and Extra Random Trees [35, 16, 60, 36]. Thus, the ensemble methods are nondeterministic, but may be more accurate and

robust depending on the dataset.

4.4 Successive Substitution and Unstable Solutions

Now, as will be seen in Chapters 5 and 8, the solutions to the systems of equations presented here may be nonunique in y. Reasons for this will be explained in Chapter 5; however, an additional point of interest is how these algorithms tend to handle the presence of multiple solutions. To illustrate, a hypothetical function $y_{out}(y)$ is plotted alongside the line for $y_{out} = y$ (i.e. the space for potential solutions to a fixed point iteration algorithm) in Figure 4.2.

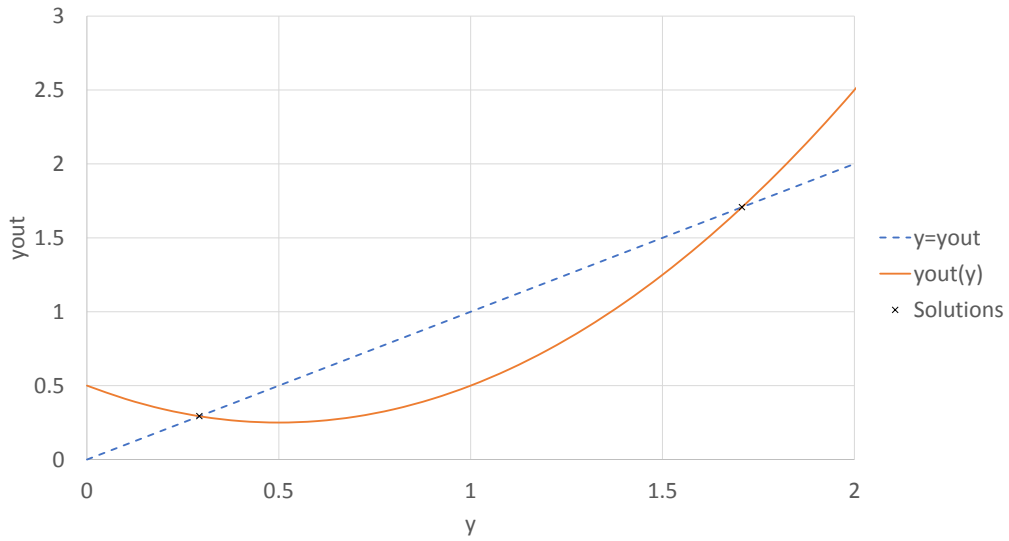

Figure 4.2: Unstable Point Example

The two intersection points at $y \approx .3$ and 1.7 are fixed points of the system, i.e. solutions. A Newton-Raphson iterative algorithm will converge to both solutions rather easily. However, it can be shown that the larger solution ($y \approx 1.7$) is unstable when using the successive substitution algorithm; as a result this methodology will either converge to the smaller solution ($y \approx .3$) or diverge to infinity. To explain, at a given iteration j, successive substitution at y_j will return $y_{j+1} = y_{out,j}$ for the next iteration. Looking at this graphically, if the curve $y_{out}(y)$ is above the 45°line ($y_{out} = y$) for a given y_j, a successive substitution iteration will return a higher value for y_j. Conversely, if the curve $y_{out}(y)$ is below the $y_{out} = y$ line at a given y_j, then a given successive substitution iteration will return a lower value for y_{j+1}. So cases to the left of $y \approx .3$ will return larger values

of y, while cases between $y \approx .3$ and 1.7 will return smaller values in y until they converge to the leftmost solution. However, starting from y values higher than 1.7 will return ever higher values for y_{j+1} until it diverges to infinity.

Now, mathematically speaking, this unstable point is not a "good" property to have. However, in nearly all cases, in the event that multiple solutions exist to the system of aircraft sizing equations, an aircraft designer is only interested in the smaller solution. As a result, while Newton-Raphson is both faster and more robust at finding solutions to a system of equations, successive substitution may actually be preferred in many cases because there may be solutions that the aircraft designer does not want to find in the first place.

Chapter 5

Analytical Cases

5.1 Problem Formulation

To help diagnose some of the key phenomena observed in this work as well as to demonstrate some of the limitations in current design methodology, several analytical "toy" aircraft sizing/optimization problems were created, defined by some decision parameter x, which represents the tradeoffs between aerodynamic and structural efficiency. Additionally, there is a dependence on the weight of the aircraft W_{in} to model the fact that structural weight fraction and aerodynamic performance do not remain fixed with aircraft size. This problem uses analytical functions to represent $\frac{L}{D}$ and f_{empty}, as shown below:

$$\frac{L}{D} = (20 - 3x + x^2) * (1 + ln(\frac{W_{in}}{10,000g}))^{\frac{1}{4}} \tag{5.1}$$

$$f_{empty} = (1.1 - .8x + .2x^2) * .35 * e^{(\frac{W_{in}}{10,000g})^{.17}}. \tag{5.2}$$

These parameters are plotted on top of each other for constant W_{in} below, illustrating some simple tradeoffs for the design choices one might make. Note that, throughout this section, the decision parameter x is kept between .01-3, to prevent spurious results such as an empty weight fraction $\geq$ 1. The weight dependent terms are plotted in Figure 5.2.

Now, in addition to the decision parameter x, there are anywhere from 1-4 sizing parameters (y) that fall out from the decision parameter, depending on the aircraft case. In the single battery case (Section A.2), the sizing vector is $\frac{[W_{in},E_1,P_{in}]}{y_{scale}}$, where W_{in} is the landing weight (a metal-air battery is assumed, so the landing weight would determine the size of components such as landing gear), E_1 refers to battery energy storage, and P_{in} is electric system power output capability. y_{scale} is a vector consisting of the constants 1E4, 1E8, and 1E5, respectively, in order to normalize these input

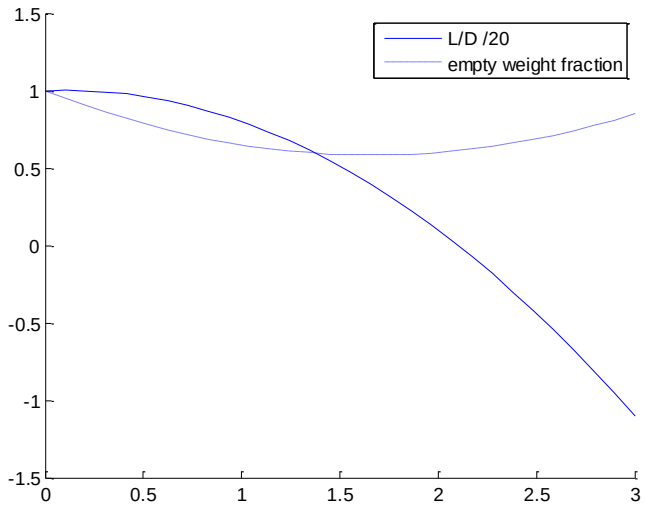

Figure 5.1: Analytical Empty Weight and L/D

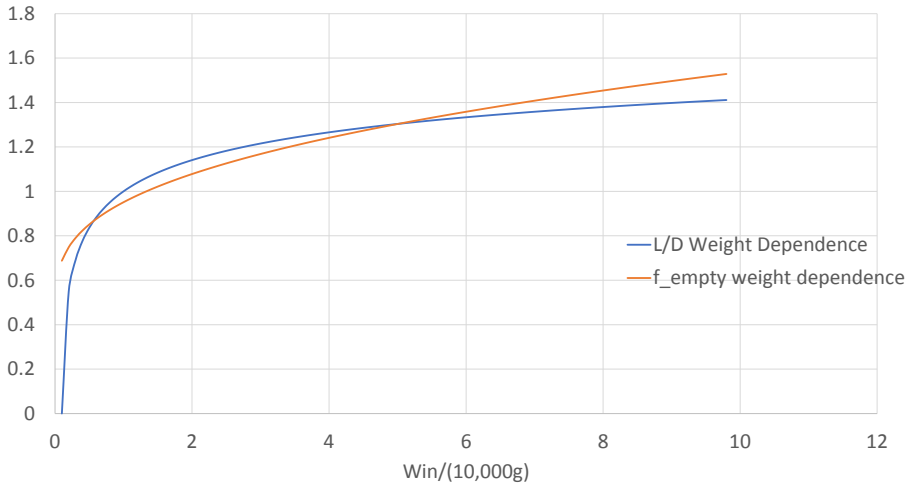

Figure 5.2: L/D and Empty Weight Fraction Dependence on GTOW

parameters to a comparable order of magnitude. For the hybrid battery case, y is a vector consisting of the parameters W_{in}, E_1, E_2, and P_{in} normalized by a different scaling vector, also called y_{scale}. In all cases, W_{empty} is simply

$$W_{empty} = f_{empty} * W_{in}. \tag{5.3}$$

In solving and defining numerical systems, function smoothness is of paramount importance. To that end, a soft max function was used in place of max statements, which is defined below [23].

$$f1 = \max(y1, y2) \tag{5.4}$$

$$f2 = \min(y1, y2) \tag{5.5}$$

$$\text{soft\_max}(y1, y2) = f1 + \ln(1 + e^{f2-f1}) \tag{5.6}$$

This function is plotted in Figure 5.3 for two lines to illustrate its smoothing characteristics.

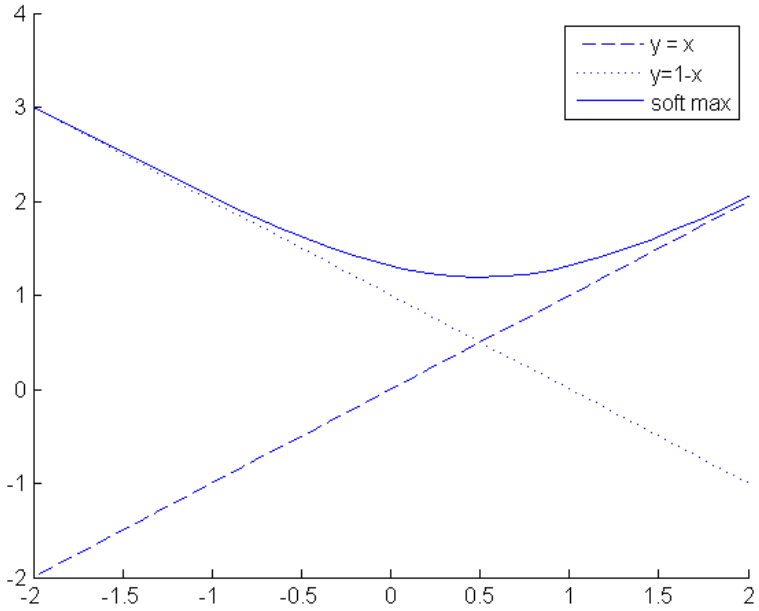

Figure 5.3: Soft Max function

Some useful features of this formulation are that it is continuous, ensuring that any computed gradients are sensible. In addition, it is conservative, so the user can be confident that any value returned by this function is at least as large as the actual max. Now, as the cases here utilize metal-air batteries, a parameter called "mgf" (mass gain factor) is used to determine the oxygen

mass that the battery consumes as it discharges.

$$W_{air} = g * mgf * E_1 \tag{5.7}$$

In a practical aircraft design process, the mgf for a particular battery chemistry may be derived analytically from the half-cell equations and Faraday's Laws of Electrolysis as in (1.2). In this chapter, the mass gain factor is set to 1E-8 for all cases. With these parameters defined, the lift L is taken as the estimated average throughout the mission, i.e.

$$L = W_{in} - \frac{W_{air}}{2}. \tag{5.8}$$

The power requirement for cruise is then

$$P_{cruise} = V * \frac{L}{\frac{L}{D}}, \tag{5.9}$$

where V is set as a constant 100. P_{out}, defined as the maximum power the aircraft experiences during the mission (such as during takeoff or hover, for VTOL aircraft), is

$$P_{out} = mpf * P_{cruise}. \tag{5.10}$$

There is often a complex tradeoff in design choices that set takeoff, climb, and hover power, but here, a constant mpf value of 1.5 is maintained for simplicity. Two different methods are used to converge the designs: successive substitution and Newton-Raphson. When the numerical method is "Newton-Raphson," successive substitution is used to converge the case until all residuals $h(y)$ ((A.5)) are 5% or lower, at which point Newton-Raphson is used to accelerate convergence. Additionally, backtracking is used in each case, according to the conditions defined below. First, define a line search direction p

$$p = y_{j+1} - y_j, \tag{5.11}$$

where y_{j+1} is determined by either a successive substitution iteration or from a step using Newton-Raphson. Then backtracking is performed using the algorithm defined in (4.8) and (4.9). Detailed formulations for four different progressively simpler aircraft design problems can be seen in Appendix A.

5.2 Results

Now for the four cases shown here, parameter sweeps of x (and for the hybrid battery, f_{aux}) were performed. Several sets of initial guesses for the set of sizing parameters y were used, comparing

the solutions obtained when using a successive substitution approach to a Newton-Raphson-based approach. This is to identify, as will be seen in Chapter 8, the cause of the non-unique solutions observed. To start with, results from the hybrid battery case (i.e. the most complicated one) will be used, followed by results from the less complicated designs.

5.2.1 Electric Aircraft (Hybrid Battery)

The exact system to be solved for this case can be seen in the Appendix Section A.3. To begin with, sweeps of landing weight divided by 10^4 for converged aircraft cases are plotted vs. the decision parameter x and auxiliary power fraction f_{aux}. Successive substitution is used to zero the sizing residuals h in Figure 5.4, while Newton-Raphson is used to zero the sizing residuals in Figure 5.5. Nonconverged cases are shown in white.

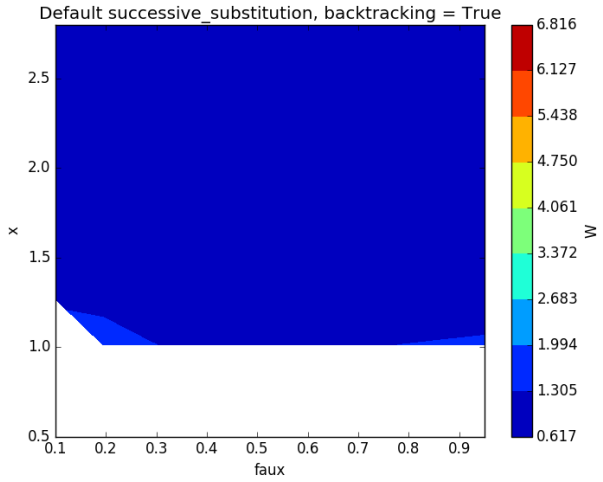

Figure 5.4: Weight (Successive Substitution, $y_{initial} = [5, 5, 5, 5]$)

As should be readily apparent, certain regions of the design space look startlingly unalike despite using the exact same underlying functions; the only difference between Figures 5.4 and 5.5 is the method used to converge the system. Furthermore, certain cases converge when using successive substitution, but diverge to infinity when using Newton-Raphson to solve the problem. Note that the color bars are identical, allowing one to more easily understand the scale of the differences. To gain a more complete understanding of these discrepancies, the normalized difference between these two carpet plots is shown in Figure 5.6.

Figure 5.6 demonstrates that both methods converge to the same result for the bulk of the design space (the contour map here is a 10x10 grid). Certain exceptions, however, exist; to see this in more detail, and compare these two solutions, Table 5.1 compares the two converged designs at a fixed

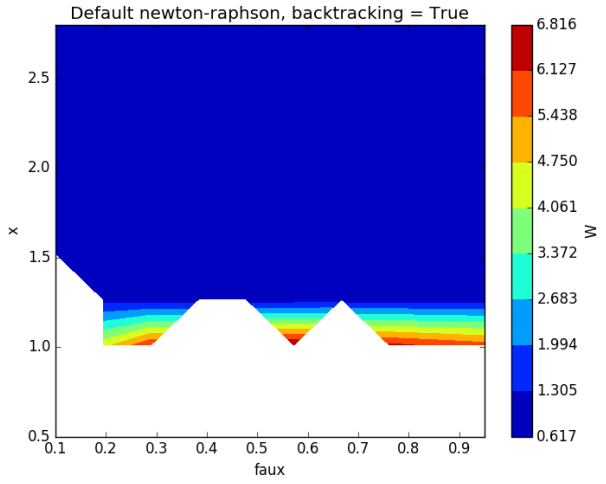

Figure 5.5: Weight (Newton-Raphson, $y_{initial} = [5, 5, 5, 5]$)

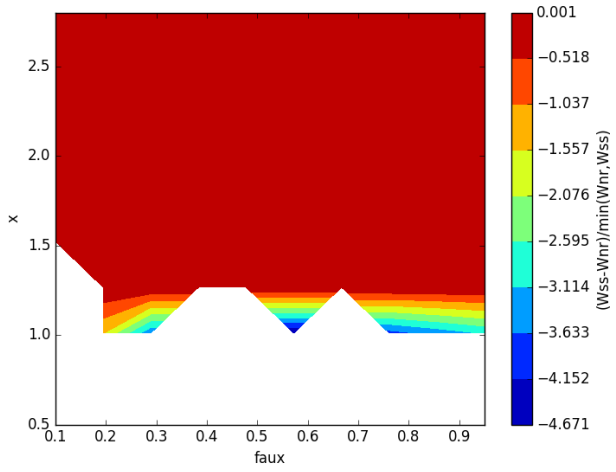

Figure 5.6: Normalized Difference, $y_{initial} = [5, 5, 5, 5]$

value of x and f_{aux}.

Table 5.1 shows that the ratio between the two separate designs is the exact same for $E1$, $E2$, and P_{in}, with a separate scaling ratio for W_{in}. This is due to the fact that the specific energy for each battery as well as the specific power of the motor are all maintained for every design; depending on the system, this may not be a realistic assumption. Additionally these two designs have an empty weight fraction that differs by a factor of about 5% due the weight dependence in (5.2). Now, the

Table 5.1: Solution Comparison (Hybrid Battery)

$x = 1$, $f_{aux} = .2$	$\frac{W_{in}}{g*10^4}$	$\frac{E1}{10^8}$	$\frac{E2}{10^8}$	$\frac{P_{in}}{10^5}$
design 1 (Newton-Raphson)	4.470	84.093	1.689	28.915
design 2 (successive substitution)	1.769	37.350	0.749	12.842
ratio	2.52	2.25	2.25	2.25

discrepancies between these designs can be erased if, rather then starting from $y = [5, 5, 5, 5]$ for all cases, having the two methods begin from $y = [1, 1, 1, 1]$, as Figure 5.7 shows.

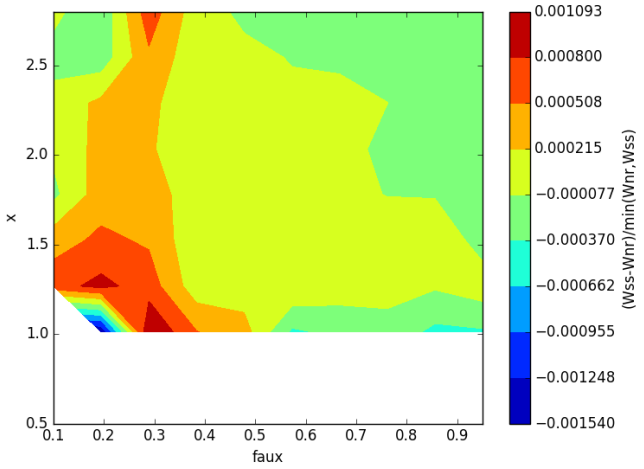

Figure 5.7: Normalized Difference, $y_{initial} = [1, 1, 1, 1]$

As Figure 5.7 illustrates, the solutions converge to the same values when using two different root-finding strategies depending on where they are started from. As a result, the nonuniqueness of these cases may remain undiscovered if the designer is careful (or in some cases, lucky). Indeed, the initial guess used for the sizing parameters from Figure 5.6 for some values of (x, f_{aux}) is higher than any converged case from the parameter sweep in Figure 5.7 (specifically, for weight). Nonetheless, there is an extreme variation in scale among most of the sizing parameters as Table 5.2 demonstrates.

Table 5.2: Solution Range (Hybrid Battery)

	$\frac{W_{in}}{g*10^4}$	$\frac{E1}{10^8}$	$\frac{E2}{10^8}$	$\frac{P_{in}}{10^5}$
max, design 1 (Newton-Raphson)	7.569	124.563	100.012	46.873
min, design 1 (Newton-Raphson)	0.617	1.324	0.202	5.826
max, design 2 (successive substitution)	1.769	37.35	29.489	12.842
min,design 2 (successive substitution)	0.617	1.327	0.202	5.826

From the results in Table 5.2, one should immediately note that the minimum converged result for each sizing parameter is identical between the Newton-Raphson and successive substitution cases. In particular, depending on the choice in parameter for x and f_{aux}, there appears to be multiple orders of magnitude in these solutions, particularly in the energy required of each battery subsystem. As a result, it may be difficult, if not impossible to estimate, a priori, appropriate values for y_{scale}. Fortunately, in all cases here (and indeed, in all cases discovered by the author thus far), this multiplicity of solutions tends to be of the sort that, at the least, values for W_{in} are quite different from each other, and that these solutions do not appear to arise solely from the "play" between the separate energy systems as the rest of this chapter will prove. Furthermore, based on experience, both in this section as well as in Chapter 8, the nonunique solutions tend to occur relatively "far" away from the globally lightest aircraft. As a result, when trying to minimize weight for these designs, this nonuniqueness property should pose little issue both in the local and global optimization cases (at least in terms of changing the final solution). Other choices for objective function (such as operating cost) may present more of a problem, but as of this writing, this has not occurred. In addition, while the difference between these designs may be significant, the designer should be able to intuitively throw out the larger result. Computational algorithms on the other hand would be less reliable in filtering and removing these extraneous solutions.

5.2.2 Electric Aircraft (Single Battery)

Now, the nonunique solutions shown in Section 5.2.1 were produced specifically in an attempt to mimic the results first discovered in the SUAVE case from Section 6.3, with nonunique sizing solutions shown in Section 8.1. An important question in either instance is the exact mechanism that is responsible for this multiplicity of solutions; initially, it was first thought that this nonuniqueness property was specific to aircraft with multiple energy systems. To determine whether this characteristic is responsible, one of the batteries was removed, simplifying the problem and yielding the aircraft design process described in Section A.2. Using this system of equations, it is relatively easy to reproduce this multiplicity of solutions as Figures 5.8 and 5.9 illustrate.

Like before, the normalized difference between the solutions is plotted. This can be seen in Figure 5.10. Again, the difference in computed landing weight between the two designs is of order three, meaning that the designer should be able to intuitively recognize and throw out the "bad," larger candidate design.

