Final Report:
Collins Aerospace Galley Thermal & Power Optimization

December 14, 2020

Prepared by:
2020 Multidisciplinary Design Project Team
Mandappa Biddanda
Nicholas Leginza
Jordan Mayor
Julian Tarazi
In Beom Yoon

Prepared for:
Collins Aerospace
Project Sponsor: Brian St. Rock
Project Mentor: Matthew Pearson

Collins Aerospace

COLLEGE OF ENGINEERING
MULTIDISCIPLINARY DESIGN PROGRAM
UNIVERSITY OF MICHIGAN

Table Of Contents

List of Figures

List of Tables

Abbreviations and Acronyms

ABS:	Acrylonitrile butadiene styrene
ACU:	Air Chilling Unit

ATLAS:	Air France, Union de Transports Aériens, Lufthansa, Alitalia, and Sabena. Widespread aircraft standard
CAD:	Computer-aided Design
CoP:	Coefficient of Performance
FAA:	Federal Aviation Administration
FEA:	Finite Element Analysis
GUI:	Graphical User Interface
HCP:	Honeycomb Paneling
MEA:	More Electric Aircraft
PC:	Power Consumption
PDE	Partial Differential Equation
PU:	Polyurethane
UI:	User Interface

Executive Summary

The galley represents the largest electrical load on an aircraft. This load leads to a major consumption of fuel, a valuable resource to airlines both for cost and emissions reduction. Additionally, reducing power demands in the galley creates room for future novel electrical systems, which is important as the industry shifts to More Electric Aircraft (MEA) - replacing hydraulic/pneumatic systems with modern electric ones. Thus, we were tasked with improving the electrical and thermal systems that support the galley to help reduce its power consumption. The impact of our work can be quantified using two main metrics: the decrease in average power consumption and the decrease in maximum power consumption.

This project provided an opportunity to not only research potential technologies to reduce galley energy usage, but to also create a platform to enable future ones to be modeled and tested. As a result of this, as well as the lack of available information regarding the current power consumption of airplane galleys, a large portion of this project was dedicated to creating a baseline power consumption model to help understand the energy profile of the galley and test our solutions. Our baseline model is modular and fully functional; this allowed us to evaluate the impact of three different technologies: aerogel insulation to mitigate heat loss, variable air flow to remove heat more efficiently, and a power controller to more intelligently manage power consumption. These solutions work with one another to reduce the average power consumption over 100 flights by 12.8%. In addition, the solutions reduced the peak power consumption over 100 flights by 27.4%. This meant we achieved our highest priority requirements, which were to reduce these power consumptions by 10% and 25% respectively.

The remainder of this report goes over an in-depth overview of our model, solutions, validation methods, results, and implications as well as next steps.

1. Baseline Model

The largest portion of the project was dedicated to the development of a baseline model of an existing galley's power consumption, integrated with a thermal model of each of the individual inserts in the galley. The model simulates insert usage on a per-second basis and records individual and collective power consumption of the galley inserts. The model takes into account the insert usage, electrical characteristics, and the thermal interaction of inserts with each other and with the galley environment to accurately model power consumption. Additionally, the creation of the baseline model provided a consistent platform for the team to integrate and test solution concepts and objectively compare results to the baseline power consumption values.

1.1 Galley Overview

The galley that we modeled for our project is the G2F galley in an Airbus A350 for Qatar Airways, Fleet QTR*MN.