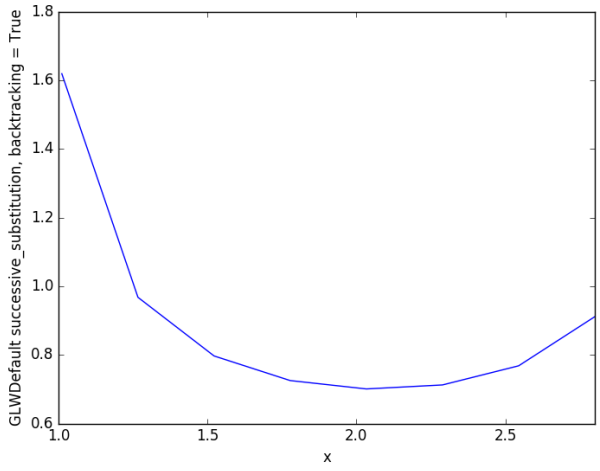

Figure 5.8: Weight (Successive Substitution, $y_{initial} = [5, 5, 5]$)

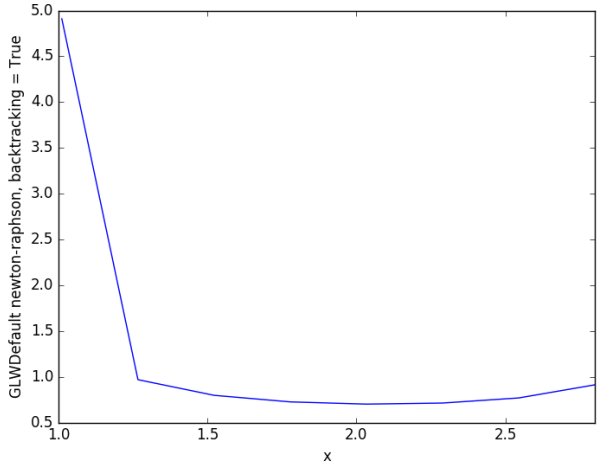

Figure 5.9: Weight (Newton-Raphson, $y_{initial} = [5, 5, 5]$)

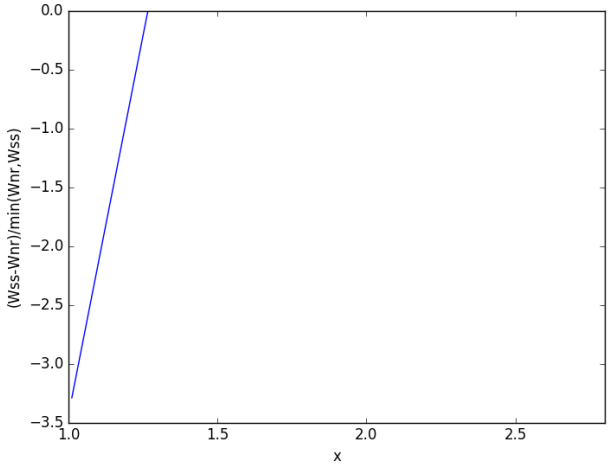

As with the hybrid battery case, reducing the initial y value for every input guess tends to ensure that the two methods converge to the same result for the values of x used here. This is shown in Figure 5.11.

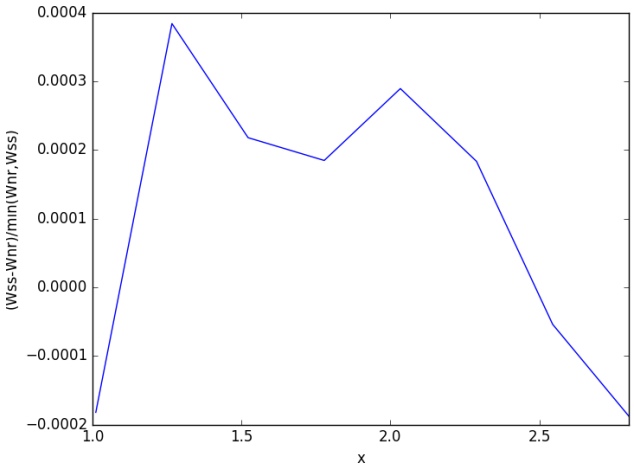

Figure 5.11: Normalized Difference, $y_{initial} = [1, 1, 1]$

Like the hybrid battery case, this result helps to show that this multiplicity of sizing solutions is not a result of a mistake or bug in the code. Additionally, it demonstrates that this undesirable property is avoidable by using the "right" initial guess to solve the problem.

5.2.3 Electric Aircraft(Single Battery, no Power Requirements)

Now, when obtaining the results in Sections 5.2.1 and 5.2.2, a couple of related phenomena were thought to be responsible. Both involve the soft_max function used to determine the battery weight in (A.6), (A.9), and (A.10). The thought was that the use of a max function, even a soft_max function introduces discontinuities that may make the design space "lumpy," which could contribute to the multiplicity of solutions observed thus far. Looking at this from a physical perspective, the possibility of moving from an energy-limited to a power-limited battery may open the potential for an additional solution to the sizing residuals system. The case here both removes any and all max functions, as well as power-related sizing characteristics, removing two potential culprits for this undesirable phenomena. When removing power requirements from the model, it is still possible to produce multiple cases in y that zero the sizing residuals. A comparison of said results can be seen in Figures 5.12 and 5.13.

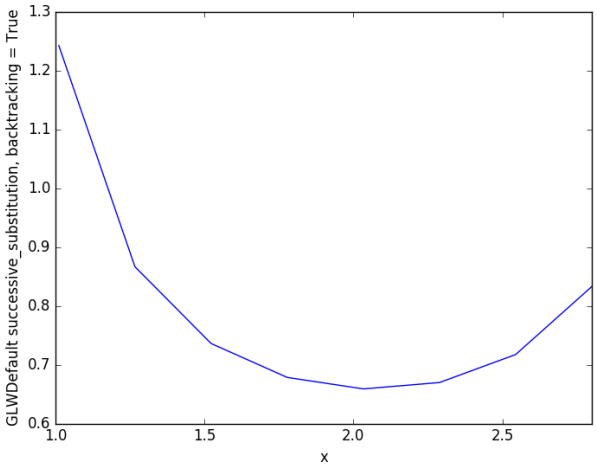

Figure 5.12: Weight (Successive Substitution, $y_{initial} = [5, 5]$)

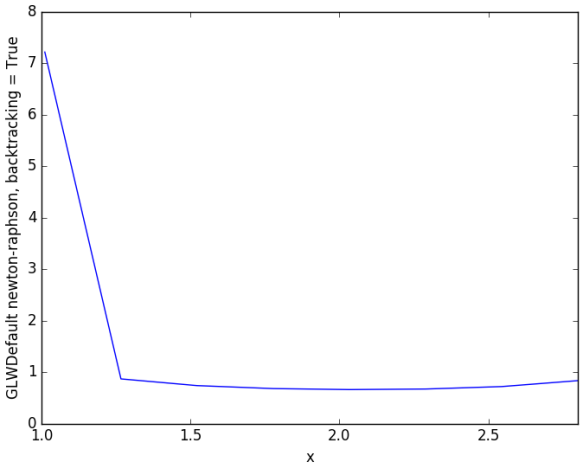

Figure 5.13: Weight (Newton-Raphson, $y\_initial = [5, 5]$)

As before, the difference between the two plots was also plotted and can be seen in Figure 5.14.

Observant readers will notice at least one important characteristic of Figures 5.12–5.14. That is, that the Newton-Raphson case, when the solution differs from the successive substitution case, may possess a value for W_{in} that is nearly an order of magnitude larger. Now, like before, one can produce the same solutions for all cases by changing the initial guess as Figure 5.15 shows, which

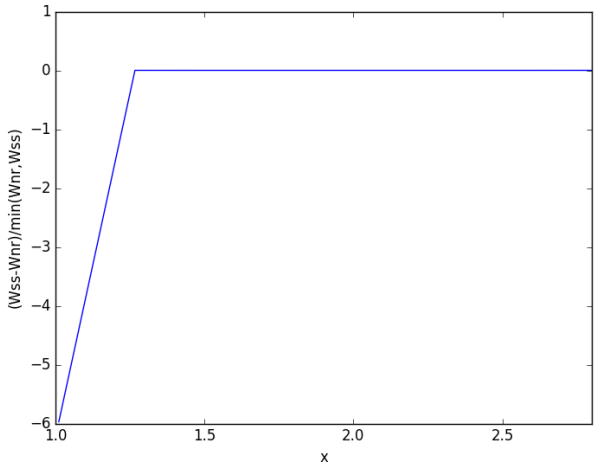

Figure 5.14: Normalized Difference, $y_{initial} = [5, 5]$

plots the normalized difference when $y_{initial}$ = [1,1].

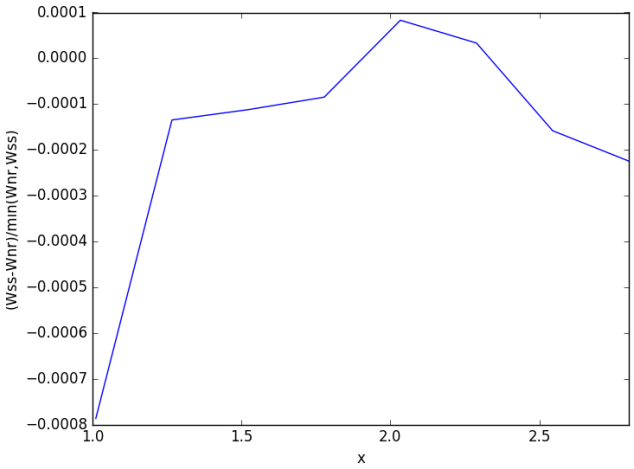

Figure 5.15: Normalized Difference, $y_{initial} = [1, 1]$

With all potential sources of discontinuities removed from the system, along with all direct modeling of power requirements in the component weight sizing, the question remains: why, in these cases, are multiple sizing solutions possible? The following, simplified case should provide some further insight and aid in offering an explanation.

5.2.4 Electric Aircraft (Single Battery, no Power Requirements, Simplified)

The system formulation of this section can be seen in Section A.4. It details what is, in the view of the author, the most basic aircraft design system here, and is essentially the exact same mathematically as those commonly used in conventional, combustion engine-powered designs. Now, as this is the simplest case presented, more time will be spent discussing the results and trends to diagnose the issues presented thus far. As before, linear sweeps of converged results for W_{in} are plotted vs. x and can be seen in Figure 5.16, along with their normalized difference.

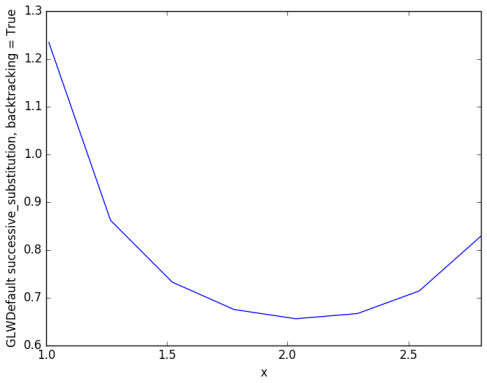

(a) Weight (Successive Substitution)

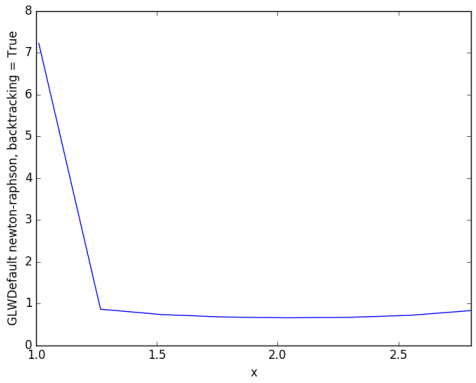

(b) Weight (Newton-Raphson)

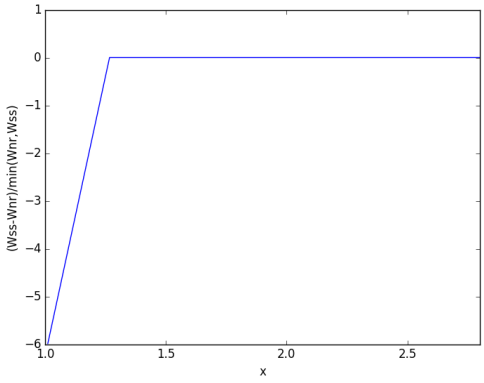

(c) Normalized Difference

Figure 5.16: $y_{initial} = [5]$

It should be some small comfort to note that Figures 5.12–5.14 are identical to the plots in Figure 5.16 i.e. that the "physical" systems described are the exact same (see Sections A.1 and A.4), and, depending on the methodologies used, converge to the same result; the only difference is in this case

the system is further simplified to iterate on only a single sizing parameter $y = W_{in}$. In addition to the results shown in Figure 5.16, this case was also run while setting $t2 = 0$ (i.e. only one mission segment was used). Multiplicity of solutions was also achieved under these conditions; it takes an extremely high initial guess for y to obtain this other converged value. Because of this, both the Newton-Raphson and successive substitution case only converged for a narrow range of x. This can be seen in Figure 5.17.

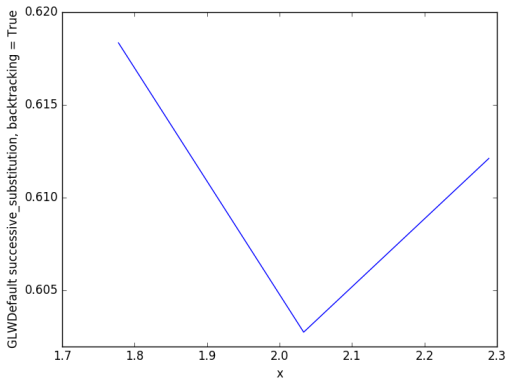

(a) Weight (Successive Substitution)

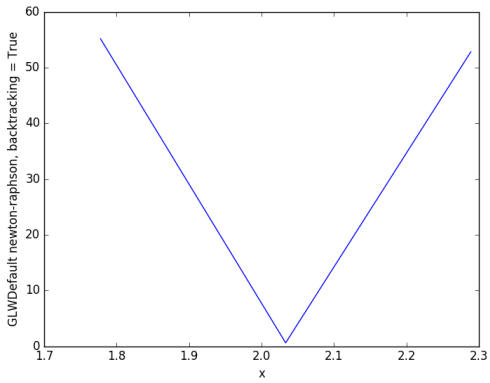

(b) Weight (Newton-Raphson)

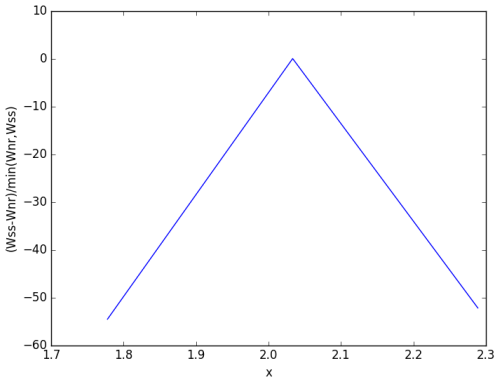

(c) Normalized Difference

Figure 5.17: $t2 = 0, y_{initial} = [50]$

As should be readily apparent, the discrepancy between the two converged systems is over an order of magnitude. It logically follows that producing the same multiplicity of solutions seen in Sections 5.2.1 and 5.2.2 becomes extremely difficult. Here, a global search algorithm (Particle Swarm Optimization, or PSO) was needed to determine values for $y_{initial}$ such that this multiplicity of results was achieved [50]. Note that the $y_{initial}$ (which, as a reminder, corresponds to $[\frac{W_{in}}{g*10^4}]$) is also extremely large compared to the previous three cases. The reason for this lack of convergence

is that, depending on the initial guess, the algorithm iterates to values for W_{in} such that the term $(1+\ln(\frac{W_{in}}{10{,}000g}))$ in (5.2) produces values for $f_{empty} > 1$ producing nonsensical results. The extreme simplicity of this case seems to make producing these "extraneous" solutions unlikely unless one is actively searching for them. Now, for the sake of completeness, the design space is plotted using $y_{initial} = 1$, to drive this point home in Figure 5.18.

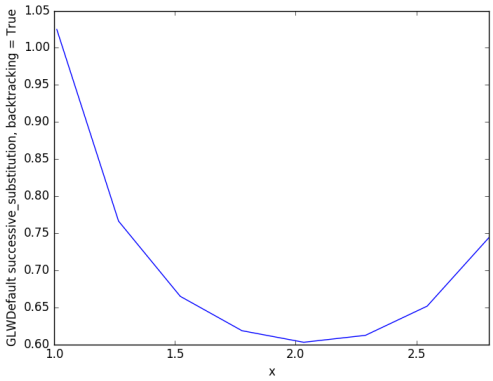

(a) Weight (Successive Substitution)

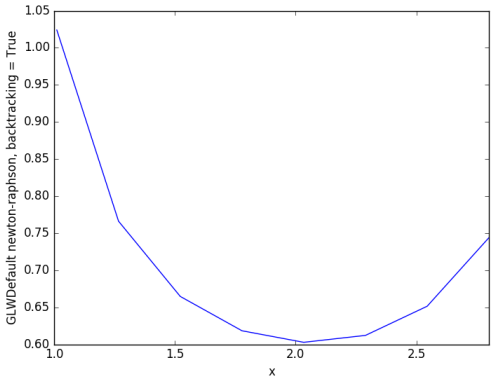

(b) Weight (Newton-Raphson)

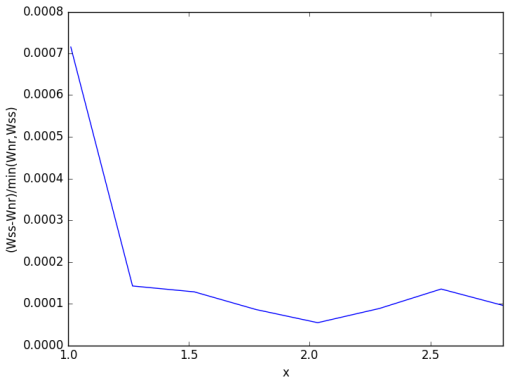

(c) Normalized Difference

Figure 5.18: $t2 = 0, y_{initial} = [1]$

Now, it has already been seen that producing these errant solutions becomes extremely difficult when the system is simplified to a sufficient degree; therefore, system complexity seems to be at least one factor. Another relevant question is what role, if any, does specific energy play in this phenomena? In other words, why are other, similar results not widely seen in looking at the conventional aircraft design process? And do the extremely low values for specific energy here play a role?

5.2.5 Electric Aircraft(Single Battery, no Power Requirements, Simplified Results (Higher Specific Energy))

Attempts at producing this sizing solution multiplicity will be tried using the system described in Section A.4 while varying the specific energy. To begin with, the parameter Esp is increased to 3,600, which corresponds to that of a Jet-A powered combustion system with a 30% overall energy conversion efficiency. Results for this, while setting $t2$ to 300 can be seen in Figure 5.19

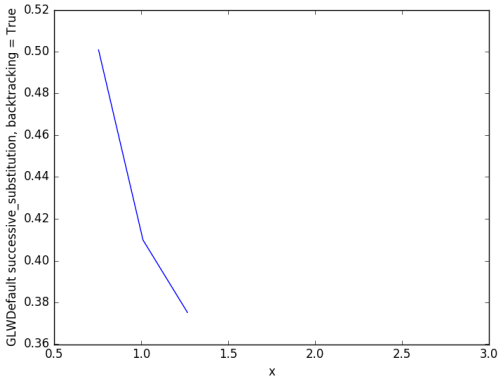

(a) Weight (Successive Substitution)

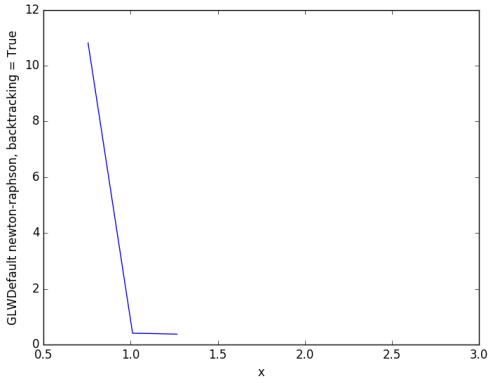

(b) Weight (Newton-Raphson)

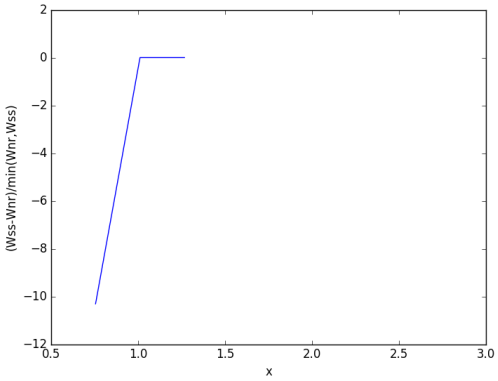

(c) Normalized Difference

Figure 5.19: $Esp = 3,600$, $y_{initial} = [10]$

Now, one will note that this case, like the one shown in Figure 5.17 only converges for a narrow range of values for x. This is because of the weight dependent term in (5.2) which causes the empty weight fraction to become greater than 1 in some cases. Finally, this same case is plotted with $t2 = 0$, as in Section 5.2.4.

Interestingly, Figures 5.19 and 5.20 are nearly identical at least qualitatively speaking. Because

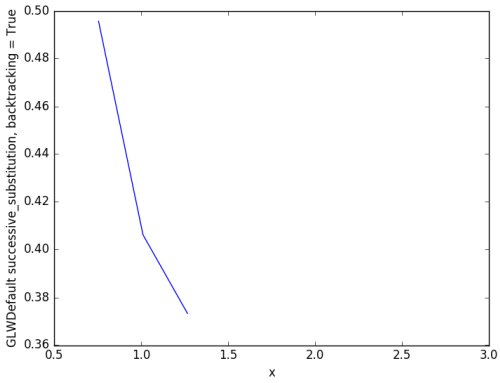

(a) Weight (Successive Substitution)

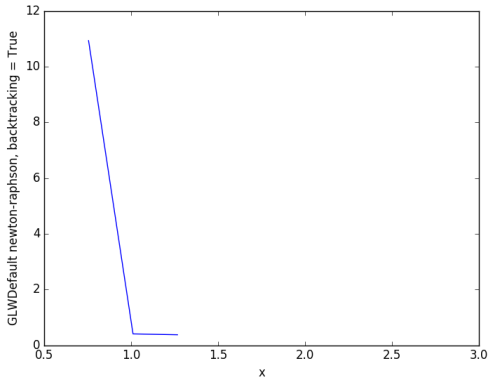

(b) Weight (Newton-Raphson)

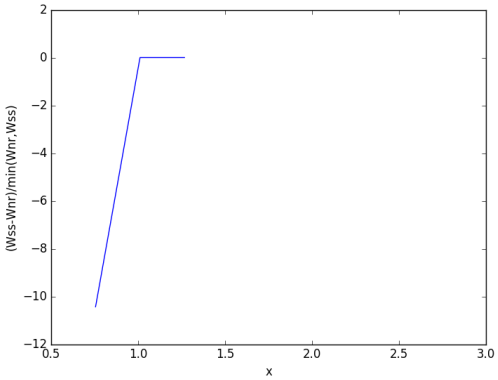

(c) Normalized Difference

Figure 5.20: $Esp = 3,600, t = 0, y_{initial} = [10]$

of the simplicity of the system, for constant mass, one can easily (and analytically) calculate that, for a fixed "airframe," the mission when $t2 = 0$ takes about 10% less energy than when $t2 = 300$ (the default case).

5.3 Summary of Important Results

Now, in producing this multiplicity of sizing solutions, several key features were observed. For example, adding more complexity in the system (including modeling power requirements in the sized aircraft), tends to make it much easier to find these additional sizing solutions. Finding additional solutions to the sizing residual problem becomes extremely difficult when power is unaccounted and $t2$ is set to 0 (i.e. only a cruise segment was used). Setting $t2 = 0$, for most of the cases (i.e. the single and hybrid batteries when power was taken into account) produced these nonunique sizing

solutions with relatively benign values for $y_{initial}$, indicating that this nonuniqueness property is not, fundamentally, because the aircraft is sized for multiple mission design points. Nonetheless, adding additional segments does make finding the extra solutions easier. Furthermore, in all four cases, separate phenomena were introduced and removed to diagnose other conditions that contribute. The most significant is that removing the weight dependence from f_{empty}, i.e., setting

$$f_{empty} = 1.1 - .8x + .2x^2, \tag{5.12}$$

produced cases where all sizing solutions appear to be unique (based on an exhaustive search in $y_{initial}$). Notably, this occurs for all four aircraft cases, not just for the cases where power requirements for the propulsion system went unmodeled. This suggests that allowing f_{empty} to vary with W_{in} creates additional opportunities for values of W_{in} that satisfy the weight sizing residual (see (A.5)). As an aside, modeling the empty weight fraction such that it is an implicit function of the total aircraft weight is typical of nearly all aircraft design programs/methods [63, 48, 87, 52, 80, 84, 95]. Thus, appearances indicate that there is no "hard or fast" rule that determines when these nonunique sizing cases exist (other than the extremely weak assumption that the empty weight fraction varies with GTOW). Rather, the inclusion of additional phenomena (such as power or multiple mission/sizing conditions) seems to introduce some "slack" into the system, allowing for the potential for regions where there are multiple solutions that set the sizing residual to 0. Based on the author's experience, these nonunique solutions also seem to often be in regions near to which the sizing loop diverges, although further study is recommended. Table 5.3 summarizes the nonunique sizing solutions encountered in this chapter for a single value of x (and for the hybrid battery cases, f_{aux}). Unless otherwise noted, $x = 1$ for all cases within Table 5.3.

Table 5.3: Analytical, Nonunique Sizing Solutions Summary

Aircraft Case	Converged y	Multiple Solutions?
Hybrid Battery	y_{ss} = [1.84, 38.7, .8, 13.3] y_{nr} =[4.2, 78.9, 1.63, 27.1]	Yes
Hybrid Battery (t2= 0) x = 1.8, f_{aux} = .4	y_{ss} = [0.61, 14.3, 1.58, 5.95] y_{nr} =[52.6, 695, 77.2, 290]	Yes
Single Battery	y_{ss} = [1.71, 37.0, 12.5]] y_{nr} = [4.53, 86.8, 29.3]	Yes
Single Battery ($t2 = 0$)	y_{ss} = [1.27, 26.0, 9.75] y_{nr} = [4.9, 94.4, 31.5]	Yes
Single Battery (no power requirements)	y_{ss} = [1.27, 29.1] y_{nr} = [6.83, 126.02]	Yes
Single Battery (simple) (no power requirements)	y_{ss} = [1.27] y_{nr} = [6.84]	Yes
Single Battery (simple, $t2$=0) (x = 1.8 no power requirements)	y_{ss} = [.616] y_{nr} = [56.3]	Yes
Single Battery (simple) (x = .75, Esp = 3,600, no power requirements)	y_{ss} = [.504] y_{nr} = [10.6]	Yes
Single Battery (simple, $t2 = 0$) (Esp = 3,600, no power requirements)	y_{ss} = [.499] y_{nr} = [10.7]	Yes
Hybrid Battery (no f_{empty} weight dependence)	y_{ss} = [1.55, 33.5, 0.69, 11.5] y_{nr} =[1.55, 33.4, 0.69, 11.5]	No
Single Battery (no f_{empty} weight dependence)	y_{ss} = [1.51, 33.3, 11.2] y_{nr} = [1.51, 33.3, 11.2]	No
Single Battery (simple) (no power requirements) (no f_{empty} weight dependence)	y_{ss} = [1.29] y_{nr} = [1.29]	No
Single Battery (simple, $t2$=0) (no power requirements) (no f_{empty} weight dependence)	y_{ss} = [1.13] y_{nr} = [1.13]	No
Single Battery (simple, $t2$=0) (no power requirements, Esp = 3,600) (no f_{empty} weight dependence)	y_{ss} = [0.48] y_{nr} = [0.48]	No

As Table 5.3 illustrates, these nonunique sizing solutions exist across a broad range of aircraft design problems. Additionally, it was found that, setting $t2 = 0$ tended to make finding these nonunique sizing solutions much more difficult, as the extremely large disparity between the converged cases for the "Hybrid Battery ($t2$= 0)" instance suggests. Interestingly, when not sizing the aircraft energy system based on power requirements, finding these undesirable "extra" solutions becomes much more difficult. Now, it should be noted that in all cases presented here, the larger solution was only found using the Newton-Raphson algorithm; successive substitution either converged to the smaller value or diverged to infinity. An explanation as to why this is the case can be seen in Section 4.4. Finally, when taking (5.12) instead of (5.2) to define the empty weight fraction

of the aircraft, none of these "spurious" solutions to the sizing residual equations were found.

5.3.1 Banach Fixed Point Iteration

Now, some of these unusual results may be at least partially explained by an application of the Banach Fixed Point Iteration Theorem to the problem here [9]. To begin with, this theorem uses a complete metric space (Y, d) where Y is the range of the iterating variables and d is some function that measures the distance between two points. Then, some operator, T, is used (which in this case is an aircraft evaluation), which returns some value for the input parameter y. Existence and uniqueness of a fixed point of this map T using this metric space is guaranteed if, for all $y1$, $y2$ in Y, that, for some q between 0 and 1

$$d(T(y1), T(y2)) \leq q * d(y1, y2). \tag{5.13}$$

Note that, to compare this to the notation used throughout this report, $T(y)$ is a function that returns the value y_{out}. A map (or operator) T is called a contraction mapping if it possesses the property shown in (5.13). The aircraft design process used in all four cases in this chapter, and indeed, nearly all aircraft design sizing loop processes does not possess this property, i.e. is not a contraction mapping, so uniqueness, or even existence of a fixed point is not guaranteed. This can be readily shown by taking an extreme case often seen from the results here; in many instances, beginning with an extremely high value for y will return increasingly high values for y_{out}, eventually diverging to an infinite value. In the terminology shown in (5.13), a given value of $T(y1)$ or $T(y2)$ may be infinite for a finite value of $y1$ or $y2$, meaning that the inequality cannot hold for any finite value of q. For nearly any aircraft design process using this sizing loop methodology, this capability of returning an infinite value is a distinct possibility, even for modest initial values of $y1$ or $y2$. This tendency to diverge is further exacerbated by the low values of specific energy, and, in some cases, specific power that are used in electric and hybrid-electric aircraft. As an aside, Banach's Fixed Point Iteration Theorem is a 1-dimensional theorem; multidimensional extensions exist, but they are much more complicated. For the morbidly curious, Brouwer's Theorem extends this proof of uniqueness to multiple dimensions, while Schauder's Theorem further extends it to infinite dimensions [31]. Regardless, the 1-D case should be sufficient to demonstrate the relevant properties here, i.e. that uniqueness to aircraft design problems such as these cannot, in fact, be proven.

5.4 Hybrid Battery Optimization Results

Now, in addition to the results exploring some of the fundamental properties of these types of systems, some optimization cases were investigated for the hybrid battery aircraft. Summaries of said optimization results for two types of optimization problems, using various methods to size

the aircraft can be seen in Tables 5.4 and 5.5. Note that, in this context, n_{size} refers to the number of aircraft evaluated for a value of (x,f_{aux},y), while n_{opt} refers to the number of design decisions made, i.e. changing values of (x, f_{aux}), including finite difference evaluations. The term "% threshold" refers to the point in terms of norm(h) at which the Jacobian is first constructed using a finite differencing technique, wherein Newton-Raphson is used to converge the system until a given tolerance value is reached (in this case 1E-4). At higher values for norm(h) (i.e. above this threshold), successive substitution is used to determine new values for y until the system becomes appropriately converged. This is because starting Newton-Raphson at highly unconverged cases introduces a tendency to diverge to either nonsense or infinite values very quickly, a tendency which should be managed. Details as to what the "Initial Step" column entails can be seen in Chapter 7. For reference, KNN(5) refers to a K Nearest Neighbors algorithm that looks at up to the 5 closest converged cases to determine an initial guess for y based on already converged results for x and f_{aux}. Optimizer decomposition refers to a reformulation of the problem detailed in Section A.3 that tends to be "nicer" when constraints are enforced by an optimizer rather than solved by an internal loop. By definition $n_{size} = n_{opt}$ for the Optimizer Decomposition case, as the optimizer "sees" every single iteration of the problem. Result for the local optimization case using SNOPT can be seen in Table 5.4 [37]. In the Optimizer Decomposition case, the finite differencing step size is set to 1E-6 while in the Sizing Loop cases (everything else), the finite differencing step size is set to 1E-2.

Table 5.4: Optimization Results (Hybrid Battery Analytical Case, $y_{initial} = [1, 1, 1, 1]$, SNOPT, initial $[x, f_{aux}] = [1.5, .5]$

Initial Step	Sizing Loop Solver	n_{size}	n_{opt}	$\frac{W_{in}}{g*10^4}$	$[x, f_{aux}]$
Constant	successive substitution	1,278	86	.63	[2.18, 0.54]
Optimizer Decomposition	N/A	135	135	.62	[2.08, 0.37]
KNN (5)	successive substitution	118	16	.67	[1.86, 0.84]
KNN (5)	Newton-Raphson (5% threshold)	1,278	86	.63	[2.18, 0.54]
KNN (5)	Newton-Raphson (2% threshold)	855	59	.64	[2.19, 0.56]

As can be seen in Table 5.4, the objective function ($\frac{W_{in}}{g*10^4}$) is relatively flat, with (nearly) equivalent local minima depending on the choice of x and f_{aux}. Additionally, one can readily see a wide disparity in terms of both n_{size} as well as n_{opt} for the various sizing loop optimization cases. In addition, for this case, the optimizer decomposition approach appears to outperform the loop methodology used here. Optimization results using the particle swarm algorithm ALPSO can be seen in Table 5.5 [50].

It should be noted that ALPSO is a stochastic method, so the results in Table 5.5 are not directly reproducible, at least not to 100 % accuracy. However, these methods were run by the author dozens of times over, so the results, while not perfectly replicable by a third party, are representative of what is typically obtained by performing these optimization using the indicated

Table 5.5: Optimization Results (Hybrid Battery Analytical Case, $y_{initial} = [1, 1, 1, 1]$, ALPSO)

Initial Step	Sizing Loop Solver	n_{size}	n_{opt}	$\frac{W_{in}}{g*10^4}$	$[x, f_{aux}]$
Constant	successive substitution	52,414	1,480	.62	[2.07, 0.37]
Optimizer Decomposition	N/A (aircraft unconverged)	48,040	48,040	1.24	[2.5, 0.2]
KNN (5)	successive substitution	29,917	1,480	.62	[2.08, 0.37]
KNN (5)	Newton-Raphson (5% threshhold)	33,138	1,480	.62	[2.07, 0.37]
KNN (5)	Newton-Raphson (2% threshhold)	23,728	1,240	.62	[2.08, 0.37]

methods. Additionally, it should be apparent from a comparison of Tables 5.4 and 5.5 that the use of global optimization algorithms comes with a significant computational cost, as the least computationally intensive case in Table 5.5 is still 200x greater than the least expensive case in Table 5.4. On the other hand, all optimization cases (with the notable exception of Optimizer Decomposition) converge to the same result, indicating that function flatness is not an issue for this global optimizer. Perhaps more importantly, it is readily apparent that, because the particle swarm algorithm does not possess the disciplined constraint handling of say, an SQP method (see Section 1.4.1), the optimizer decomposition approach does not perform adequately here due to the relatively complicated interplay of the energy system constraints.

Chapter 6

Aircraft Cases

Four separate aircraft classes will be investigated here, with a variety of different sizing techniques employed. The aluminum-air aircraft and the diesel-electric aircraft in particular are highlighted to demonstrate some of the potential algorithmic and numerical challenges in sizing and designing aircraft that employ more than one energy system. For each aircraft case, the fixed point iteration methods from Chapter 4 will be applied, comparing speed and robustness. In addition, the effectiveness of various machine learning techniques to increase code speed according to the methodology outlined in Chapter 7 will also be explored. A hypothetical problem setup of each aircraft using optimizer decomposition approaches is outlined, for the sake of completeness, as well as illustrating that a variety of valid, equivalent system setups may be employed in designing said aircraft. The regional jet and the lithium-air jet were used to compare the effectiveness of these techniques when applied to aircraft with simpler energy systems.

6.1 Regional Jet (Conventional)

For an aircraft sized for a single mission, this aircraft can close with a simple sizing variable of $y = GTOW$, assuming that variables such as jet engine thrust and wing area are either set as a constant or normalized by the $GTOW$. This conventional case will be used as a benchmark, and is most representative of aircraft in service today. The propulsion system consists of two turbofan engines powered by Jet-A. The aircraft itself can carry 125 passengers and has a cruise range of 2000 nautical miles. A drawing of this conventional aircraft using OpenVSP can be seen below in Figure 6.1 [33].

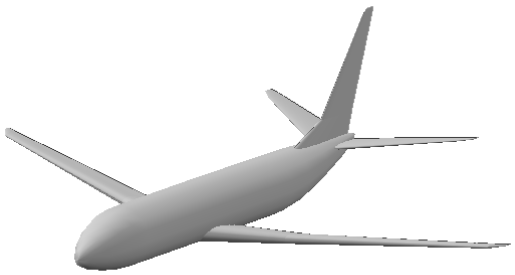

Figure 6.1: Conventional Turbofan-Powered Aircraft

Here, a "guess" for the gross takeoff weight ($GTOW$) is assigned, and fuel weight carried by the aircraft falls out from the gap between gross takeoff weight and payload and empty weight, as shown in (6.1).

$$W_{fuel} = GTOW - (W_{empty} + W_{payload}) \tag{6.1}$$

Empty weight is computed based on a component weight buildup of the fuselage, wing, tail, propulsion system and other associated systems, while payload weight is maintained as a constant. Due to the iterative nature of the mission setup, the actual fuel burned by the aircraft throughout ($W_{fuel,out}$) is calculated and an associated output weight computed.

$$W_{out} = W_{fuel,out} + W_{empty} + W_{payload} \tag{6.2}$$

A residual is then calculated to determine whether the aircraft is converged.

$$h = \frac{GTOW - W_{out}}{GTOW} \tag{6.3}$$

If one remembers the framework from the generalized case in (1.17)-(1.21), it will be immediately clear that the system here is identical. Due to its simplicity, it is relatively easy to converge using a successive substitution approach, where W_{out} is assigned as a new guess for $GTOW$ until $W_{out} \approx GTOW$. A plot of the computed residual $h(y)$ (labeled $R(y)$ in Figure 6.2) using said method shows a log-linear relationship convergence rate as shown in Figure 6.2.

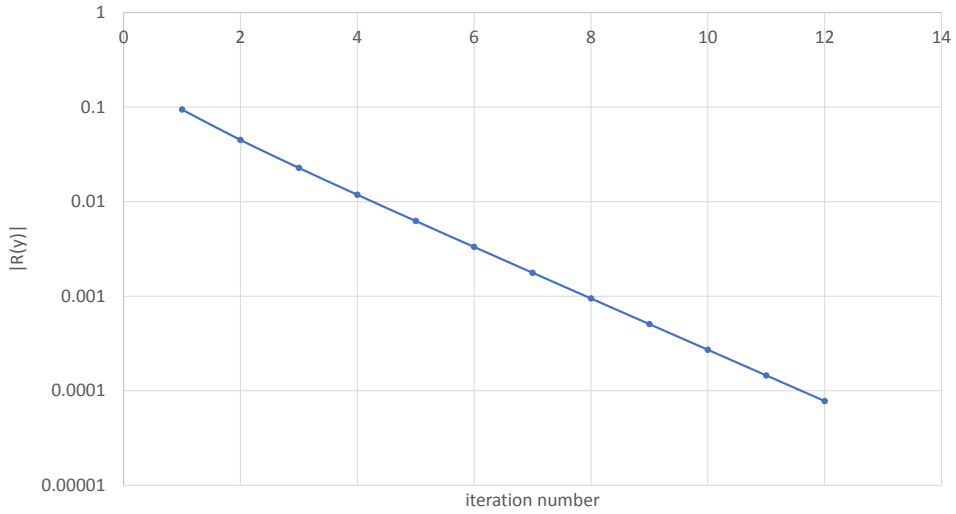

Figure 6.2: Convergence Results for Conventional Aircraft

6.1.1 Optimizer Decomposition in Conventional Aircraft

Optimizer decomposition is a technique that has been widely used in the design and optimization of conventional, gas-turbine powered aircraft. As the simplest case, its application here is relatively straightforward. $GTOW$ may easily be assigned by the optimizer. Empty weight and fuel weight are computed as before, and the associated constraint becomes

$$W_{fuel,constraint,opt} = \frac{W_{fuel} - W_{fuel,out}}{C_{fuel}} \tag{6.4}$$

where C_{fuel} is a constant chosen such that the value of $W_{fuel,constraint}$ is of the same order as the rest of the constraints assigned for the optimization problem (often of order 1). Note that (6.4) differs from (1.21) in that the fuel mass is normalized by a constant, rather than by the $GTOW$. Additionally, it may be used as an inequality constraint, which is numerically easier to solve. This makes the problem more numerically stable, especially when gradients are calculated. On the other hand, for cases where the designer knows little about the aircraft, a poorly-chosen value for C_{fuel} may be inadvertently used, which can lead to numerical issues in the optimization process (such as making the Hessian close to singular, reducing its usefulness).

6.2 Regional Jet (Lithium-Air)

Present-day battery technology does not allow for the development of commercial-scale all-electric aircraft, due to limitations in the specific energy, and in some cases specific power of lithium-ion batteries. Nonetheless, lithium-air batteries have projected specific energies ranging from 1,000-2,000 W-h/kg, and specific powers from .4-.67 kW/kg [51, 92]. In addition to energy constraints, one has to account for power constraints in designing electric aircraft. Lithium-air batteries, for instance, are expected to have a relatively low specific power, on the order of .67 kW/kg. The feasibility of these designs can be significantly improved by shaping both the flight profile and the wing to account for this. Additionally, the aircraft grows heavier throughout the mission based on the parameters seen in (1.2) and (1.3). The sizing variables for this configuration can be seen below, along with a CAD drawing of an example aircraft (with a box representing the estimated required lithium-air battery volume) in Figure 6.3.

$$y = [W_{in}, E_{bat}, P_{in}] \tag{6.5}$$

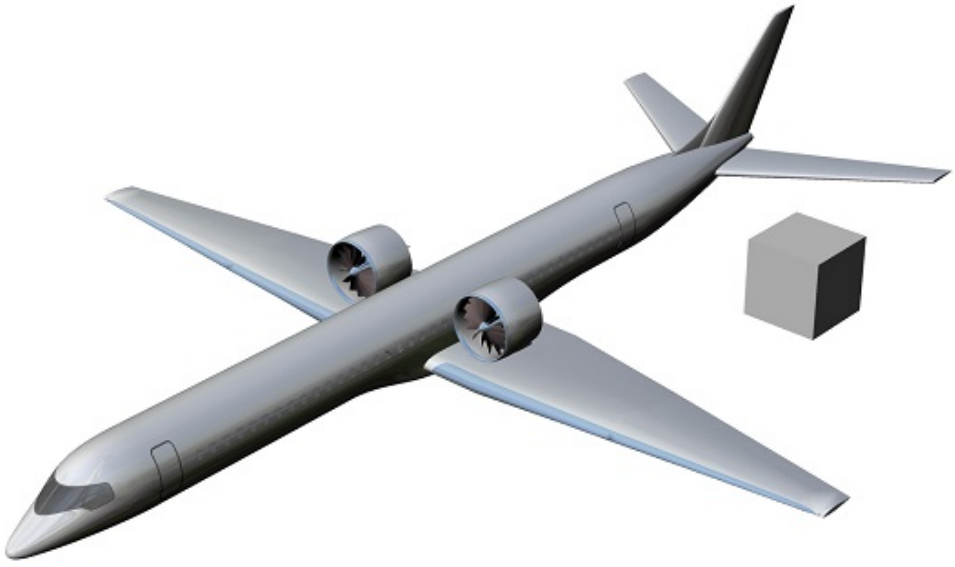

Figure 6.3: Lithium-Air Aircraft

Now, from (6.5), one can define the battery weight based on the maximum of the energy requirements or the power requirements, i.e.

$$m_{bat} = \max(\frac{Ebat}{Esp}, \frac{P_{in}}{Psp}) \tag{6.6}$$

Now, in most numerical algorithms, discontinuities tend to result in difficulties when standard methods are used to solve them. To that end, the "soft_max" function was used, which is defined in (5.4)-(5.6). The aircraft evaluated here is a regional passenger jet, with a design payload of 114 passengers and design range of about 2,400 nautical miles. Now, a plot of the convergence of the residuals for mass, energy, and power for a representative design can be seen in Figure 6.4.

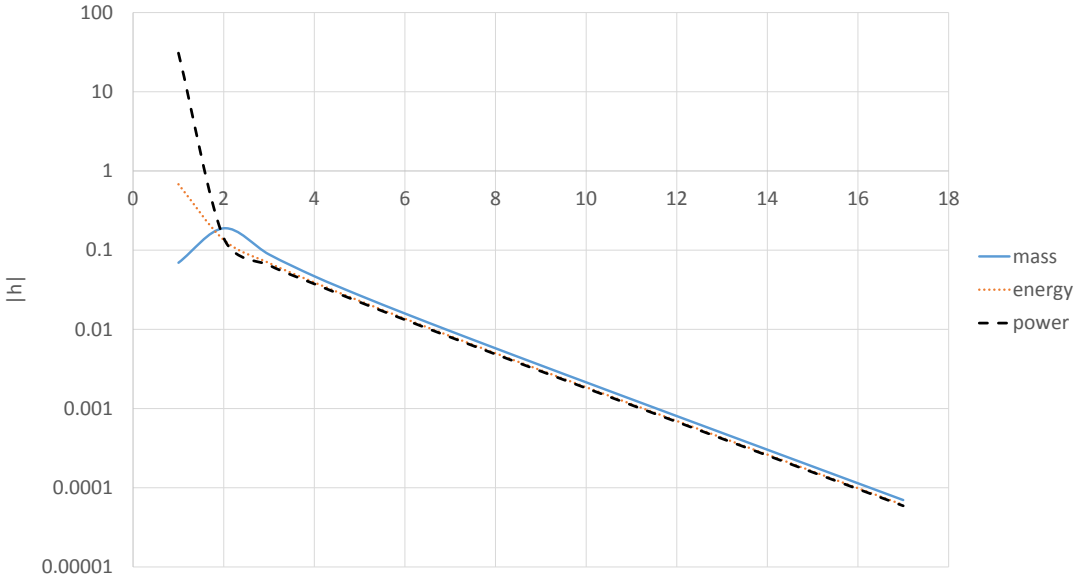

Figure 6.4: Convergence Results for Lithium-Air Aircraft

Note that, like the conventional, fossil-fuel powered aircraft, the residuals all show log-linear convergence. An additional point of interest is the fact that these aircraft become heavier as they fly. To illustrate, Figure 6.5 shows the vehicle mass, altitude, power, and state of charge plotted vs. mission time to see how they vary.