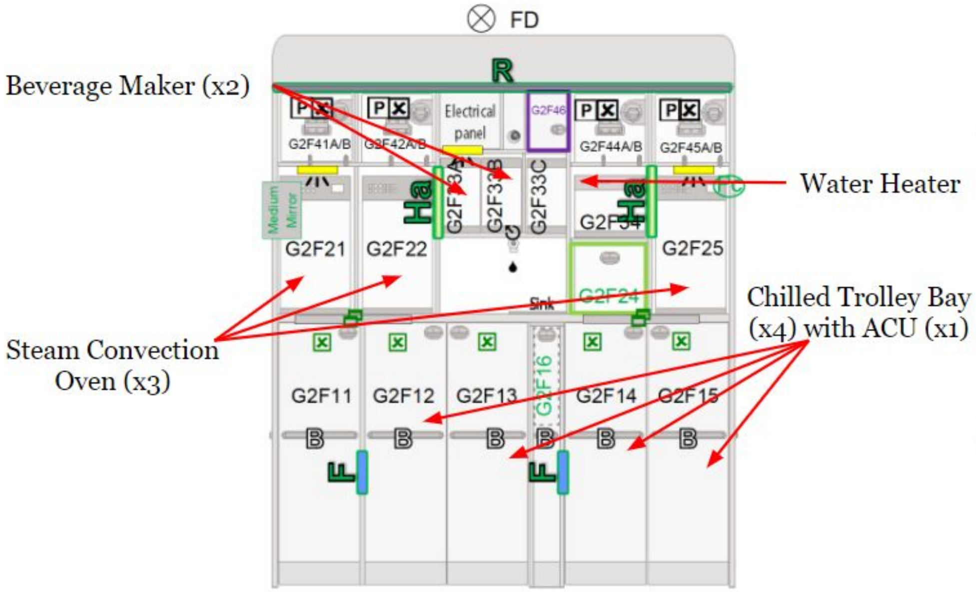

Figure 1.1-1: Galley G2F configuration

The components modeled are three steam ovens, two beverage makers, one water heater and one air chilling unit (ACU) servicing the four chilled cart bays. The configuration of these inserts can be seen in Figure 1.1-1. Although the G2F galley can also contain a bun warmer, it was decided early on in the project with the sponsors at Collins that our focus should be on the aforementioned inserts.

1.2 Assumptions

While some information used in the model was pulled from spec sheets provided by our sponsor at Collins Aerospace and partner companies, many assumptions had to be made throughout the project to fill in areas without concrete information. These assumptions were based on research using online resources, conversations with our project sponsor at Collins, and personal experience.

1.2.1 General Flight

When modeling a typical flight, we decided to keep the flight time constant for every simulation, and settled on an 8-hour flight as this is the length of the average international flight [1]. We modeled the number of passengers on board the plane as a normal distribution with a mean of 300 people and a standard deviation of 10 people (we round up a decimal number of passengers). 88% of A350 aircraft in operation today are the A350-900 model, which has a capacity ranging from 300-350 people [2, 3]. Thus, we went with the middle of this range, 325, and assumed that an average flight had 25 seats unfilled, bringing the mean down to 300. The assumption on the number of unfilled seats came from us, using personal experience with international flights, as the statistics on the internet regarding the fullness of flights is affected quite heavily by domestic flights on small planes. The standard deviation chosen makes it so that the chance that the number of passengers above 325 is 0.6%, and the chance that it's above 350 is essentially zero.

We also model two meal service periods, which aligns with our experience on international flights. The first meal tends to be served two hours after takeoff and the second meal within two hours of landing [4]. Therefore we picked the second meal to be served 90 minutes before landing, and decided that the meal preparation would begin 30 minutes before a meal is served. In terms of galley environment conditions, we assume the galley temperature is 21.1ºC, which is standard room temperature. We also assume that there is a ¾" honeycomb fiberglass panel between inserts in the galley, which isn't insulated.

1.2.2 Oven

When it comes to the frequency of oven use, we made the assumption that all passengers eat a hot meal during the first meal serving period. In our experience, the second meal on international flights tends to be smaller and oftentimes either does not include any heated items, or maybe has one at most, and thus the oven use is not needed for the second meal period and is not included in the model. The ovens on board are able to heat 32 standard ATLAS meals in 24 minutes according to the Essence inserts brochure distributed by our sponsor at Collins. There are 15 steam ovens on board, three in galley G2F, and we assume all of the ovens on board are used to heat meals.

When it comes to the power demands of the oven, we assume that the ovens are preheated once at the beginning of the meal preparation period to 176.67°C, and assume they remain on (while cycling power to keep the internal temperature between 168.33°C and 185°C) for as long as it takes for all the meals to be prepared. The preheat start time (when an oven first turns on) of an oven is randomly generated within the first 15 minutes of meal preparation.