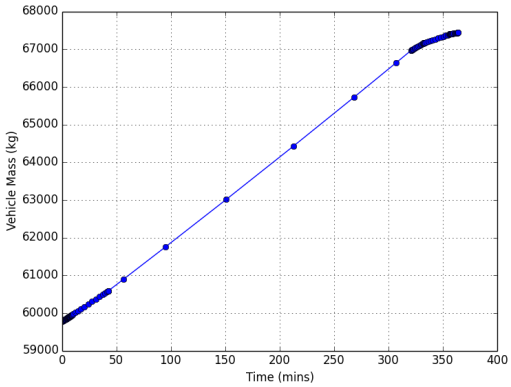

(a) Vehicle Mass

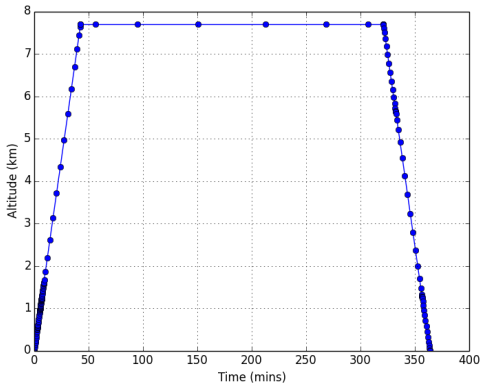

(b) Altitude

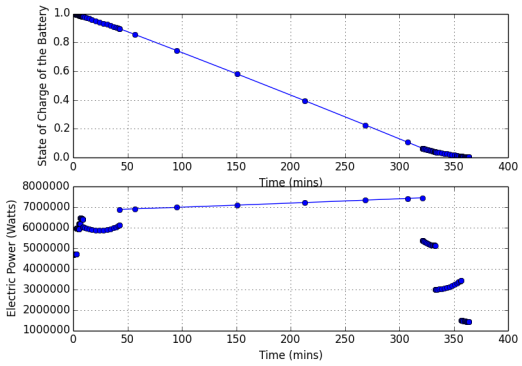

(c) Power

Figure 6.5: Lithium-Air Jet Mission Outputs

In this case, the aircraft grows about 10% heavier as it flies the mission, while the battery is completely drained at the end of the mission. A more realistic design study would ensure some buffer in state of charge at the end, (often about 20%) to ensure that the battery maintains capacity after cycling. More details into the physics of the aircraft can be seen in previous work [98].

6.2.1 Optimizer Decomposition in Metal-Air Aircraft

Now, in many cases, the formulation for defining battery weight shown in (6.6) is a poor choice when attempting to use an optimizer to simultaneously optimize and size the aircraft. While the soft max formulation shown in (5.4)-(5.6) smooths out all discontinuities, it may not necessarily be the best approach. Instead, it is often more computationally efficient to split the max function and use multiple constraints. This necessitates the assignment of battery mass directly, i.e. add m_{bat} as a design variable, and ensure that the actual energy and power of the battery meet all mission

requirements, i.e.

$$E_{bat} = m_{bat} * Esp \tag{6.7}$$

$$P_{bat} = m_{bat} * Pbat \tag{6.8}$$

$$E_{constraint} = \frac{E_{bat} - E_{mission}}{C_{energy}} \geq 0 \tag{6.9}$$

$$P_{constraint} = \frac{P_{bat} - P_{mission}}{C_{power}} \geq 0 \tag{6.10}$$

where C_{energy} and C_{power} are constants chosen to ensure that the values of the constraints themselves are of the same order of magnitude as the other constraints. In addition, for a metal-air battery, one must add another design variable, such as $E_{discharge}$, to account for the mass change of the aircraft and size the relevant components (such as the landing gear). Now remembering from (1.2), the mass gain rate of the battery based on the power requirements of the mission can easily be calculated. One can easily integrate (1.2) to determine the overall mass accumulated by the battery throughout the mission. As a result, another associated constraint needs to be applied to ensure that the components sized prior to running the mission can meet the mission requirements. This requirement is shown in (6.11).

$$E_{constraint,mass} = \frac{E_{discharge} - E_{mission}}{C_{energy}} \geq 0 \tag{6.11}$$

6.3 Hybrid Battery Aircraft

The most-complicated case is an aluminum-air powered aircraft, augmented with a lithium-ion battery for high-power operation. Aluminum-air batteries are a recent topic of interest due to their impressive specific energy characteristics (∼1,300 W-h/kg) and high technology readiness level, but possess poor specific power (∼.2kW/kg). The given specific energy and specific power are representative of batteries produced in 2002 [102].

Any aircraft powered by aluminum-air batteries needs to carry the necessary water supplies, which must be accounted for in the sizing of the vehicle (see (1.4) in the Introduction). Practically, aluminum-air batteries would function as primary batteries without some significant technological breakthrough, and thus, would not be rechargeable. Therefore, they would have to be either recycled onsite or shipped to some processing facility. Additionally, there is still a significant specific energy penalty vis. a vis. fossil fuels, which results in larger aircraft than turboprops of comparable performance. Furthermore, their poor specific power necessitates the incorporation of less energy dense, but more power dense lithium-ion batteries for high-power operation (such as climb), which results in additional design complexity, not just in terms of the physical aircraft, but also system

modeling.

This is because when defining these systems and simulating the mission, the user must decide on the "rules" as to how each battery operates (this was avoided in Chapter 5 because of the simplistic nature of the mission/sizing problem). Here, because the specific energy of the aluminum-air (or "primary") battery is much higher, this battery runs until its maximum power output is reached, at which point the auxiliary battery "kicks" in and meets whatever additional power output is needed. Some pseudocode can be seen for this in (6.12) and (6.13)

$$\begin{aligned} &\text{If } P_{mission} \leq P_{bat,prim,max} \\ &P_{bat,prim} = P_{mission} \end{aligned} \tag{6.12}$$

$$\begin{aligned} &\text{Else If } P_{mission} > P_{bat,prim,max} \\ &P_{bat,prim} = P_{bat,prim,max} \\ &P_{bat,aux} = P_{mission} - P_{bat,prim,max} \end{aligned} \tag{6.13}$$

where $P_{bat,prim}$ refers to the power drawn from the aluminum-air battery at a given point in time, and $P_{bat,aux}$ similarly refers to the power drawn from the lithium-ion battery in the same instance. Now, in most aircraft missions the highest-power operation is during takeoff, which, for a minimum-weight aircraft, would likely leave the auxiliary battery drained. In these cases, the primary battery would charge the auxiliary battery, as shown in (6.14)

$$\begin{aligned} &\text{If } P_{mission} < P_{bat,prim,max} \\ &P_{bat,aux} = -(P_{bat,prim,max} - P_{mission}) \end{aligned} \tag{6.14}$$

where, by convention, the "$-$" sign in front of the $(P_{bat,prim,max} - P_{mission})$ term indicates that the battery is being charged. The sizing variable vector y for this configuration can be seen in (6.15).

$$y = [W_{in}, E_{prim}, E_{aux}, P_{in}] \tag{6.15}$$

The aluminum-air design possesses the most complex propulsion system investigated here, and in addition, was thought to be the most representative of the power/energy tradeoffs of next-generation hybrid-electric vehicles. Indeed, this aircraft was originally chosen as a "toy" problem before tackling a hybrid-electric aircraft where off-design performance is more difficult to model. However, in investigating this design, it was found that this aircraft class possesses a number of interesting characteristics in its own right, so it will be explored somewhat more extensively than the other cases. Now, the choice for problem formulation shown in (6.15) is relatively arbitrary. For instance,

one suggested approach is to only iterate on total energy, rather than on the energy of each energy system. However, this is likely to converge to physically unrealizable designs; the auxiliary battery often only consumes a small amount of the total energy, yet accounts for a substantial fraction of the power; thus, using total energy as a sizing variable results in highly infeasible aircraft, and needs to be handled separately. Another suggested approach is to iterate on the mass of each subsystem, rather than energy. This approach has its merits, but lacks the structure needed to make use of Ragone-plot optimums as the cases here do (see Figure 2.4). Nonetheless, they are a suggested avenue of future work. Payload corresponds to 30 passengers with a cruise speed of 300 kts. Unless otherwise specified, range is set to 500 nautical miles. A CAD image of an example aircraft from a previous paper can be seen in Figure 6.6 (which includes the required volumes of the energy subsystems within the wing, where the water is blue, the the lithium-ion battery is red, and the aluminum-air battery is green) [97].

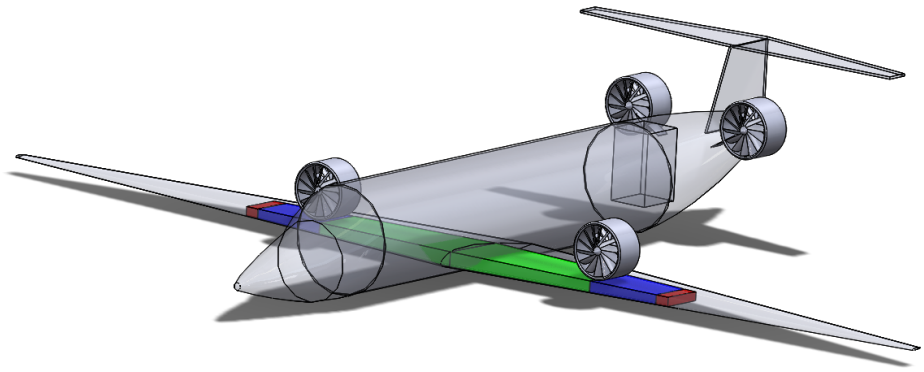

Figure 6.6: Aluminum-Air/Lithium Ion Powered Aircraft

Additionally, in defining each propulsion system, one must first determine the relative size of the two energy sources. In hybrid cars, this is variously called an "electrification factor," or "hybridization factor," among others, and like in Chapter 5 is defined as

$$f_{aux} = \frac{P_{aux}}{P_{in}} \tag{6.16}$$

where P_{aux} is the maximum power output that the auxiliary power system (in this case, the lithium battery) can output at a time, while P_{in} refers to the maximum power that the entire energy system can deliver. Here, f_{aux} is also called the "auxiliary power fraction." Now, as stated previously, the use of more than one energy system introduces additional numerical complexity when attempting to converge the system. To illustrate, Figure 6.7 compares the value of the sizing residuals at each iteration of an aircraft converged using successive substitution.

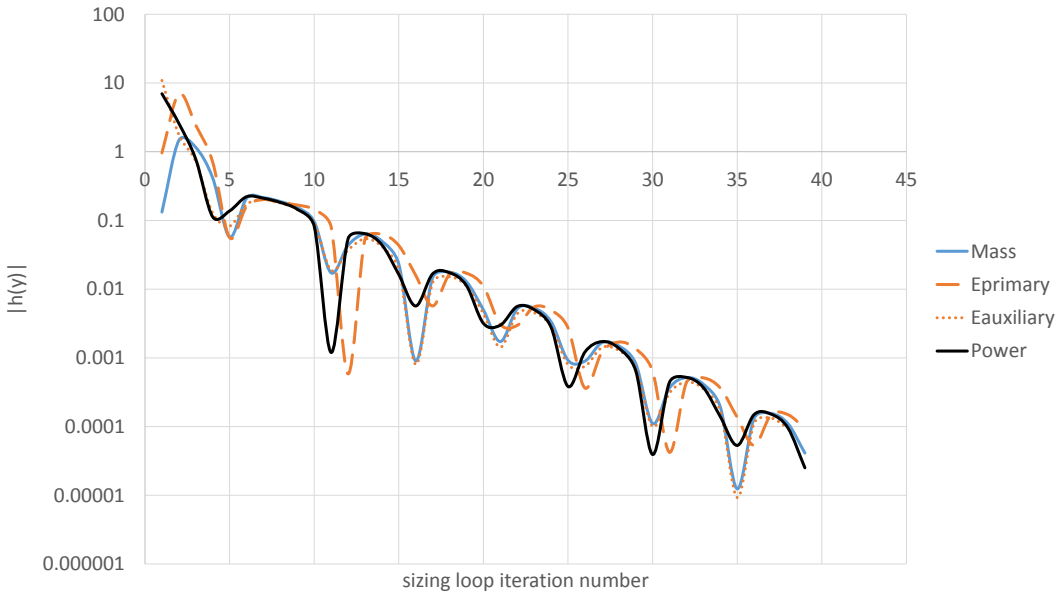

Figure 6.7: Convergence Results for Higher-Dimensional y

Unlike the previous two aircraft cases seen in Figures 6.2 and 6.4, the convergence plot here is not nearly as "clean," i.e. it does not resemble the log-linear plots in the previous sections. From a numerical standpoint, as the solution proceeded (for fixed design variables), there was significant oscillation in the error; furthermore, the different sizing errors were out of phase with each other, which significantly slowed the convergence rate. Furthermore, to converge all of the required parameters, some of the other parameters may become over-converged, which, due to the higher range factor, may affect the design enough to result in a serious reduction in gradient accuracy.

Physically, what is happening is, based on the guess for y, a given $y_{out} = [W_{out}, E_{prim,out}, E_{aux,out}, P_{out}]$ is returned. Now, in the absence of the other energy system, $E_{prim,out}$ would return the required energy for the mission; however, in conjunction with the returned value $E_{aux,out}$, overshoot occurs in each system. To put it more concretely, say $E_{prim,out} > E_{prim}$, and $E_{aux,out} < E_{aux}$. Using $E_{prim,out}$ and $E_{aux,out}$ as new guesses for the primary and auxiliary battery energies will cause the

primary battery to have too much energy and the auxiliary battery too little energy. Now, it should be noted that the weight of the converged aircraft, and by extension, shape of most likely objective functions vary markedly depending on the range and specific energy/power of each energy system. This can be seen in Figure 6.8. When set to a constant value, $Esp_{prim} = 1{,}300\ \frac{W-h}{kg}$, $Psp_{prim} = .2\ \frac{kW}{kg}$, $Esp_{aux} = 300\ \frac{W-h}{kg}$ and $Psp_{aux} = 1\ \frac{kW}{kg}$. The "Opt" plots use the Ragone plots shown in Figure 2.4 to determine the minimum weight battery that meets the required energy and power based on the given y and f_{aux}. Note that the lithium-sulfur battery was chosen here as representative of a futuristic lithium-ion battery. Figure 6.8 displays a 1-dimensional sweep of landing weight vs the "f_{aux}" term for a 500 nautical mile design with fixed wing loading, thrust loading, and fan pressure ratio as shown.

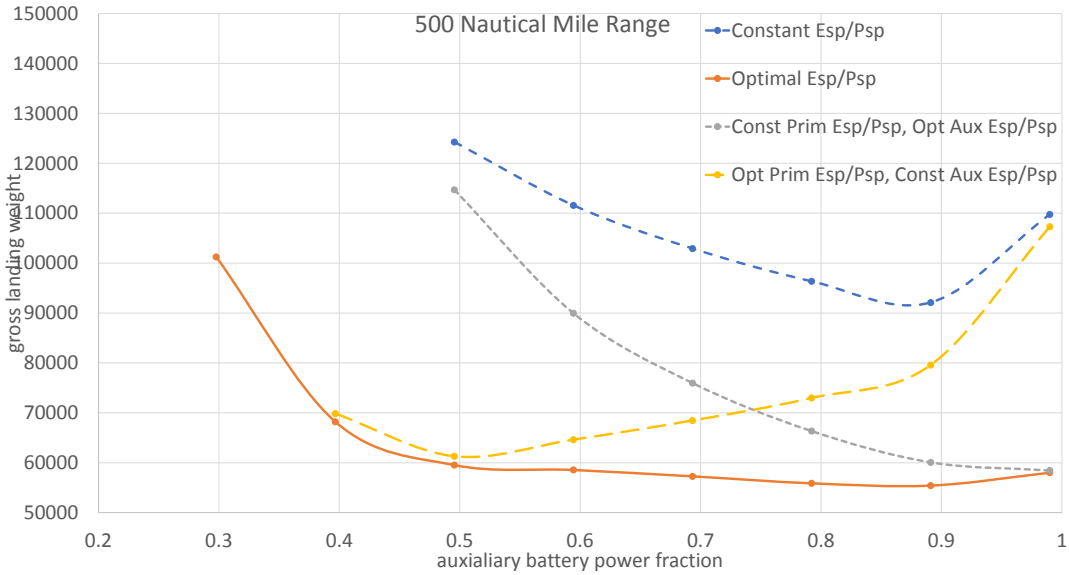

Figure 6.8: Ragone-Optimal vs. Constant Specific Energy Aircraft

Optimizing the specific energy/power of both battery systems results in a very flat optimum for this particular range. In addition, note that, when taking into account discharge properties rather than sizing the aircraft based on nominal specific energies and specific powers, the design space appears completely different. When comparing the weight of the aircraft where both energy systems account for off-design discharge capabilities to the single-point discharge cases, the single-point discharge were at times twice as heavy. This illustrates that very significant performance penalties may be experienced when not accounting for the specific energy/power variation at other discharge points. Thus, accurate modeling of battery performance is of paramount important in the

earliest stages of these and similar aircraft. Another complication is that as each battery switches from being energy to power limited (or vice versa), discontinuities happen in the derivatives of the aircraft weight which, depending on the way the problem is set up, may cause serious issues in an optimization problem. Table 6.1 displays the four separate cases where these discontinuities can occur.

Table 6.1: Multiple Battery Case Table

Case Number	primary battery	auxiliary battery
1	energy limited	energy limited
2	power limited	energy limited
3	energy limited	power limited
4	power limited	power limited

Practically, there are a number of reasons why one would want to design an aircraft with this energy system configuration. Because of the ability to tailor the energy system via changing the relative size of each battery, they can enable significant weight savings. This has the potential to enable all-electric aircraft at some measure of range that would be impossible with lithium batteries alone, which is further explained in Chapter 2. To illustrate, a set of four pareto fronts for aircraft optimized for weight, cost, as well as an additional set of curves for when the field length constraint is relaxed are plotted in Figure 6.9 vs. design range. The *GTOW* of an Embraer 120 (an aircraft with a similar seating arrangement) is also plotted for reference.

As Figure 6.9 illustrates, despite significant penalties in specific energy, the combination of the two energy systems allows for the design of all electric-aircraft at commercial ranges without prohibitive weight penalties. Additionally, depending on the objective function, composition of the propulsion system, and in some cases, overall weight of the aircraft varies markedly. Below a certain range, it is generally cheaper from an operating cost standpoint to design an aircraft using lithium batteries as long as the system closes robustly.

In addition to the environmental benefits associated with these all-electric aircraft, these aircraft have the potential to become cost competitive with modern transport aircraft. Figure 6.10 illustrates that these aircraft, while suffering from a significant weight penalty over more traditional turboprop aircraft, can offer substantial economic benefits, even when including the cost of recycling the aluminum.

Outside of the lithium-dominated ranges ($f_{aux} \approx 1$), the aircraft operating cost/passenger-mile curves for the most part scale flatly vs. design range. The exception is the weight optimal, longer field length curve, which shows favorable scaling at longer ranges. This suggests that these aircraft may be economically profitable at a variety of design ranges. Thus, one could potentially design an aircraft for a longer range, while replacing the battery system at the airport with whatever combination of aluminum-air and lithium batteries minimizes the operating cost for that particular mission. A plot of the operating cost when the aluminum is purchased for each flight rather than

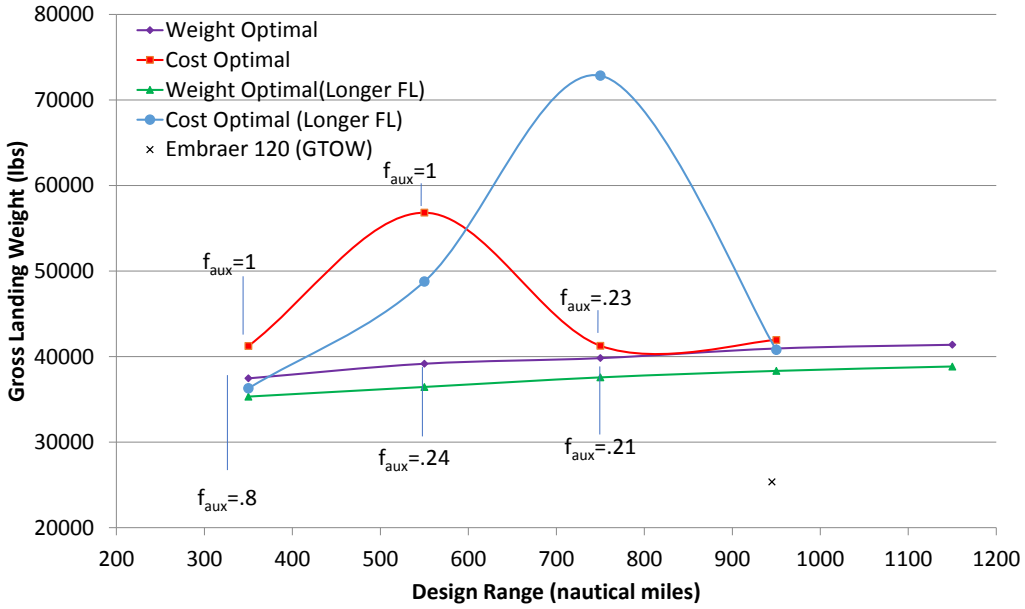

Figure 6.9: Weight Pareto Fronts

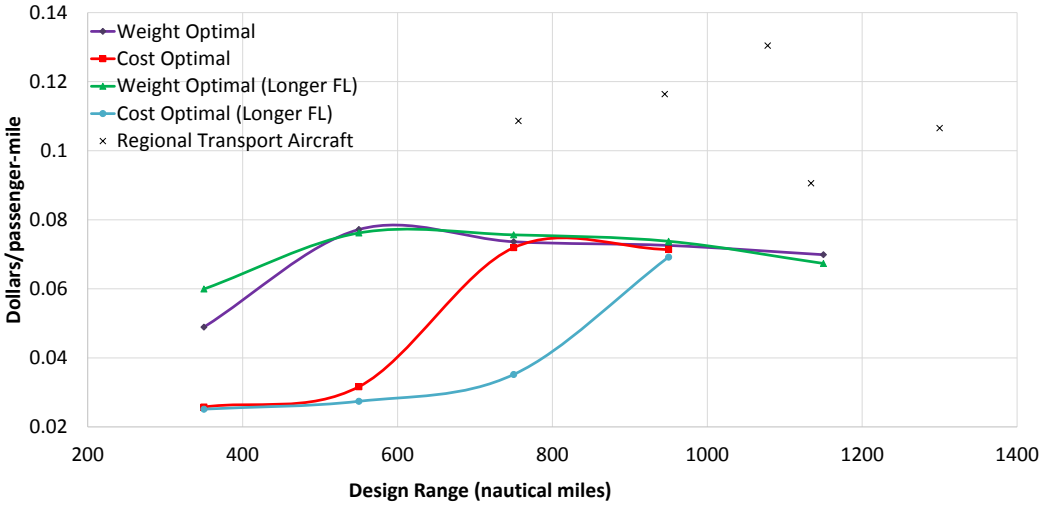

Figure 6.10: Cost per Passenger Mile (With Recycled Aluminum)

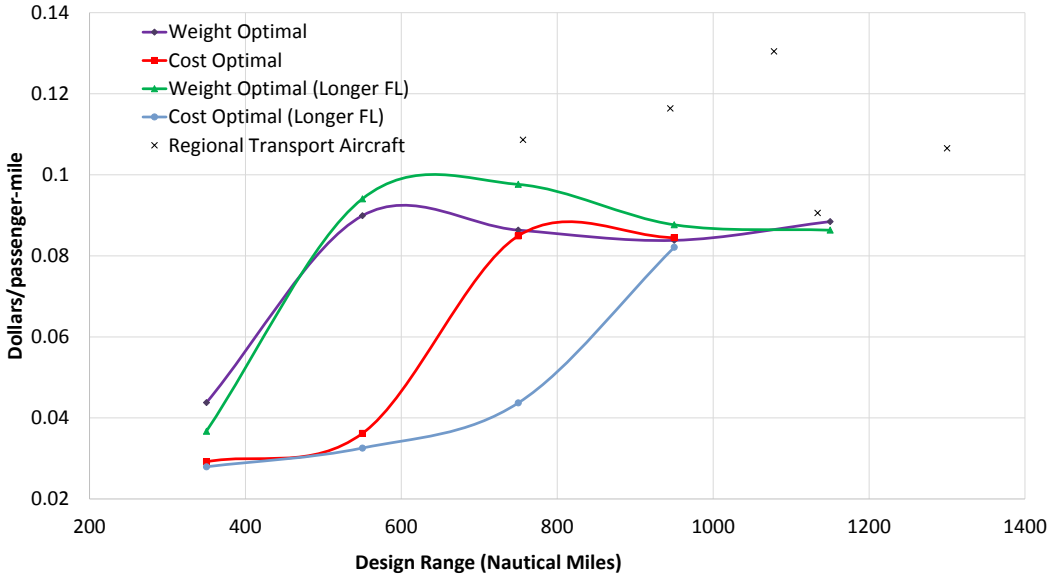

Figure 6.11: Cost per Passenger Mile (No Recycling)

recycled can be seen in Figure 6.11.

Note that the plotted cost per per passenger mile is not exactly a fair comparison, as the cost for jet fuel is computed using the generic price, rather than the discounted price that airlines pay for buying in bulk. Additional taxes that fund runway infrastructure are built into the price of jet fuel that are not accounted for in the electricity or aluminum price [30]. Nonetheless, it can at least act as a decent first-order comparison, showing that the price difference is not obviously prohibitive. Nonetheless, an early market entry of these designs, even without the recommended aluminum recycling infrastructure, may allow for lower operating costs over current aircraft that may be dependent on government policy as much as economic questions. However, substantial weight penalties exist over comparable turboprops, which suggests that their unit price would be higher than similar gas-powered aircraft. The new technology and novel configuration would only compound this. Furthermore, the data here do not include labor, which would bring the operating cost curves closer to the other transport aircraft.

6.3.1 Optimizer Decomposition in Hybrid Battery Aircraft

Now, due to the additional complexity required in modeling aircraft with multiple propulsion systems, using an optimizer decomposition-based approach becomes significantly more complicated than in the single battery case (See Section 6.2.1). Each individual battery's properties must be

carefully bookkept to ensure consistency. That is, the constraints shown in (6.7)-(6.10) must be repeated for each battery.

$$E_{bat,prim} = m_{bat,prim} * Esp_{prim} \tag{6.17}$$

$$E_{bat,aux} = m_{bat,aux} * Esp_{aux} \tag{6.18}$$

$$P_{bat,prim} = m_{bat,prim} * Pbat_{prim} \tag{6.19}$$

$$P_{bat,aux} = m_{bat,aux} * Pbat_{aux} \tag{6.20}$$

$$E_{constraint,prim} = \frac{E_{bat,prim} - E_{mission,prim}}{C_{energy}} \geq 0 \tag{6.21}$$

$$E_{constraint,aux} = \frac{E_{bat,aux} - E_{mission,aux}}{C_{energy}} \geq 0 \tag{6.22}$$

$$P_{constraint,prim} = \frac{P_{bat,prim} - P_{mission,prim}}{C_{power}} \geq 0 \tag{6.23}$$

$$P_{constraint,aux} = \frac{P_{bat,aux} - P_{mission,aux}}{C_{power}} \geq 0 \tag{6.24}$$

In addition, because a metal-air battery is used, one must include some form of (6.11) as a constraint to account for mass accumulation throughout the mission. Additional bookkeeping required for the water that also needs to be carried onboard. Now, in practice, (6.23) would never be active for the cases show in this work, because of the battery discharge rules in (6.12) and (6.13) . Other, more complicated designs may have more complex rules in battery operation which may allow for violation of this constraint that need to be solved and accounted for at the design level.

6.4 Diesel-Electric Aircraft

The diesel-electric hybrid aircraft is, conceptually, very similar to the aluminum-air aircraft. It uses a diesel combustion engine, connected in series with a lithium-ion battery to keep it charged throughout the mission. This means that all of the power requirements have to be met by the battery, but the diesel engine may only need to provide a small fraction of the power, based on a parameter called "$ic\_power\_fraction$" here, defined as

$$ic\_power\_fraction = \frac{P_{comb}}{P_{total}}. \tag{6.25}$$

Note that, as a series hybrid design, using the electrification factor from (2.6) to parametrize the system would be nonsensical, as again, the electrical system by definition provides all vehicle power. The diesel-electric case also has the same number of sizing parameters to iterate on, but substitutes fuel mass for the primary battery energy. For the diesel-electric case, like the multiple battery aircraft, one needs to define how exactly how each energy system operates. For the aircraft here, the logic can be seen in (6.26) and (6.27)

$$\begin{aligned} &\text{If } P_{mission} \leq P_{comb,max} \\ &P_{comb} = P_{mission} \end{aligned} \tag{6.26}$$

$$\begin{aligned} &\text{Else If } P_{mission} > P_{comb,max} \\ &P_{comb} = P_{comb,max} \\ &P_{bat} = P_{mission} - P_{comb,max} \end{aligned} \tag{6.27}$$

Now, again, to allow for high-power maneuvers, some logic is included so that the battery becomes charged during the low-power mission segments. For completeness, some pseudocode defining how each energy system operates can be seen in (6.28)

$$\begin{aligned} &\text{If } P_{mission} < P_{comb,max} \\ &P_{bat} = -(P_{comb,max} - P_{mission}), \end{aligned} \tag{6.28}$$

where like before, the "−" indicates that the battery is being charged. Additionally, because of the nature of combustion engines, they tend to run more efficiently near their maximum design power. This allows the designer to play "games" with the relative operation of each system shown in (6.26), (6.27), and (6.28) to allow for a more energy efficient system. As an aside, this also means that accurately modeling off-design performance becomes much more important here than in the all-electric cases. Furthermore, these phenomena are less important in the commercial fixed-wing realm, where the power gap between cruise and takeoff is less severe than in VTOL or military aircraft. This can be illustrated by referring back to Figure 2.5. When a serial-hybrid design is used (such as in this dissertation), the battery and motor must be able to supply all the power needs of the aircraft. Practically, this means that one must simulate running $P_{mission}$ through the battery and motor systems to account for these efficiency losses (unless constant efficiencies are assumed). Sizing parameters for this design can be seen in (6.29) below.

$$y = [GTOW, m_{fuel}, E_{bat}, P_{in}] \tag{6.29}$$

Note that this aircraft possesses a simpler mission profile, with two climb segments, a cruise velocity of 300 knots along with 3 descent segments to return to sealevel. Cruise range is 65 nautical miles. The payload weight consists of three passengers and a pilot. Like the aluminum-air case, significant oscillation can be seen as one iterates towards the solution, albeit in converged cases, this oscillation tends to be less severe than in the aluminum-air case as shown in Figure 6.12

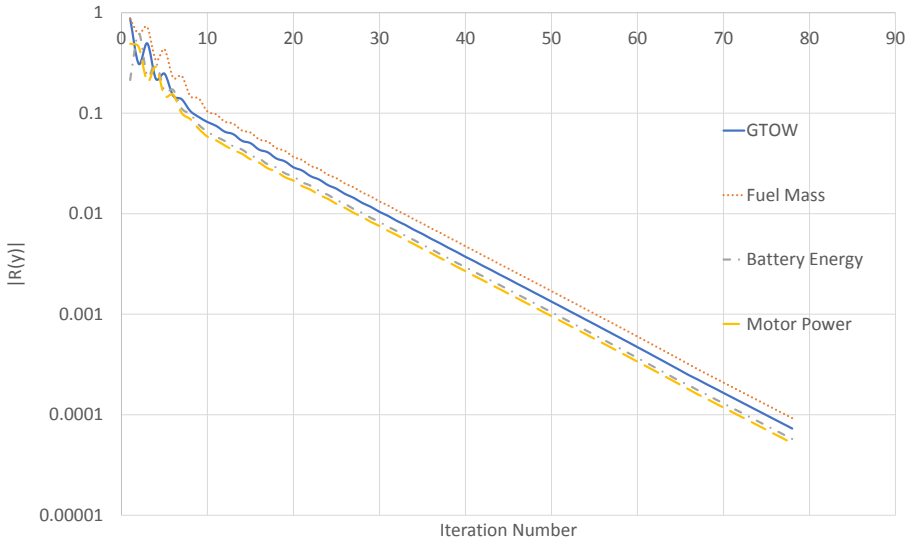

Figure 6.12: Convergence for Diesel-Electric Aircraft

On the other hand, because of the scale and mission, as well as due to the fact that the energy systems are in series, certain pathological cases arise which make the problem unsolvable using a simple successive substitution iteration method. An example of this can be seen in Figure 6.13.

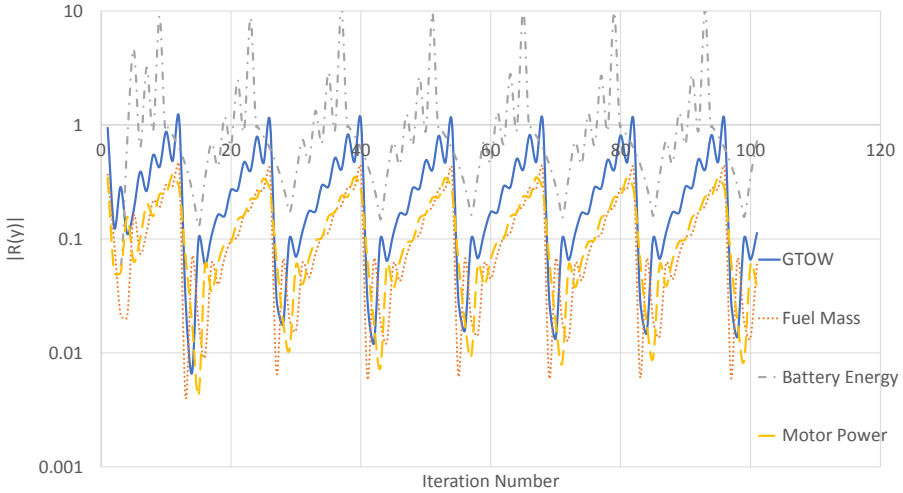

Figure 6.13: Convergence for Diesel-Electric Aircraft

Cases such as that shown in Figure 6.13 occur when, in the course of the solution process, the battery remains fully charged at a given iteration, returning a value of 0 for *Ebat* (see (6.29)). However, based on the required updated values for $GTOW$, m_{fuel},E_{bat}, P_{in} (along with the chosen value for $ic\_power\_fraction$), this may result in a nonzero value for the next (i.e. second) iteration on E_{bat}, causing the successive substitution algorithm to oscillate between these cases, preventing convergence to any reasonable tolerance level. Practically, this also means there is a strong likelihood that the aircraft sizing loop can close using only the diesel engine.

Now, as in the aluminum-air case, something should be said about the effect of allowing the user to account for other design points when designing/sizing the aircraft. Figure 6.14 shows the case when the battery weight is minimized based on the energy and power requirements along the lithium Ragone plot from Figure 2.4

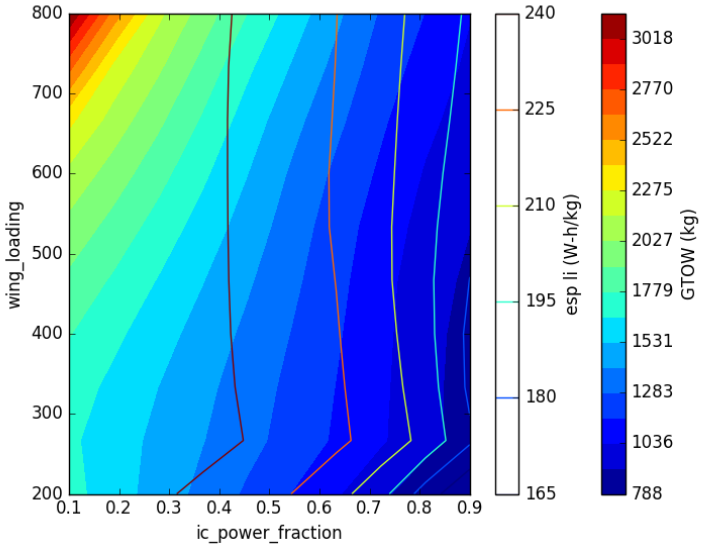

Figure 6.14: Diesel-Electric Gross Landing Weight Contour Plot (Newton-Raphson)

Note that like in the aluminum-air cases, the optimal chosen specific energy determined by the bounded Ragone optimization subproblem is shown in Figure 6.14. The constant specific energy (set to 200 $\frac{W-h}{kg}$) case is shown in Figure 6.15.

As should be readily apparent from the color bar scales, this serial-hybrid case also exhibits an extreme sensitivity to the specific energy/power assumption used. The heaviest cases, in particular, are doubled when specific energy is set to a constant. Taken in conjunction with the results from the aluminum-air case, this illustrates that accurate assumptions in battery performance from both

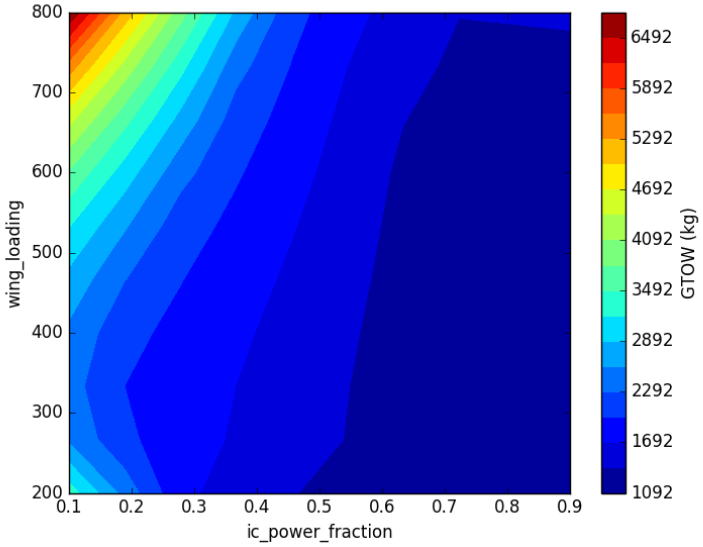

Figure 6.15: Diesel-Electric Gross Landing Weight Contour Plot (Constant Specific Energy)

power and energy-related perspectives are critical in the earliest stages of the design process, particularly when tradeoffs between multiple energy systems are leveraged as in the hybrid-electric and hybrid battery cases.

6.4.1 Optimizer Decomposition in Diesel-Electric Aircraft

As one might surmise, the optimizer decomposition-based approach for the diesel-electric aircraft is also quite similar to the multiple-battery aircraft. In this case, many of the constraints are identical to those shown in (6.32)-(6.35), which are repeated below for the sake of completeness.

$$E_{bat} = m_{bat} * Esp \tag{6.30}$$

$$P_{bat} = m_{bat} * Pbat \tag{6.31}$$

$$W_{constraint} = \frac{GTOW - W_{out}}{C_{weight}} \geq 0 \tag{6.32}$$

$$m_{fuel,constraint} = \frac{m_{fuel} - m_{fuel,out}}{C_{fuel}} \geq 0 \tag{6.33}$$

$$E_{constraint} = \frac{E_{bat} - E_{mission}}{C_{energy}} \geq 0 \tag{6.34}$$

$$P_{constraint} = \frac{P_{bat} - P_{mission}}{C_{power}} \geq 0 \tag{6.35}$$

Chapter 7

Regression Augmented Sizing

7.1 Regression Algorithm Comparison

Because the aluminum-air aircraft was the most complicated aircraft class, training data from a full optimization case was generated and used to evaluate several regression algorithms to determine a "good" initial guess in an attempt to generalize their use here. This aircraft has 4 sizing variables as seen in (6.15), along with twelve input variables shown in Table 7.1.

Table 7.1: Aluminum-Air Design Variables

variable	Lower Bound	Upper bound
f_{aux}	0	1.
Wing Loading $(\frac{kg}{m^2})$	200	800
Thrust Loading $(\frac{N}{N})$	.05	.6
Fan Pressure Ratio	1.1	2.4
$\frac{t}{c}$	.07	.2
$V_{climb_1}(\frac{m}{s})$	50	140
$V_{climb_2}(\frac{m}{s})$	50	140
$V_{climb_3}(\frac{m}{s})$	50	140
$h_{cruise}(ft)$	20,000	30,000
$h_{frac,climb_1}$	.1	1
$h_{frac,climb_2}$	.2	1
$h_{frac,descent_1}$	.1	1

f_{aux} refers to the parameter defined in (6.16), while thrust loading is defined here by the engine system thrust (in Newtons) normalized by the aircraft landing weight (in Newtons). The fan pressure ratio refer to the design total pressure ratio across the ducted fan and is a measure of the amount of work performed by the propulsor; lower values correspond to a larger ducted fan with a higher on-design propulsive efficiency. The thickness-to-chord ratio of the main wing is $\frac{t}{c}$. The terms V_{climb_1},

V_{climb_2} V_{climb_3}, refer to the total magnitude of the velocity at each climb segment while h_{cruise} is the altitude at cruise. $h_{frac,climb_1}$ and $h_{frac,climb_2}$ refer to the fraction of cruise altitude that the aircraft flies before beginning the next segment (there are three climb segments). $h_{frac,descent_1}$ refers to the altitude that the aircraft descends to as a fraction of the cruise altitude (there are two descent segments). Now, as an aside, it should be noted that changing these trajectory variables proved crucial in determining the size of the actual aircraft seen in Chapter 9 as one or both batteries had a tendency to frequently be power limited. As the optimization process proceeded, the algorithm was trained on iteration 1 through j-1, and tested on iteration j, where the norm of the error was output. In other words

$$\begin{aligned} &\text{for j in converged results} \\ &\quad \text{train on 1:j-1} \\ &\quad \text{test on j.} \end{aligned} \tag{7.1}$$

Some statistics were then calculated to compare the effectiveness of each algorithm over the entire dataset; i.e. data is saved from a representative optimization case, and (7.1) is run for each data point, with the value for each sizing parameter estimated based on the previous data points. The error is then tabulated at each step, and the mean, median, maximum, and minimum error are recorded along with the standard deviation. Table 7.2 compares some representative linear-regression algorithms (many others were tested, but these proved to be the most effective on this set of data). Note that, the mean, median, and maximum errors were the most significant metrics for algorithm performance. The mean and median both represent the performance the algorithm is most likely to have when evaluating a particular optimization point, while the maximum error, if too high (meaning a poor initial guess), could cause the sizing loop to diverge, which could result in a breakdown of the optimization process (when run in-the-loop). Even when run outside of the loop, spurious divergence based on a poor initial guess can grossly misinform the designer as to the feasibility of a given candidate aircraft.

Table 7.2: Aluminum-Air Aircraft Outer Loop Performance (Linear Regression Algorithms)

algorithm	mean err	median err	std err	min err	max err
SVR (RBF Kernel)	0.830	0.432	1.576	0.070	6.935
Gaussian Process	0.462	0.300	0.433	0.025	2.297
Table	0.784	0.062	2.183	0.001	11.166
Kneighbors (5), distance weighted	0.454	0.131	1.037	0.006	6.505

In addition, several ensemble methods were tested on this dataset. Note that, the ensemble methods are stochastic, so the results will vary slightly each time the method is run. However, after running each several times, it was determined that the results in Table 7.3 were representative of the overall performance of each method.

Table 7.3: Aluminum-Air Aircraft Outer Loop Performance (Ensemble Algorithms)

algorithm	mean err	median err	std err	min err	max err
Random Forest	0.403	0.293	0.402	0.003	2.110
Gradient Boosting	0.324	0.147	0.395	0.002	1.898
Bagging Regressor	0.400	0.361	0.366	0.008	1.721
Extra Trees Regressor	0.316	0.137	0.380	0.009	1.901

For these cases, the simple table lookup model possessed the lowest median error, but a higher maximum error (and with it, a higher standard deviation), than the other methods. The max error occurs when the optimizer takes large steps, which tends to happen early on in the optimization process. Compared to the table-lookup method, K Nearest Neighbors (KNN) has lower maximum, mean, and median errors than the table lookup method. Gaussian Process Regression (GPR) reduced the maximum error by a factor of three, but with lower average accuracy. This is indicative of a more robust but less accurate algorithm. The ensemble methods, on the other hand, all appear to maintain the robustness of GPR, with the addition of increased accuracy, with Extra Random Trees edging the other algorithms out in terms of performance. With the algorithm performance estimated for the most complicated case, the table-lookup, K Nearest Neighbors (using distance-weighted coefficients, and up to 5 neighbors), Gaussian Process Regression, and Extra Random Trees were evaluated and compared for three aircraft classes (the regional jet, lithium-air jet, and the aluminum-air aircraft). This can be seen in Chapter 9.

Now, it should be noted that, realistically, using these predictive algorithms to extrapolate too far from calculated data, or in an unbounded fashion, can lead to spurious divergence. An extreme example of this can be seen in Figure 7.1, which uses an unbounded Gaussian Process Regression algorithm to predict new values for each y value (for a 2-D sweep of wing loading vs. f_{aux}).

Due to the low specific powers and energies experienced here, poor initial guesses can lead to divergence in an otherwise feasible aircraft candidate. The "fix" in this cases is to use a nonextrapolating algorithm (such as K Nearest Neighbors) if certain criteria are not met, such as "closeness" to existing designs, or more relevantly, reasonable values for an initial guess. An important aspect of this is that each individual guess for weight, energy, and power may, by themselves appear to be reasonable; however, their combination may result in values for y that go to infinity. Thus, changing the initial guesses for sizing parameters only on an individual basis results in plots with a similar number of "holes" as Figure 7.1 shows. Figure 7.2 shows the same case when all sizing variables are bounded such that if one of $[W_{in}, E_{prim}, E_{aux}, P_{in}]$ appears to be unreasonable, a nonextrapolating algorithm (K Nearest Neighbors) is used to determine an informed guess for all of $[W_{in}, E_{prim}, E_{aux}, P_{in}]$. The results can be seen in Figure 7.2.