From a thermal standpoint, we assume there is a 0.5mm thick Stainless Steel 304 sheet with porcelain enamel lining the inside of the oven, and a 0.5mm thick Stainless Steel 304 sheet on the outside [5]. We assume ceramic fiber is used for insulation, as it is the first result that appears when researching oven insulation. The insulation thickness for the front (16mm), top (86mm), bottom (48mm), and back (48mm) sides of the oven were measured from the spec sheets provided by our sponsor at Collins, assuming the schematic is to scale. The insulation thickness for the sides (22mm) was calculated by taking the outer insert dimensions and subtracting it by the inner dimensions. We also assume that steam is not present during the preheating process, and is rather injected once the meals are in the oven [6]. The heating mechanism in the oven is assumed to be a nichrome heating element with a surface area of 0.07m^2. The temperature profile of the heating element was taken from observations from household ovens. The oven model uses conductive, convective, and radiative heat transfer equations to model the temperature change of the oven.

1.2.3 Hot Beverage Maker

For the beverage makers, we modeled the percentage of passengers that require the use of a beverage maker as a normal distribution with mean 29% and standard deviation of 5%. This was chosen because an average of 29% of people drink coffee on any given day [7], and without data on which beverages are most commonly prepared using the hot beverage maker, we decided to use coffee since in our experience, it's the most common hot beverage ordered on a flight alongside tea. The 5% standard deviation was decided upon to give some variation in the percentage of passengers ordering coffee without deviating too far from the mean. In addition, we assumed from personal experience that most airlines use "picco" sized cups, which have a capacity of 6.5oz (195mL), so each passenger ordering coffee consumes that much in volume. For reference, the capacity of the beverage maker is 4.3L according to its spec sheet, and there are 12 beverage makers on board the aircraft. Its brew cycle time is 2 minutes and 45 seconds according to the Essence inserts brochure.

Unlike the oven, we assume hot drinks are served primarily over both meal periods, so the vast majority of coffee consumption happens around those meals. However, it is not uncommon that some passengers order coffee during the middle stretch of the flight, and those orders tend to be placed by individuals at random times. Thus, we model this by generating a random percentage between 0-25% and saying that this percentage of the passengers who order coffee during meals also order a cup in between the meal periods. We assume that one passenger is served with each use of the beverage maker in between meals, so a random brew time is generated for each of the orders.

In order to avoid the problem of two uses of the hot beverage maker coinciding by accident (since random number generators cannot be predicted), we divided up the flight time into 15-minute blocks where coffee can be brewed and served. The start of the brewing time is randomly generated within the first 10 minutes of these blocks.

With respect to the power demands of the hot beverage maker, we assume that coffee is brewed to 87.8°C to be served immediately, and thus there is not any power cycling to maintain a certain temperature over a significant period of time. Its temperature increase each second during brewing is hardcoded to reflect the brew cycle time of 2 minutes 45 seconds. The beverage maker is either "on" and drawing maximum power, or in its standby mode. The beverage maker lacks more thorough thermal modeling due to a lack of information. Its spec sheet provides information on the flow rate of water, but no information on the heating mechanism. Another large missing piece of information was whether there is a tank of water that maintains hot water or if the beverage maker heats up water as it dispenses a beverage. Since the maximum flow rate of water for the beverage maker is 0.60L/min - roughly half a tablespoon per second or 1/20th of the average household sink faucet flow rate - we made the assumption that the beverage maker heats up the water as it dispenses it and that the heating element or heating plate is small enough to be negligible in thermal interactions between inserts.

1.2.4 Water Heater

For the water heater, we modeled the percentage of passengers that drink tea (make use of the water heater) as a normal distribution with mean 49% and standard deviation of 5%. This was chosen because an average of 49% of people drink tea on any given day [8]. The 5% standard deviation was decided upon to give some variation in the percentage of passengers ordering tea without deviating too far from the mean. In addition, we assumed from personal experience that most airlines use “picco” sized cups, which have a capacity of 6.5oz (195mL), so each passenger ordering tea is served that much in volume. However, at meal time, tea tends to be kept in insulated containers and brought through the cabin, so we assume that half the tank is drained per “use” of the water heater during mealtimes. For reference, the capacity of the water heater is 3L according to the water heater specification document, and there are four water heaters on board the aircraft.

We also assume hot drinks are served primarily over both meal periods, so the vast majority of tea consumption happens around those meals. However, for passengers who order between meals, we model usage by generating a random percentage between 0-25% and say that this is the percentage of passengers who order tea during meals that also order it in between meal periods. For these “uses” of the water heater in between meals, only a cup’s worth is drained from the tank. Also, we assume one passenger is served with each use of the water heater during non meal time.