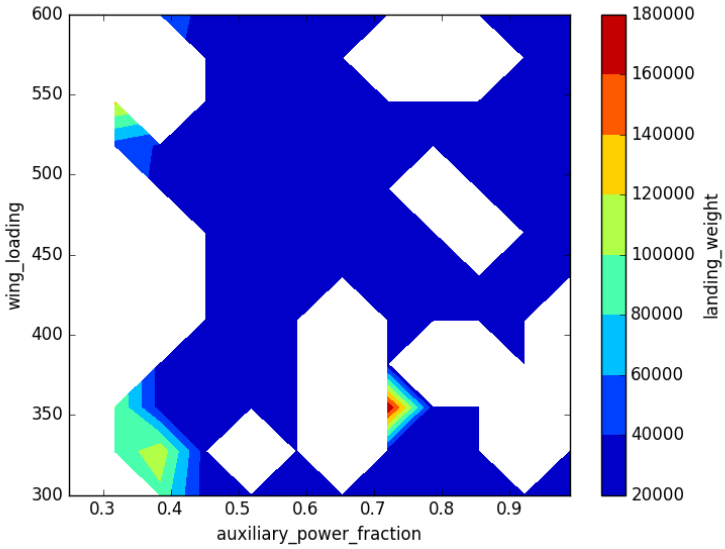

Figure 7.1: Aluminum-Air Gross Landing Weight Contour Plot (Successive Substitution, Gaussian Process Regression with Holes)

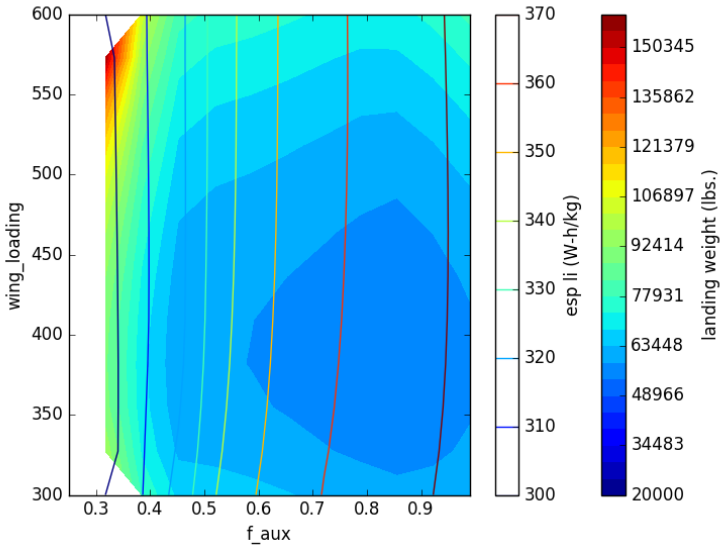

Figure 7.2: Aluminum-Air Gross Landing Weight Contour Plot (Successive Substitution, Gaussian Process Regression with Bounded Initial Guesses)

As Figure 7.2 shows, properly bounding these algorithms such that they do not venture into nonsensical regions, or too far from cataloged data is crucial to ensure robustness in the results. This can be seen as analogous to the "trust region" inside the Trust Region Model Management Algorithm mentioned in Section 1.4.2 [3].

Chapter 8

Sweep Comparison

Now, because most global optimization schemes utilize stochastic methods with comparably large distances between the candidate designs, an attempt was made to mimic this behavior in a deterministic manner in order to compare the algorithms employed/developed thus far. To that end several sweeps of the design space were performed for the aluminum-air and diesel-electric cases to illustrate the relative performance of the algorithms shown in Chapters 4 and 7. Qualitatively, this reveals some very interesting phenomena that can arise when multiple propulsion systems are utilized. Section 8.1 exhibits the potential for a multiplicity of sizing solutions, explored in some detail in Chapter 5. Section 8.2, on the other hand displays the shortcomings that sometimes arise when attempting to solve the system without using a robust backtracking algorithm.

8.1 Hybrid Battery Aircraft Sweeps

Sweeps of landing weight vs. auxiliary power fraction (or f_{aux}) and wing loading (in $\frac{kg}{m^2}$) can be seen in Figures 8.1 (using successive substitution) and 8.2 (using Newton-Raphson); in this case, a 12X12 grid of values for wing loading and auxiliary power fraction were computed, with the resulting aircraft weights shown. The choice of units of $\frac{kg}{m^2}$ for wing loading is based off of the convention from Raymer [80]. Note that these aircraft determined the optimal point on the Ragone plots for both batteries (see Figure 2.4) and a line plot of the lithium battery specific energy is overlayed.

As should be readily apparent, Figures 8.1 and 8.2 are not identical (in some regions of the design space), despite the fact that the systems used to model them are the exact same. Reasons for this can be seen in some detail in Chapter 5, which explores the exact, somewhat complicated mechanism behind this phenomena; what can be said with certainty is that uniqueness of these systems is unprovable, so aircraft designers should be aware of the potential for such issues when performing large-scale searches. One culprit is the low specific energy and power of the energy systems. This means that the iterative algorithms used to solve the system move through comparably large domains

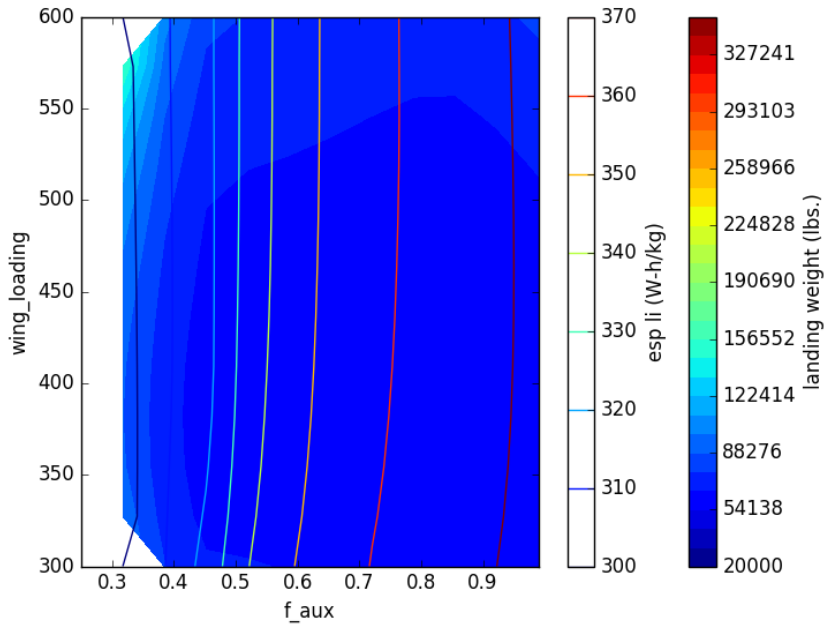

Figure 8.1: Gross Landing Weight Contour Plot (Successive Substitution)

which increases the possibility that additional solutions are found. To drive home the significance, the aircraft from Figures 8.3a and 8.3b (created using OpenVSP) possess identical fuselages, wing loading as well as f_{aux}, but differ wildly in terms of weight and performance [33].

A human may readily be able to decide that the aircraft from Figure 8.3a is much more credible than the one in Figure 8.3b a posteriori; however, in the context of a large sweep of the design space, using the "wrong" methodology to converge the system could result in the information from the design in Figure 8.3a being lost, which could likely lead to a heavier, lower-performance aircraft. Now, it should be noted that using Newton-Raphson rather than successive substitution to converge the system significantly increases the likelihood that these errant solutions are found. One fix to this problem is to increase the threshold at which the Newton iterations begin; the plot constructed in Figure 8.2 used a threshold of 5% convergence to switch to Newton-Raphson. Using a 2% threshold seems to solve this issue here, as Figure 8.4 shows.

As an aside, this 2% number was arrived at heuristically. An additional benefit is that the Newton-Raphson algorithm here is able to solve one case that did not converge using successive substitution (the lower left hand corner where $f_{aux} \approx .3$, and $WL \approx 325\ \frac{kg}{m^2}$). Now in addition to this increased robustness, the 2% case shows a nontrivial increase in code speed. Table 8.1 compares the relative computational cost of a variety of methods used to both determine the initial step as well

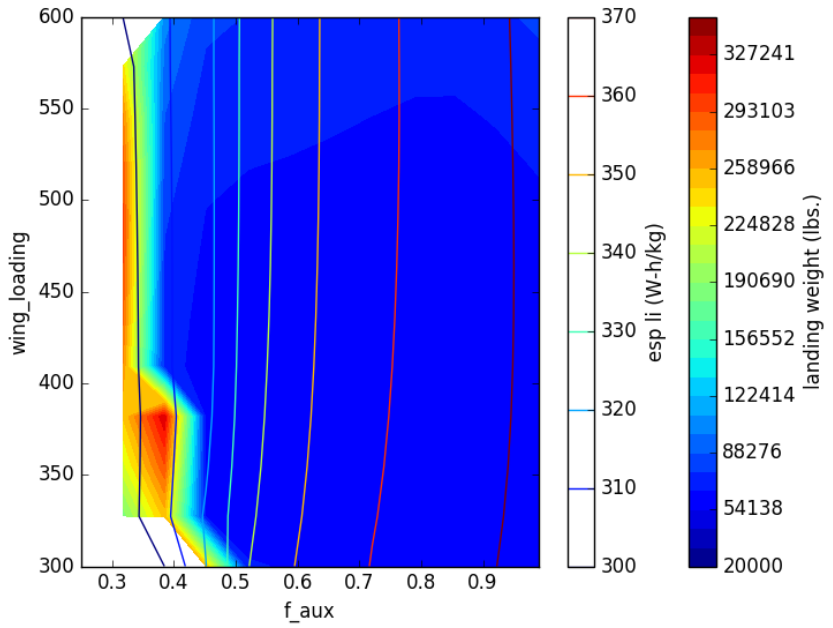

Figure 8.2: Gross Landing Weight Contour Plot (Newton-Raphson, Aluminum-Air)

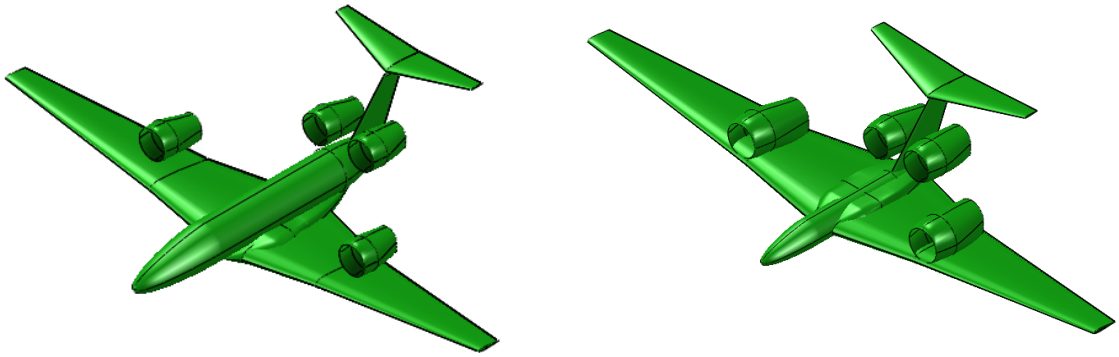

(a) Successive Substitution Solution

(b) Newton-Raphson Solution

Figure 8.3: Aluminum-Air Aircraft Solution Multiplicity Comparison

as converge the system. Note that a major reason for the improved code speed for the 2% Newton-Raphson case over the 5% case is that the 2% case does not converge to the larger, extraneous solutions which, in addition to being less desirable from a design standpoint, also takes longer to converge.

The use of machine learning algorithms from Chapter 7 shows a code speed improvement of about

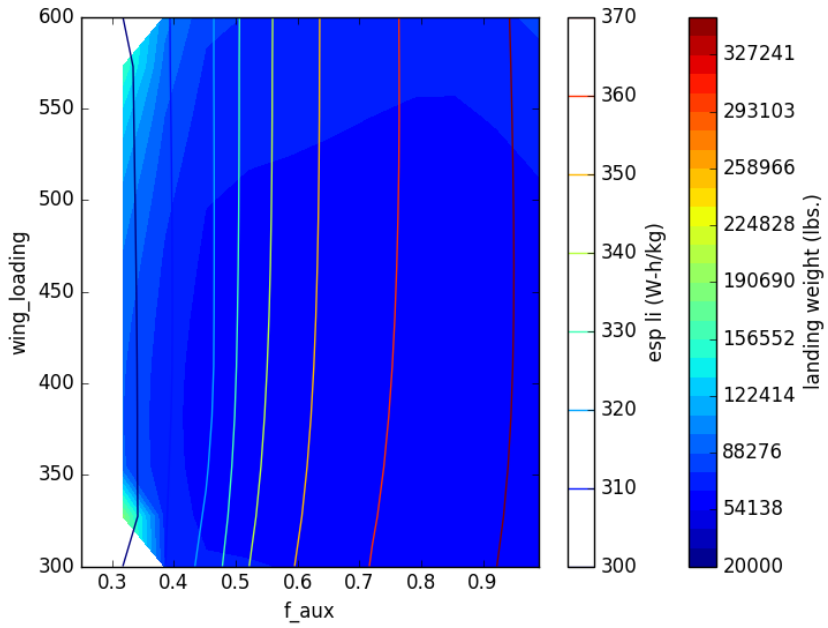

Figure 8.4: Gross Landing Weight Contour Plot (Newton-Raphson (2% threshold), Aluminum-Air)

Table 8.1: Aluminum-Air Sweep Results

loop evaluation	initial step	$n_{converged}$ (possible 144)	n_{size}	n_{size} (for converged aircraft)
Successive Substitution	Same Point	129	4,829	3,288
Successive Substitution	KNN(5)	129	3,925	2,088
Successive Substitution	GPR	129	3,960	2,388
Newton-Raphson 5% threshold	Same Point	130	4,020	2,645
Newton-Raphson 2% threshold	Same Point	130	3,957	2,566
Newton-Raphson 2% threshold	KNN(5)	130	3,460	2,088
Newton-Raphson 2% threshold	GPR	130	3,418	2,033

30% over the baseline case. Similarly, using Newton-Raphson, rather than successive substitution to converge the systems improves code speed by another 20%. Interestingly, there is one design that the successive substitution algorithms proved unable to converge within the limit of 100 iterations; Newton-Raphson on the other hand had no trouble with said case. All entries within Table 8.1

utilized a backtracking algorithm to ensure that each new iteration was more converged than the previous iteration.

8.2 Diesel-Electric Sweeps

Like in Section 8.1 for the aluminum-air case, sweeps of $GTOW$ vs. wing loading and $ic\_power\_fraction$ were calculated and plotted for the diesel-electric aircraft. Here, a 10x10 grid was solved (or at least, attempted) using a simple successive substitution algorithm and is plotted in Figure 8.5.

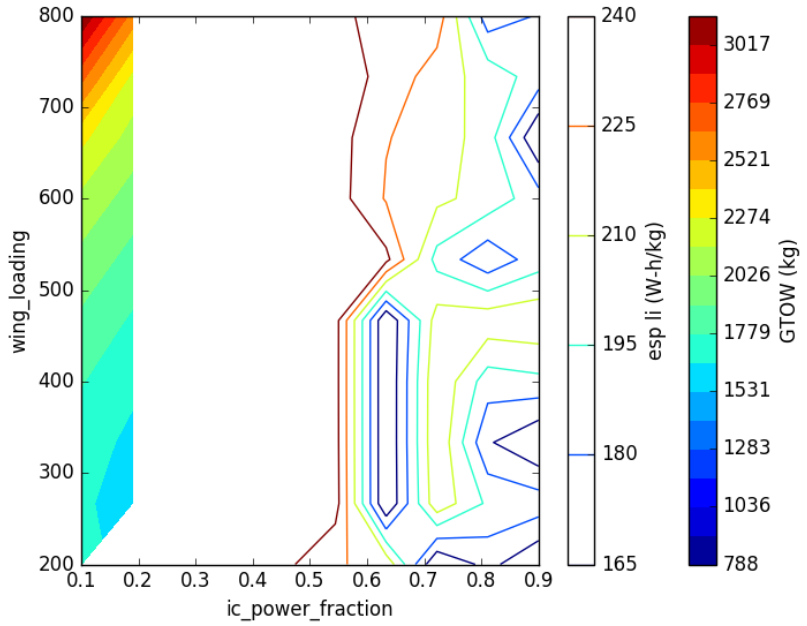

Figure 8.5: Diesel-Electric Gross Landing Weight Contour Plot (Successive Substitution)

As can readily be seen, most of the designs attempted in Figure 8.5 did not converge. This is due to the issues highlighted back in Figure 6.13, wherein depending on the guess for each propulsion system, the battery may not discharge at all, resulting in oscillation until the algorithm terminates. For higher values of the $ic\_power\_fraction$, which correspond to a greater proportion of diesel engine power, convergence does not happen. This is because, for the indicated guesses, the diesel engine can provide all of the requested power for the mission. Notably, the cases in Section 8.1 diverged due to low values of specific energy and power, whereas the nonconverged results here were purely an algorithmic issue. Figure 8.6 shows how this problem is solvable with the simple application of backtracking, wherein rather than immediately accepting the returned values for the successive

substitution algorithm, the solver checks if this guess is more converged than the previous guess. If not, a line search is performed along the direction of the (initial) returned values until the residual is decreased.

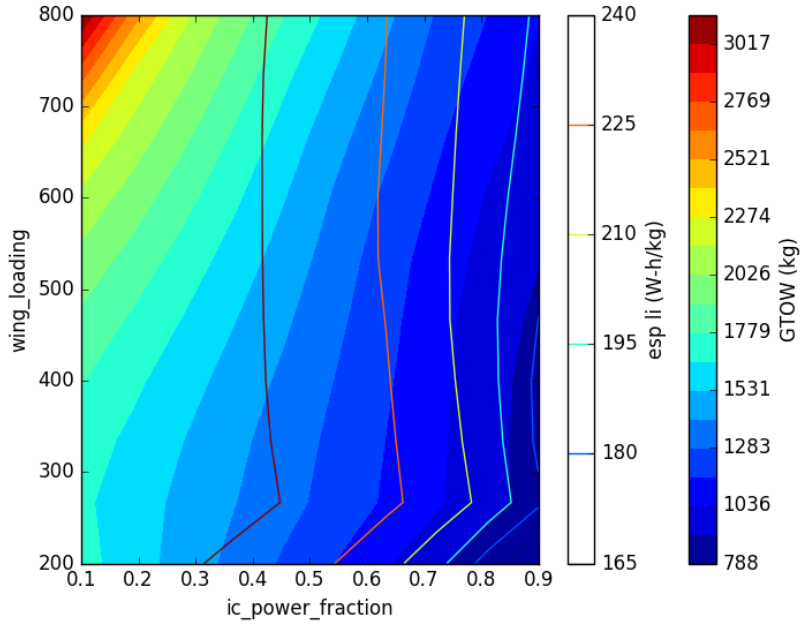

Figure 8.6: Diesel-Electric Gross Landing Weight Contour Plot (Successive Substitution, Backtracking)

Figure 8.6 shows a number of interesting characteristics. First, note how at a wing loading of about 250 $\frac{kg}{m^2}$, there is a "kink" in the calculated specific energy curve. This can be explained by a change in the mission profile; remembering that the battery has to supply all of the power needs of the mission, decreasing the wing loading can result in cases where the internal combustion engine is less able to handle the power draw of the mission, meaning that the chosen specific energy on the Ragone curve is increased. As an aside, this aircraft does not exhibit any obvious solution multiplicity issues that the aluminum-air aircraft does. This can readily be seen in Figure 8.7 which plots the same sweep of $GTOW$ vs. wing loading and $ic\_power\_fraction$ where Newton-Raphson is used to converge the system.

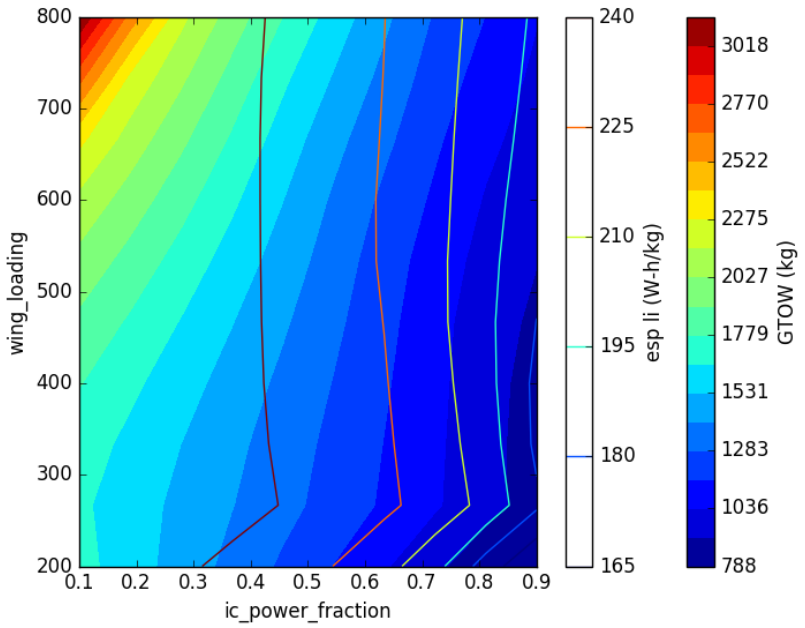

Figure 8.7: Diesel-Electric Gross Landing Weight Contour Plot (Newton-Raphson, repeat)

Figures 8.6 and 8.7 appear identical, illustrating that this aircraft system seems less prone to the solution multiplicity issue than the aluminum-air case. This may be partially a result of the fact that different component weight buildup correlations were used; the aluminum-air case used commercial aircraft empirical fits from Shevell, while the diesel-electric aircraft used weights correlations for a general-aviation aircraft that were created by Raymer [87, 80]. As an aside, naive implementation of Newton-Raphson to converge the system results in qualitatively similar "holes" in the problem where the system diverges, albeit for entirely different reasons. In this case, it is because, due to the small size of the aircraft, the returned value from Newton-Raphson sometimes returned negative values for the sizing parameters, causing the system to break down. This is solved in a similar way, with the use of backtracking; here, if any sizing parameter oversteps a prescribed bound (such as, for instance, making $GTOW <$ the payload weight), the solver backtracks along the search direction until all of the sizing parameters move to within sensible bounds. Note that, as in Chapter 7, all of the sizing parameters need to be changed simultaneously, as this is a multidimensional root finding problem. Table 8.2 lists the relative computational cost and robustness of running this 10x10 sweep using a variety of different methods.

Table 8.2: Diesel-Electric Sweep Results

loop evaluation	initial step	$n_{converged}$ (possible 100)	n_{size}	n_{size} (for converged aircraft)
Successive Substitution	Same Point	100	2,734	2,734
Successive Substitution	Same Point (no backtracking)	22	9,126	1,248
Successive Substitution	GPR	100	2,022	2,022
Newton-Raphson 5% threshold	Same Point	100	2,298	2,298
Newton-Raphson 2% threshold	Same Point	100	2,252	2,252
Newton-Raphson 5% threshold	GPR	100	1,910	1,910
Newton-Raphson 2% threshold	GPR	100	1,973	1,973

As should be readily apparent, backtracking for all algorithms, including the successive substitution cases, is critical in ensuring robust convergence of these designs. Furthermore, using Newton-Raphson to accelerate convergence of these systems shows a clear improvement in computation cost, especially when used in conjunction with a machine learning algorithm that "learns" well-informed initial guesses. Note that, compared to the aluminum-air aircraft, this case does not show a tendency to diverge because of specific energy/power limitations, so in that regard, it is an easier problem to solve robustly, at least once a backtracking algorithm is implemented.

Chapter 9

Optimization Results

This chapter presents some representative optimization results using the methodology outlined in Chapter 7 and compares them to a standard aircraft design framework. Results using an optimizer decomposition strategy should be included in a followup work. Now, the choices in optimization variables and bounds for the first two cases can be seen in Table 9.1.

Table 9.1: Design Variables (Regional Jet and Lithium-Air Jet)

variable	Lower Bound	Upper bound
Wing Loading $(\frac{kg}{m^2})$	200	800
Thrust Loading	.05	.6
Fan Pressure Ratio	1.1	2.4
$\frac{t}{c}$	.07	.2
$V_{climb_1}(\frac{m}{s})$	50	140
$V_{climb_2}(\frac{m}{s})$	50	140
$V_{climb_3}(\frac{m}{s})$	50	140
$h_{cruise}(ft.)$	20,000	30,000
$h_{frac,climb_1}$	.1	1
$h_{frac,climb_2}$	.2	1
$h_{frac,descent_1}$	.1	1

An explanation of their terminology can be seen in Chapter 7. Additionally, the regional jet, lithium-air jet, and the aluminum-air jet each possessed the same mission flight profile and constraints. Design variables were constrained to limit required field lengths, as well as ensure a reasonable angle of attack and consistency in the flight profile. The constraints can be seen in (9.1) - (9.5).

$$h_{frac,climb_1} < h_{frac,climb_2} \tag{9.1}$$

$$-15^\circ \leq AoA \leq 15^\circ \tag{9.2}$$

$$TOFL \leq 2500ft \tag{9.3}$$

$$LFL \leq 2500ft \tag{9.4}$$

$$throttle \leq 1 \tag{9.5}$$

All optimization cases in this chapter use SNOPT, a gradient-based optimizer. Global optimization results for the analytical case can be seen in Chapter 5. In the context of this work, only gradient-based optimization cases will be examined for the full aircraft, as time constraints preclude collecting statistically significant results for the full range of designs and numerical methods. Based on the results from Section 5.4, however, the use of machine learning and Newton-Raphson's method or Broyden's method to converge these types of systems appears promising in increasing code speed. Nonetheless, appropriate countermeasures need to be in place to account for some of the numerical difficulties seen in Chapters 5, 6, and 8.

9.1 Regional Jet

Sizing results using the methodology in Chapter 7 for a regional jet can be seen in Table 9.2 where successive substitution is used to converge the aircraft fuel consistency residual and the methodology used to determine a "good" initial guess is varied.

Table 9.2: Regional Jet Optimization Results (Successive Substitution)

loop evaluation	initial step	n_{size}	n_{opt}	$\frac{n_{size}}{n_{opt}}$	landing weight (lbs.)
Successive Substitution	Same Point	406	57	8.8	83,282
Successive Substitution	Table	149	46	3.2	83,454
Successive Substitution	KNN(5)	150	46	3.3	83,454
Successive Substitution	GPR	150	46	3.3	83,454
Successive Substitution	Extra Trees	149	46	3.2	83,457

The regional jet case was interesting in that the problem was simple enough that the optimization case converged after only $\approx$ five major iterations, meaning that the regression algorithms only ran for one or two iterations (three converged data points were chosen as the threshold before regression methods other than a table lookup were run, to prevent extrapolation to extremely unphysical results). Because of the swift convergence of this optimization problem, the extra complication in setting up the more elaborate algorithms may not necessarily be cost-effective in a more applied setting. For completeness' sake, results using Newton-Raphson to converge the inner loop can be

seen in Table 9.3, while results using Broyden's method to update the Jacobian can be seen in Table 9.4.

Table 9.3: Regional Jet Optimization Results(Newton-Raphson)

loop evaluation	initial step	n_{size}	n_{opt}	$\frac{n_{size}}{n_{opt}}$	landing weight (lbs.)
Newton-Raphson	Same Point	312	46	6.8	83,259
Newton-Raphson	Table	177	46	3.8	83,252
Newton-Raphson	KNN(5)	238	46	5.2	83,274
Newton-Raphson	GPR	220	46	4.8	83,267
Newton-Raphson	Extra Trees	229	46	5.0	83,271

Table 9.4: Regional Jet Optimization Results(Broyden)

loop evaluation	initial step	n_{size}	n_{opt}	$\frac{n_{size}}{n_{opt}}$	landing weight (lbs.)
Broyden's Method	Same Point	312	145	6.8	83,261
Broyden's Method	Table	211	57	3.7	83,255
Broyden's Method	KNN	205	57	3.6	83,249
Broyden's Method	GPR	196	57	3.4	83,251
Broyden's Method	Extra Trees	269	68	4.0	83,321

The more exotic loop-evaluation methods (Newton-Raphson and Broyden) proved to be somewhat more computationally expensive compared to the successive substitution updates; this is because they had a tendency to "overshoot" the solution when SNOPT performed the finite difference iterations, resulting in less accurate gradients. Furthermore, because most of the sizing loop iterations were taken during finite difference steps in the optimizer, the initial guess tended to be "near" the solution value (for the regression methods). As a result, the extra sizing loop iterations to construct the Jacobian may actually be slower than successive substitution in any given case.

9.2 Lithium-Air

Tables 9.5 shows results for the lithium-air-powered aircraft, which, conceptually, is somewhat more complicated than the regional jet due to the fact that power requirements are handled "in the loop."

Note that, like the regional jet case, the lithium-air jet optimization problem was found to converge within only a few optimizer iterations, so the effectiveness of each algorithm is difficult to compare. Nonetheless, for successive substitution, the overall number of sizing calls was lower for each of these methods by an order of magnitude than by starting from the same initial y values. Results for the lithium-air jet using Newton-Raphson and Broyden's method to converge the inner sizing loop can be seen below.

Table 9.5: Lithium-Air Jet Optimization Results (Successive Substitution)

loop evaluation	initial step	n_{size}	n_{opt}	$\frac{n_{size}}{n_{opt}}$	landing weight (lbs.)
Successive Substitution	Same Point	3,112	326	9.55	131,272
Successive Substitution	Table	192	62	3.1	134,208
Successive Substitution	KNN(5)	171	50	3.4	134,277
Successive Substitution	GPR	206	62	3.3	134,443
Successive Substitution	Extra Trees	197	62	3.1	134,231

Table 9.6: Lithium-Air Jet Optimization Results(Newton-Raphson)

loop evaluation	initial step	n_{size}	n_{opt}	$\frac{n_{size}}{n_{opt}}$	landing weight (lbs.)
Newton-Raphson	Same Point	3,894	350	11.1	130,918
Newton-Raphson	Table	765	134	5.7	131,210
Newton-Raphson	KNN(5)	2,632	640	4.1	133,056
Newton-Raphson	GPR	1,052	218	4.8	130,944
Newton-Raphson	Extra Trees	3,415	74	5.7	132,951

Table 9.7: Lithium-Air Jet Optimization Results Summary(Broyden)

loop evaluation	initial step	n_{size}	n_{opt}	$\frac{n_{size}}{n_{opt}}$	landing weight (lbs.)
Broyden's Method	Same Point	2,845	314	6.8	131,371
Broyden's Method	Table	867	182	3.7	131,101
Broyden's Method	KNN	422	112	3.6	133,419
Broyden's Method	GPR	214	62	3.4	133,715
Broyden's Method	Extra Trees	714	170	4.0	130,969

The lithium-air case demonstrates the "overshooting" problem with Newton-Raphson, as it took a substantially larger number of optimizer calls to converge than the other cases. Additionally, from inspection, none of the initial step methods (Table, KNN, GPR, or Extra Trees) appears to perform significantly better than the others. Note that, like the regional jet case, the lithium-air jet optimization problem was found to converge within only a few optimizer iterations (at least for successive substitution), so the effectiveness of each algorithm is difficult to compare. Nonetheless, for successive substitution, the overall number of sizing calls was lower for each of these methods by an order of magnitude than by starting from the same initial y values. Broyden's method mitigates some of the numerical noise issues seen in Newton-Raphson, but still appears to be marginally outperformed by the successive substitution methods (on average).

9.3 Hybrid Battery

The choice in optimization variables along with their bounds can be seen in Table 7.1, which appends f_{aux} as a design variable to the variables listed in Table 9.1. Optimization results for the aluminum-air aircraft can be seen in Table 9.8.

Table 9.8: Aluminum-Air Optimization Results(Successive Substitution)

loop evaluation	initial step	n_{size}	n_{opt}	$\frac{n_{size}}{n_{opt}}$	landing weight (lbs.)
Successive Substitution	Same Point	95,849	4,396	21.8	44,976
Successive Substitution	Table	7,227	1,343	5.4	45,040
Successive Substitution	KNN(5)	11,838	2,409	4.9	45,134
Successive Substitution	GPR	3,415	524	6.5	48,454
Successive Substitution	Extra Trees	10,555	2,045	5.2	44,979

The aluminum-air aircraft case, like the lithium-air aircraft, has an order of magnitude reduced computational cost when intelligent initial guesses were used (see Chapter 7. Furthermore, numerical noise became a more significant issue here, which substantially increased optimization time compared to the lithium-air case; note that for each of these algorithms, the optimizer design variables moved very close to the solution output (generally at $\approx$ 1,000 optimization calls (which included finite difference steps)), but was unable to satisfy the optimality conditions, so the optimizer spent a significant amount of time searching for better points. Results using the more complicated inner-loop methods can be seen below.

Table 9.9: Aluminum-Air Optimization Results(Newton-Raphson)

loop evaluation	initial step	n_{size}	n_{opt}	$\frac{n_{size}}{n_{opt}}$	landing weight (lbs.)
Newton-Raphson	Same Point	11,678	433	27.0	48,221
Newton-Raphson	Table	17,398	2,357	7.4	45,412
Newton-Raphson	KNN	11,208	1,695	6.6	48,330
Newton-Raphson	GPR	3,652	550	6.6	47,309
Newton-Raphson	Extra Trees	28,353	4,580	6.2	45,016

Table 9.10: Aluminum-Air Optimization Results(Broyden)

loop evaluation	initial step	n_{size}	n_{opt}	$\frac{n_{size}}{n_{opt}}$	landing weight (lbs.)
Broyden's Method	Same Point	10,562	328	32.2	45,754
Broyden's Method	Table	14,169	3,194	4.4	44,669
Broyden's Method	KNN	1,582	225	7.0	48,504
Broyden's Method	GPR	3,292	342	9.6	50,460
Broyden's Method	Extra Trees	1,620	173	9.4	50,435

The aluminum-air aircraft case, like the lithium-air aircraft, had an order of magnitude reduced computational cost when not starting from the same point. Additionally, solving the sub-problem proved more troublesome for the inner-loop solver, due to the increased complexity (as seen in Figure 4.1). Furthermore, numerical noise became a more significant issue here, which substantially increased optimization time compared to the lithium-air case. For some of the cases, Broyden's method struggled to fully converge, (due to accumulated numerical error), which significantly increased the number of sizing-loop iterations. This may be mitigated by using finite differencing to reinitialize the Jacobian. For the sake of completeness, local optimization cases for the diesel-electric case were also run.

9.4 Diesel-Electric Aircraft

This aircraft optimization problem is, in many ways, simpler than the other designs investigated in this work. Here, no mission variables are changed in the course of the optimization problem. Table 9.11 list the optimization variables and their bounds.

Table 9.11: Aircraft Design Variables (Hybrid)

Variable	Lower Bound	Upper Bound	Initial
Wing Loading (WL, $\frac{kg}{m^2}$)	200	800	400
Thrust Loading (TL, $\frac{kg}{kg}$)	0.1	0.80	0.32
Aspect Ratio (AR)	5	10	6.5
wing taper ratio	0.10	0.30	.28
wing thickness to chord ($\frac{t}{c}$)	0.07	0.20	.11
ic power fraction	0.1	.9	.4

As no design variables in Table 9.11 change the mission profile, fewer constraints are employed. Said constraints can be seen in (9.6) and (9.7).

$$throttle \leq 1 \tag{9.6}$$

$$TOFL \leq 1200m \tag{9.7}$$

Optimization results using five different regression strategies to determine the initial guess for y where successive substitution is used to converge the sizing residuals can be seen in Table 9.12.

Here, the Table-lookup, K Nearest Neighbors, and Extra Random Trees-based algorithms become stuck in local minima where, to the numerical tolerance of the optimizer, the optimality conditions are satisfied. In the other cases, the optimizer moves to points along the way that do not satisfy the optimality conditions, but instead terminates when it determines that the point can no longer be improved. As a result, comparing overall code effectiveness becomes difficult, save to say that

Table 9.12: Diesel-Electric Optimization Results(Successive Substitution)

loop evaluation	initial step	n_{size}	n_{opt}	$\frac{n_{size}}{n_{opt}}$	landing weight (lbs.)
Successive Substitution	Same Point	2,916	107	27.3	2,001
Successive Substitution	Table	1,000	121	8.26	2,153
Successive Substitution	KNN(5)	351	37	9.5	2,163
Successive Substitution	GPR	3,862	422	9.16	2,001
Successive Substitution	Extra Trees	356	37	9.62	2,161

the K Nearest Neighbors approach appears to be more computationally efficient than either the table-lookup method or when using unintelligent initial guesses. Table 9.13 shows the same results but Newton-Raphson is used to converge the problem, this time with a threshold of 2%, rather than the 5% employed in other parts of this dissertation.

Table 9.13: Diesel-Electric Optimization Results(Newton-Raphson)

loop evaluation	initial step	n_{size}	n_{opt}	$\frac{n_{size}}{n_{opt}}$	landing weight (lbs.)
Newton-Raphson	Same Point	1,960	79	24.8	2,007
Newton-Raphson	Table	9,384	898	10.5	2,003
Newton-Raphson	KNN(5)	875	72	12.2	2,034
Newton-Raphson	GPR	725	65	11.2	2,005
Newton-Raphson	Extra Trees	2,545	265	9.50	2,003

Results where Broyden's method is used to converge the system can be seen in Table 9.14. Again, a 2% threshold is used to switch from successive substitution to the Jacobian-based method here.

Table 9.14: Diesel-Electric Optimization Results(Broyden)

loop evaluation	initial step	n_{size}	n_{opt}	$\frac{n_{size}}{n_{opt}}$	landing weight (lbs.)
Broyden	Same Point	2,608	116	22.5	2,001
Broyden	Table	1,839	207	8.88	2,103
Broyden	KNN(5)	1,836	207	8.87	2,103
Broyden	GPR	1,841	207	8.89	2,103
Broyden	Extra Trees	1,836	207	8.87	2,103

An important facet of these cases is that, when Broyden's method works, it tends to be much faster than Newton-Raphson in converging the system due to the elimination of expensive finite difference calculations to compute the Jacobian. Unfortunately, in around 5% of the cases shown in Table 9.14, the solver becomes stuck and unable to fully converge the system; this is because, with too many iterations, the approximation to the Jacobian based on the Sherman-Morrison formula becomes less accurate; reinitializing the Jacobian should fix this particular issue and allow for more dependable results here. Additionally, this becomes a major issue for this case, because the unconverged aircraft

design occurs earlier on in the optimization process, when the optimizer makes the largest steps.

Note the identical values for the K Nearest Neighbors algorithm and Extra Random Trees; this is because both algorithms only take around 1 major optimization step where the regression algorithm is called; in this case, the steps are close enough to each other that the data are functionally identical. Each of the other regression techniques in Table 9.14 becomes stuck in the same local minimum. These results suggest that, without intelligently reinitializing the Jacobian in the iteration process, the loss of robustness from the use of Broyden's method may not make up for the potential increased code speed. Further study is recommended.

As an aside, comparing the relative computational cost reveals that the diesel-electric case is, in general, cheaper computationally than the aluminum-air case; this is because it uses a diesel engine, which possesses a high specific energy; from Chapter 8, one can readily see from the 2D sweep of wing loading vs combustion power (Figure 8.7), that, when the right numerical methods are applied, all designs converge. For the aluminum-air case, on the other hand, many of the design diverged to infinity which both consumed a large number of function evaluations as well as exacerbated the numerical difficulties experienced by the optimizer.

Chapter 10

Conclusions

Aircraft that utilize multiple energy systems possess not just fundamental physical challenges, but also numerical challenges that may not be immediately apparent to the men and women looking to model them. In particular, these designs seem to be prone to having multiple sizing solutions, even in the most basic cases when using commonly employed aircraft design procedures. Now, as Chapter 5 illustrates, this phenomena is not unique to these hybrid aircraft, and indeed, are applicable to a wide range of aircraft design cases; however, in practice, the low specific energy and power associated with these energy systems greatly exacerbates this; broad searches of the design space are associated with a wide range of aircraft sizes which increase the likelihood that extraneous solutions are found, at least when they exist. Furthermore, successive substitution, a method widely employed in solving aircraft design problems, proved incapable of converging certain systems when multiple energy systems are in play, especially when they are arranged in a series-hybrid format. This is because of oscillation that occurs in the solution process; the interaction between the two energy systems prevents the convergence of either. The solution to this problem is to either implement a robust backtracking algorithm, or preferably, solve the system using a root-finding solver such as Newton-Raphson that also utilizes a backtracking approach.

In addition to these issues, including another energy system into the design process significantly slows convergence. This is because of the interplay between the energy systems, which can be mitigated by a multidimensional zero-finding method such as Newton-Raphson or Broyden's method, as the work here demonstrates. Additionally, many aircraft design problems, including those in the work performed here can benefit greatly from the use of machine learning algorithms that learn the design space to predict intelligent initial guesses to what mzay be expensive iterative systems. Finally, this work revealed, that, when leveraging differences in specific energy and power between the energy systems, mere point performance is not enough; rather, a robust off-design performance model must be used to capture the full range of aircraft operating conditions. Merely using nominal values for specific energy and power often resulted in designs that were at times twice as heavy,

showing that they are inadequate in estimating aircraft size and making high level decisions. In summary, key contributions include:

- Identifying nonunique aircraft design solutions that may arise in a component weight buildup approach
  - Illustrating methodologies that possess a tendency to converge to undesirable results and ways to utilize them while avoiding the "bad," larger solution
- Demonstrating that, especially for serial-hybrid aircraft, systems often become unsolvable using fixed point iteration techniques
  - Requires a robust backtracking algorithm.
- Showing that code speed slows significantly (≈2-3x) when solving aircraft with multiple energy systems vs. just a single system as a result of oscillation in the numerics
  - Mitigated by use of multidimensional zero-finding methods (≈ 20% code speed improvement)
- Developing a framework to leverage machine learning regression techniques to learn the design space and start iterating from a more converged guess
  - Increased code speed by ≈ 30%
- Revealing that in many hybrid-electric and hybrid-battery cases, single point estimations of specific energy and power are inadequate in estimating aircraft weight (at times 2x heavier than when full range of battery performance incorporated)
  - With a battery Ragone plot, can perform an easy 1-D optimization problem to account for other sections of the discharge profile
- The framework developed can be readily used to enable robust, computationally tractable, large-scale explorations of the design space such as global optimization

All of the methodologies and frameworks for this system have been released into an open source aircraft design tool. Now, it should be said that many alternative, valid formulations exist for the aircraft design problems here that should in theory converge to the same result. Chapter 6 illustrates this in more detail. This work does not make the claim that the formulations are the best possible in any given situation. Rather, it takes methodology commonly used throughout academia as well as industry, applies it to novel problems, and thereby illustrates both its shortcomings in certain instances as well as avenues of significant improvement while maintaining the same overall structure, easily allowing its integration into current aircraft design tools.

Appendix A

Aircraft Sizing Problems for the Analytical Case

A.1 Electric Aircraft (Single Battery, no Power Requirements)

This case corresponds to an aircraft where the designer does not account for power requirements when sizing the battery, motor, and other parts of the energy system, and is representative of a "back-of-the-envelope" type approach. It is developed here to illustrate some of the mathematical characteristics of these problems. The battery weight is

$$W_{bat} = g * \frac{E1}{Esp}, \tag{A.1}$$

where Esp is set to 200*3,600 (the factor of 3,600 to convert from $\frac{W-h}{kg}$ to $\frac{J}{kg}$). Motor weight here is set to a constant 2,000; this may be unreasonably small depending on the aircraft. Again, this case is not meant to accurately design physical aircraft, but rather, to illustrate what can happen mathematically when solving these systems. With the battery and motor weights defined, and air mass associated with flying the mission taken from (5.7), the computed required weight associated with the input values is given in (A.2)

$$W_{out} = W_{empty} + W_{bat} + W_{air} + W_{mot} + W_{pl}, \tag{A.2}$$

where W_{pl} is set to 19,620 (or 2,000g). With weight defined, the computed output energy is then

$$E_{out,1} = t1 * P_{cruise} + t2 * P_{out}, \tag{A.3}$$

where $t1$ and $t2$ are chosen to be 4,000 and 300, respectively. Note that, for conventional VTOL aircraft, $t2$ would typically be much smaller than $t1$. With the output sizing values W_{out} and $E_{out,1}$ calculated, the associated output sizing vector is simply

$$y_{out} = \frac{[\frac{W_{out}}{g}, E_{out,1}]}{y_{scale}}, \tag{A.4}$$

where y_{scale} = [1E4, 1E8].The resulting residual that needs to be converged to zero for a sized aircraft is

$$h = \frac{y - y_{out}}{y}. \tag{A.5}$$

A.2 Electric Aircraft (Single Battery)

The problem formulation for this class of aircraft is virtually identical to that from Section A.1, with a couple of exceptions. Essentially, this design problem formulation takes into account power requirements, so rather than defining the battery weight from (A.1), (A.6) is used

$$W_{bat} = g * soft\_max(\frac{E_1}{Esp}, \frac{P_{in}}{Psp}), \tag{A.6}$$

where Esp is 200*3,600 and Psp is 1,000. Motor weight, rather than being set to a constant is specified from (A.7)

$$W_{mot} = \frac{P_{in}}{Psp_{mot}} \tag{A.7}$$

where $Psp_{mot} = 1600$, with the assumption that internal losses in the motor and propulsion system are taken into account via the specific power term. The computed output weight then is (A.2), only with different computed values for W_{bat} and W_{mot}. In addition, P_{out} is taken from (5.10), making the output sizing vector

$$y_{out} = \frac{[\frac{W_{out}}{g}, E_{out,1}, P_{out}]}{y_{scale}}, \tag{A.8}$$

where y_{scale} = [1E4, 1E8, 1E5].

A.3 Electric Aircraft (Hybrid Battery)

When adding an additional battery system, even the most basic sizing problems become considerably more complicated. In addition to the dependence on x and W_{in} on performance, an additional parameter called "f_{aux}" was defined here, which specifies the relative amount of power of each

energy system. Now, the specific energy (Esp) and specific power (Psp) are defined a-priori for each system, and the weight of each battery, using the soft max function ((5.4)-(5.6)) are then simply.

$$W_{bat,1} = g * soft\_max(\frac{E_1}{Esp_1}, \frac{P_{in} * (1 - f_{aux})}{Psp_1}) \tag{A.9}$$

$$W_{bat,2} = g * soft\_max(\frac{E_2}{Esp_2}, \frac{P_{in} * f_{aux}}{Psp_2}) \tag{A.10}$$

Specific energy and specific power values for the two batteries are representative of practically attainable aluminum air and lithium ion chemistries, respectively. To be more specific, for this aircraft class, Esp_1 = 1,300*3,600, Psp_1 = .2, Esp_2 = 200*3,600, Psp_2 = 1,000, where, like before, the factor of 3,600 represents a conversion from $\frac{W-h}{kg}$ to $\frac{J}{kg}$. The computed output weight W_{out}, like before, is a summation of the component weights

$$W_{out} = W_{empty} + W_{bat,1} + W_{bat,2} + W_{air} + W_{mot} + W_{pl}, \tag{A.11}$$

where, as before, W_{pl} is set as a constant 2,000. Now, in order to determine how much energy is drawn from each battery subsystem, the maximum power output of each subsystem must be defined, which are simply

$$P1 = \frac{W_{bat,1} * Esp_1}{g} \tag{A.12}$$

and

$$P2 = \frac{W_{bat,2} * Esp_2}{g}. \tag{A.13}$$

For a power limited battery 1, $P1$ is exactly $P_{in} * (1 - f_{aux})$, while a power limited battery 2 has power $P2 = P_{in} * f_{aux}$. Output parameters for the energy of each battery are defined below, where the minus signs are used to avoid having to define a "$soft\_min$" function.

$$E_{out,1} = t1 * -soft\_max(-P_{cruise}, -P1) + t2 * -soft\_max(-P_{out}, -P1) \tag{A.14}$$

$$E_{out,2} = t1 * soft\_max(P_{cruise} - P1, 0) + t2 * soft\_max(Pout - P1, 0) \tag{A.15}$$

where $t1$ and $t2$ are the same as before (see Section A.1). The auxiliary battery (battery 2 here) typically runs for a substantial portion of the mission depending on the value of f_{aux}. With this simple mission simulation, the required values for the sizing parameters is defined as

$$y_{out} = \frac{[\frac{W_{out}}{g}, E_{out,1}, E_{out,2}, P_{out}]}{y_{scale}} \tag{A.16}$$

where y_{scale} = [1E4, 1E8, 1E8, 1E5]. Now, an alternative formulation exists that removes these soft

max statements used above, and instead uses battery mass as an input. In this case, the values E_1 and E_2 are replaced with $W_{bat,1}$, $W_{bat,2}$, and E_1 (decoupling the relationship between air mass used and battery weight for the metal-air battery), making the sizing vector

$$y = [W_{in}, W_{bat,1}, W_{bat,2}, P_{in}, E_1]. \tag{A.17}$$

Here, one can use this formulation for y and allow them to be controlled by the optimizer, simultaneously optimizing and sizing the aircraft. This particular case will be explored more in Chapter 6.

A.4 Electric Aircraft (Single Battery, no Power Requirements, Simplified)

Now, in creating these aircraft design methodologies, it should be noted that the formulations shown in Sections A.1- A.3 are painfully nonunique. A large number of alternative formulations that would, in theory, converge to the exact same result exist. To help demonstrate this, the following aircraft design system, equivalent to that shown in Section A.1, is used. In this case only a single sizing parameter $y = [W_{in}]$ is used. Additionally, instead of treating energy as an output and defining a corresponding residual, battery weight is instead determined by backing out from the gap between empty weight and previously defined weights. Start by defining a term W_{energy}

$$W_{energy} = Wout - W_{empty} - W_{mot} - W_{pl} = W_{bat} + W_{air}, \tag{A.18}$$

which in turn can be used to determine the energy required via the equation

$$W_{energy} = g * E_1 * (mgf + \frac{1}{Esp}). \tag{A.19}$$

At this point, it becomes trivially easy to solve for E_1. W_{air} is then defined as in the previous sections.

$$W_{air} = g * mgf * E_1 \tag{A.20}$$

With W_{mot} defined as a constant as in A.1, the output weight is then, like before. A subtle difference here is that in order to ensure that the mission requirements are met in the converged case, battery weight is determined based on the output, rather than the input energy, i.e.

$$W_{bat} = g * \frac{E_{out,1}}{Esp} \tag{A.21}$$

where the parameter $E_{out,1}$ is determined from (A.3) as before. With battery and air weight determined, then the output weight can be found as in the previous cases as

$$W_{out} = W_{empty} + W_{bat} + W_{air} + W_{mot} + W_{pl}, \tag{A.22}$$

which is exactly the output y determined by the fixed point iteration scheme

$$y_{out} = \frac{\lceil \frac{W_{out}}{g} \rceil}{y_{scale}}, \tag{A.23}$$

where y_{scale} is simply [1E4]. As stated before, the primary difference between this case and the case outlined in Section A.1 is that the problem is formulated in such a way as to eliminate energy as an input parameter, reducing the dimensionality of the problem, and likely making it easier to solve.

Bibliography

[1] Intergovernmental panel on climate change. `http://www.ipcc.ch/`. Accessed: May 2014.

[2] Airbus. Airbus global market forecast 2017–2032 growing horizons, 2017.

[3] N.M. Alexandrov, J.E. Dennis, R.M. Lewis, and V. Torczon. A trust-region framework for managing the use of approximation models in optimization. *Structural optimization*, 15(1):16–23, 1998.

[4] K.R. Antcliff and F. Capristan. Conceptual design of the parallel electric-gas architecture with synergistic utilization scheme (PEGASUS) concept. In *18th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Denver, Co, 2017. AIAA AVIATION Forum.

[5] K.R. Antcliff, M.D. Guynn, T.V. Marien, D.P. Wells, S.J. Schneider, and M.T. Tong. Mission analysis and aircraft sizing of a hybrid-electric regional aircraft. In *AIAA Scitech*, San Diego, CA, January 2016.

[6] N. E. Antoine. *Aircraft Optimization for Minimum Environmental Impact*. PhD thesis, Stanford University, ProQuest Information and Learning Company 300 North Zeeb Road P.O. Box 1346 Ann Arbor, Ml 48106-1346, 2004.

[7] L. Armijo. Minimization of functions having Lipschitz continuous first partial derivatives. *Pacific J. Math.*, 16(1):1–3, 1966.

[8] S. Ashcroft, A. Padron, K. Pascioni, G. Stout, and D. Huff. Review of propulsion technologies for n+3 subsonic vehicle concepts. NASA Glenn Research Center, 2011.

[9] S. Banach. Sur les operation dans les ensembles abstraits et leur application aux equations integrales. *Fundamenta Mathematicae*, 3:133–181, 1922.

[10] A.D. Belegendu and T.R. Chandrupatla. *Optimization Concepts and Applications in Engineering*. Cambridge, Cambridge, New York, 2nd edition edition, 2011.

[11] G. Biros and O. Ghattas. Parallel Lagrange–Newton–Krylov–Schur methods for PDE-constrained optimization. part II: The Lagrange–Newton solver and its application to optimal control of steady viscous flows. *SIAM Journal on Scientific Computing*, 27(2):714–739, 2005.

[12] Boeing Commercial Airplanes. Current market outlook: 2012–2031. `http://www.boeing.com/cmo`, 2012.

[13] E. Botero, A.D. Wendorff, T. MacDonald, Anil Variyar, J.M. Vegh, J.J. Alonso, T.H. Orra, and C. Ilario da Silva. SUAVE: An open-source environment for conceptual vehicle design and optimization. In *AIAA Scitech*, San Diego, CA, January 2016.

[14] M. Bradley and C. Droney. Subsonic ultra green aircraft research: Phase I final report. Technical report, Boeing, April 2011.

[15] M. Bradley and C. Droney. Subsonic ultra green aircraft research Phase II: N+4 advanced concept development. Technical report, Boeing, May 2012.

[16] L. Breiman. Bagging predictors. *Mach. Learn.*, 24(2):123–140, August 1996.

[17] P.G. Bruce, S. A. Freunberger, J. J. Hardwick, and J. M. Tarascon. Li-O2 and Li-S batteries with high energy storage. *Nature Materials*, 11(1), 2012.

[18] D. Buecherl, I. Bolvashenkov, and H.G. Herzog. Verification of the optimum hybridization factor as design parameter of hybrid electric vehicles. In *IEEE Vehicle Power and Propulsion Conference*, Dearborn, MI, September 2009.

[19] V. Caramia and B. Bozzini. Materials science aspects of zinc–air batteries: a review. *Materials for Renewable and Sustainable Energy*, 3(2):28, 2014.

[20] A. Cauchy. Methode generale pour la resolution des systemes dequations simultanes. *C.R. Acad. Sci. Par*, 25, 1847.

[21] M. Chen and G. A. Rincon-Mora. Accurate electrical battery model capable of predicting runtime and iv performance. *IEEE transactions on energy conversion*, 21(2):504–511, 2006.

[22] T. Christen and M.W. Carlen. Theory of Ragone plots. *Journal of Power Sources*, 91(2):210 – 216, 2000.

[23] J. Cook. How to compute the soft maximum. John D. Cook Consulting, 2010.

[24] T. Cover and P. Hart. Nearest neighbor pattern classification. *IEEE Trans. Inf. Theor.*, 13(1):21–27, September 2006.

[25] N. Daidzic, L. Piancastelli, and A. Cattini. Diesel engines for light-to-medium helicopters and airplane (editorial). volume 1, August 2014.

[26] A. Datta and W. Johnson. Requirements for a hydrogen powered all-electric manned helicopter. Moffett Field, CA, 2011. NASA Ames Research Center.

[27] M. Drela. Development of the D8 transport configuration. In *29th AIAA Applied Aerodynamics Conference*, Honolulu, Hi, Jun 2011. American Institute of Aeronautics and Astronautics.

[28] M. Duffy, S. Wakayama, and R. Hupp. A study in reducing the cost of vertical flight with electric propulsion. In *17th AIAA Aviation Technology, Integration, and Operations Conference*, Denver, Co, 2017. AIAA AVIATION.

[29] M.S. Eldred and D.M. Dunlavy. Formulations for surrogate-based optimization with data-fit, multifidelity and reduced-order models. In *Proceedings of the 11th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, No. 2006–7117 in AIAA Paper*, 2006.

[30] Federal Aviation Administration. Current aviation excise tax structure, 2017.

[31] P.J.S.G. Ferreira. Fixed point problems - an introduction. *Revista do Detua*, 1(6):505–513, 1996.

[32] F.R. Fink. Airframe noise prediction method. Technical report, Federal Aviation Administration, 1977.

[33] W. Fredericks. Aircraft conceptual design using vehicle sketch pad. In *AIAA Scitech*, Orlando, FL, January 2010.

[34] W. Fredericks, M. Moore, and R. Busan. Benefits of hybrid-electric propulsion to achieve 4x increase in cruise efficiency for a vtol aircraft. Hampton, VA, 2013. AIAA Aviation Technology, Integration, and Operations (ATIO) Conference.

[35] J. H. Friedman. Greedy function approximation: A gradient boosting machine. *Annals of Statistics*, 29:1189–1232, 2000.

[36] P. Geurts, D. Ernst, and L. Wehenkel. Extremely randomized trees. *Machine Learning*, 63(1):3–42, 2006.

[37] P.E. Gill, W. Murray, and M.A. Saunders. SNOPT: An SQP algorithm for large-scale constrained optimization. *SIAM Journal on Optimization*, 47(1):99–131, 2002.

[38] J.C. Gladin, C. Perullo, J.C.M Tai, and D.M. Mavris. A parametric study of gas turbine propulsion as a function of aircraft size class and technology level. In *55th AIAA Aerospace Sciences Meeting*, Grapevine, Texas, 2017.

[39] M. Green, B. Schiltgen, and A. Gibson. Analysis of a distributed hybrid propulsion system with conventional electric machines. In *AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit, Joint Propulsion Conferences*, Atlanta, GA, January 2012.

[40] G. Grewal, S. Coros, D. Banerji, and A. Morton. Comparing a genetic algorithm penalty function and repair heuristic in the DSP application domain. In *IASTED International Conference on Artificial Intelligence and Applications*, pages 31–39, Innsbruck, Austria, January 2006.

[41] A. Griewank, , and C. Faure. Reduced functions, gradients and Hessians from fixed-point iterations for state equations. *Numerical Algorithms*, 30(2):113–139, Jun 2002.

[42] S. Gunther, N.R. Gauger, and Q. Wang. Simultaneous single-step one-shot optimization with unsteady PDEs. *Journal of Computational and Applied Mathematics*, 294:12 – 22, 2016.

[43] M. Hazen. *Search Strategies for Global Optimization.* PhD thesis, University of Washington, 2008.

[44] M. Hepperle. Electric flight- potential and limitations. German Aerospace Center, 2012.

[45] J. Heywood. *Internal Combustion Engine Fundamentals.* McGraw-Hill, New York, NY, 2nd edition, 1988.

[46] J. Holden and N. Goel. Uber elevate: Fast-forwarding to a future of on-demand urban air transportation. Technical report, Uber, April 2017.

[47] J. Hruska. New aluminum air battery could blow past lithium-ion, runs on water. Extreme Tech. Accessed: Mar. 15, 2015.

[48] R Shevell I. Kroo. AA241A course notes: Aircraft design, synthesis and analysis. `http://adg.stanford.edu/aa241/AircraftDesign.html`. Accessed: Nov. 11, 2014.

[49] A.T. Isikveren, S. Kaiser, C Pornet, and P.C. Vratny. Pre-design strategies and sizing techniques for dual-energy aircraft. *Aircraft Engineering and Aerospace Technology: An International Journal*, 86(6):525–542, 2014.

[50] P.W. Jansen and R.E. Perez. Constrained structural design optimization via a parallel augmented Lagrangian particle swarm optimization approach. *Computers and Structures*, 89(13):1352 – 1366, 2011.

[51] L. Johnson. The viability of high specific energy lithium air batteries. Technical report, Excellatron Solid State LLC., October 2010.

[52] W. Johnson. NDARC - NASA design and analysis of rotorcraft: Validation and demonstration. In *American Helicopter Society Aeromechanics Specialists Conference*, 2010.

[53] W. Johnson, S. Silva, and E. Solis. Concept vehicles for vtol air taxi operations. In *AHS Technical Conference on Aeromechanics Design for Transformative Vertical Flight*, San Francisco, CA, January 2018.

[54] JPDO. Next generation air transportation system - integrated plan. Technical Report, Joint Planning and Development Office (JPDO), 2011.

[55] J. Kennedy and R.Eberhart. Particle swarm optimization, 1995.

[56] G. Khalil. Challenges of hybrid electric vehicles for military applications. In *IEEE Vehicle Power and Propulsion Conference*, Dearborn, MI, September 2009.

[57] I. Kroo, S. Altus, R. Braun, P. Gage, and I. Sobieski. Multidisciplinary optimization methods for aircraft preliminary design. Hampton, VA, January 1994. AIAA Paper -94-4325-CP.

[58] L. Lasdon, R. Fox, and M. Ratner. Nonlinear optimization using the generalized reduced gradient method. *Rev. Franc. Automat. Inform. Rech. Operat., V*, 8(3):73–104, 1974.

[59] S. Levine. Two big labs step back from the most promising next-generation battery. Quartz, May 2014. Accessed: September 2014.

[60] A. Liaw and M. Wiener. Classification and regression by randomForest. *R News*, 2(3):18–22, 2002.

[61] Y. Liu, Q. Sun, W. Li, K.R. Adiar, J. Li, and X. Sun. A comprehensive review on recent progress in aluminum-air batteries. *Green Energy and Environment*, pages 246–277, June 2017.

[62] D.G. Luenberger and Y. Ye. *Linear and Nonlinear Programming*. Springer Publishing Company, Incorporated, 2015.

[63] T. Lukaczyk, A. Wendorff, E. Botero, T. MacDonald, T. Momose, A. Varyar, J.M. Vegh, M. Colonno, T Orra, C. Illaria da Silva, and J. Alonso. SUAVE: An open-source environment for multi-fidelity conceptual vehicle design. Dallas, TX, 2015. AIAA Aviation Forum 2015.

[64] H. Ma, Y.and Zhang, B. Wu, M. Wang, X. Li, and H. Zhang. Lithium sulfur primary battery with super high energy density: Based on the cauliflower-like structured c/s cathode. *Nature*, 2015.

[65] J.R.A. Martins and R. Lambe. Multidisciplinary design optimization: A survey of architectures. *AIAA Journal*, 51(9):2049–2075, 2013.

[66] M. Marwa, S.M. Martin, C.M. Borja, and R.P. Anderson. Analytic and numeric forms for the performance of propeller-powered electric and hybrid aircraft. In *55th AIAA Aerospace Sciences Meeting*, Grapevine, Tx, January 2017. AIAA.

[67] R.A. McDonald. Fundamental sizing implications of constant or increasing weight aircraft. In *12th AIAA Aviation Technology, Integration, and Operations (ATIO) Conference and 14th AIAA/ISSM*, Indianapolis, Indiana, 2012.

[68] M. Moktar, M.Z.M. Talib, E.H. Majlan, S.M. Tasirin, W.M.F.W. Ramli, W.R.W. Daud, and J. Sahari. Recent developments in materials of aluminum air batteries: A review. *Journal of Industrial and Engineering Chemistry*, 2015.

[69] M. Moore and B. Fredericks. Misconceptions of electric propulsion aircraft and their emergent aviation markets. In *AIAA SciTech Forum*, National Harbor, Maryland, 2014. NASA Langley Research Center.

[70] L.F. Nazar, M. Cuisinier, and Q. Pang. Lithium-sulfur batteries. *Materials Research Society Bulletin*, 39, May 2014.

[71] J. Nocedal and S.J. Wright. *Numerical Optimization*. Springer, 2nd edition, 2006.

[72] T.E. Noll, J.M. Brown, M.E. Perez-Davis, Ishmael S.D., G.C. Tiffany, and M. Gaier. Investigation of the Helios prototype aircraft mishap. Technical report, NASA, January 2004.

[73] N. Omar, P.V. Bossche, and J.V. Mierlo. Peukert revisited - critical appraisal and need for modification for lithium-ion batteries. *Energies*, 6(11), 2013.

[74] M. Patterson, M. Daskilewiczy, and B. German. Conceptual design of electric aircraft with distributed propellers: Multidisciplinary analysis needs and aerodynamic modeling development. In *AIAA Scitech*, National Harbor, MD, January 2014.

[75] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

[76] J. Penner, D. Lister, D. Griggs, D. Dokken, and McFarland M. Aviation and the global atmosphere. Technical report, Intergovernmental Panel on Climate Change, 1999.

[77] R.E. Perez, P.W. Jansen, and J.R.R.A. Martins. pyOpt: A Python-based object-oriented framework for nonlinear constrained optimization. *Structures and Multidisciplinary Optimization*, 45(1):101–118, 2012.

[78] C. Pornet, C. Gologan, P.C. Vratny, A. Seitz, O. Schmitz, A.T. Isikveren, and M. Hornung. Methodology for sizing and performance assessment of hybrid energy aircraft. *Journal of Aircraft*, 52(1), 2015.

[79] C.E. Rasmussen and C.K.I. Williams. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press, 2005.

[80] D.P. Raymer. *Aircraft Design: A Conceptual Approach*. AIAA, Playa del Ray, California, 4th edition, 2006.

[81] Eviation. Alice commuter. `https://www.eviation.co/alice/`. Accessed: 2018-05-04.

[82] R. Revel, T. Audichon, and S. Gonzalez. Non-aqueous aluminium-air battery based on ionic liquid electrolyte. *Journal of Power Sources*, pages 415–421, November 2014.

[83] J.F. Rodriguez, J.E. Renaud, B.A. Wujek, and R.V. Tappeta. Trust region model management in multidisciplinary design optimization. *Journal of Computational and Applied Mathematics*, 124(1-2):139 – 154, 2000. Numerical Analysis 2000. Vol. IV: Optimization and Nonlinear Equations.

[84] J. Roskam. *Airplane Design: Part VI, Preliminary Calculation of Aerodynamic, Thrust and Power Characteristics*. Design, Analysis and Research Corporation, Lawrence, Kansas, 2008.

[85] S. Salcedo-Sanz. A survey of repair methods used as constraint handling techniques in evolutionary algorithms. *Computer Science Review*, 3(3):175 – 192, 2009.

[86] J. Sherman and W.J. Morrison. Adjustment of an inverse matrix corresponding to changes in the elements of a given column or a given row of the original matrix. In *Annals of Mathematical Statistics*, volume 20.

[87] R.S. Shevell. *Fundamentals of Flight*. Prentice-Hall, 2nd edition, 1989.

[88] J.D. Sinsay, B. Tracey, J.J. Alonso, D.K. Kontinos, J.E. Melton, and S. Grabbe. Air vehicle design and technology considerations for an electric vtol metro-regional public transportation system. In *12th AIAA Aviation Technology, Integration, and Operations (ATIO) Conference and 14th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Indianapolis, IN, 2012. AIAA.

[89] A. J. Smola and B. Schölkopf. A tutorial on support vector regression. *Statistics and Computing*, 14(3):199–222, August 2004.

[90] C. Snyder, J. Berton, G. Brown, and et. all. Propulsion investigation for zero and near-zero emissions aircraft. In *NASA STI Program*. NASA Glenn, 2009.

[91] A.M. Stoll, J. Bevert, P.P. Pei, and E.V. Stilson. Conceptual design of the Joby S2 electric vtol pav. In *Aviation Technology, Integration, and Operations Conference*, Atlanta, GA, June 2014. AIAA.

[92] S. Stückl, J. van Toor, and H. Lobentanzer. Voltair: The all electric propulsion concept platform-a vision for atmospheric friendly flight. In *28th International Congress of the Aeronautical Sciences*. EADS, 2011.

[93] K. I. Swartz. Charging forward: New eVTOL concepts advance. *Vertiflight*, July/August 2014.

[94] W. Tacke and S. Gao. Kitty Hawk and more e-flight at airventure. *eVtol*, May 2017.

[95] E. Torenbeek. *Advanced Aircraft Design: Conceptual Design, Technology and Optimization of Subsonic Civil Airplanes*. Wiley Press, 2013.

[96] K. Varesi and A. Radan. A novel methodology proposed for optimizing the degree of hybridization in parallel hevs using genetic algorithm. *World Academy of Science, Engineering and Technology International Journal of Computer, Electrical, Automation, Control and Information Engineering*, 5(8), 2011.

[97] J.M. Vegh and J. Alonso. Design and optimization of short-range aluminum-air powered aircraft. In *54th AIAA Aerospace Sciences Meeting*, San Diego, CA, 2016. AIAA Scitech.

[98] J.M. Vegh, J. Alonso, T. Orra, and C. Ilario da Silva. Flight path and wing optimization of lithium-air battery powered passenger aircraft. In *AIAA 2015-1674 53rd AIAA Aerospace Sciences Meeting*, Kissimee, FL, 2015. AIAA Scitech.

[99] J.M. Vegh, J. Alonso, and J. Sinsay. Modeling of diesel and diesel-electric hybrid propulsion systems for conceptual design of rotorcraft. In *Technical Meeting on Aeromechanics Design for Vertical Lift*, San Francisco, CA, 2016. American Helicopter Society.

[100] J. Voelcker. Phinergy 1000-mile aluminum-air battery: On the road in 2017? `http://www.greencarreports.com/`, 2013. Accessed: Oct. 26, 2015.

[101] T. Weise. Global optimization algorithms – theory and application. `http://www.it-weise.de/projects/book.pdf`, 2008.

[102] S. Yang and H. Knickle. Design and analysis of aluminum/air battery system for electric vehicles. Journal of Power Sources, 2002.