We assume the water heater is turned on during the first meal period, and then remains on until the second one has ended. The water heater has two power modes depending on how it's configured, and we assume it’s operating in its lower power mode. Also, when turned on, the water heater continuously maintains a temperature of 90°C, and power cycles as necessary to keep the temperature relatively constant. We assume that the water heater refills when the tank is 50% full, and we estimate that it refills by 0.06L (2% of the tank capacity) per second. The water coming in is at the same temperature as the galley environment, and the temperature of the heated water in the tank is lowered accordingly as they mix.

From a thermal standpoint, we assume that the inner tank is lined with 1mm of stainless steel 304 and that the outer tank is 0.5mm galvanized steel. We also assume there is 25mm thick fiberglass insulation for the tank. The water heater model uses conductive and natural convective heat transfer equations to model the temperature change of the water.

1.2.5 Air Chilling Unit

We assume the ACU in the galley is turned on at the beginning of the flight and remains functional for the duration of the flight. When food is loaded onto a plane, it is already cooled and remains on an auxiliary power unit, so the cart bay and all of its contents is already at 4°C at the beginning of the flight. In terms of power, the ACU cycles power to maintain a temperature between 2°C and 4°C, and we made the assumption that its coefficient of performance (CoP) is 0.8; this is corroborated with our sponsor at Collins. In addition, every 10 minutes the ACU runs a defrost cycle which sees it draw maximum power and the defrost cycle runs for 1.5 mins. When the ACU is turned on or running a defrost cycle it draws 410W. When the ACU is turned on, it has a volumetric flow rate of 120L/s which is then distributed into 3 cart bays--one of the bays stores two carts--through 4 inlet vents. With a CoP of 0.8, the ACU can remove a total of 328W of heat at any given time. Since the ACU is sized for approximately the maximum amount of heat it can remove, we can make the assumption that at any given time, the heat entering the cart bays should not exceed 328W. Using this we can validate the accuracy of our model by examining the amount of heat entering the cart bays per second on a flight to ensure that the number is comparable to 328W. We allowed a minor amount of instances in which the heat can surpass this number, such as if the carts' and their contents were below a certain temperature or when the cart bay doors were opened. The ACU measures the average of the return air from the cart bays, and uses this measurement to decide when to cycle power.

We assume the volume of the cart bay is calculated from standard cart size [9], plus clearances of 25mm to the front, rear, top, and each side, and 76mm (3") on the bottom because of the wheels. The CAD model shows a close size, but was not used as the only assumption source due to the model's omissions of doors and the floor. One ACU cools all the cart bays in the galley; in our case, that means 3 cart bays with 4 inlet and return air vents where one cart bay holds two carts and two cart bays hold one cart. The cart bay is lined with ABS plastic and is layered directly on top of the honeycomb structural paneling and a ¼" layer of high density polyurethane (PU) insulation that lines the top of the cart bay, as well as inside of the cart bay doors. We also assume that there is leakage of cabin air into the cart bay, and the total area that is exposed to the cabin air is approximately 2mm multiplied by the width of the cart bays. When the carts are removed from the cart bay, we assume there is an instantaneous influx of cabin air into the cart bay.

The dimensions of the cart were taken from standard cart size [9], and the carts are made fully of aluminum with a 1cm thickness. We modeled our carts as a hollow box with one open side. Note that some airlines use carts that have insulation or other materials for the sides and the geometry of the carts may differ with added features.

The drinks are often served before meals are and as such, we model their service 45 minutes to 30 minutes before the first meal prep period and during the second meal prep period. This allows the drinks to be fully served and gives time for meals to be prepared. Two carts are removed during the first meal while the other two are removed during the second meal. One galley services a quarter of economy class, so using the seating plan for an aircraft in this fleet, we subtract 36 passengers and divide by four to find the number of passengers that will be serviced by one galley - or two carts - during a service period. Using the assumption that it takes on average 15 seconds to serve a passenger we can find the time that a cart is in service.

2. Solution Concepts

Once the baseline model was fully functional, we began implementing solutions that were directed towards decreasing the average power consumption and peak power consumption. The solutions are a power controller, aerogel insulation, variable air flow, and multiple ACUs. The related assumptions are mentioned below.

2.1 Power Controller

The purpose of the power controller is to add a layer of intelligence to power distribution in the galley. In the galley's current state, power requests made by inserts are automatically granted. This allows for large power peaks and little limit on power consumption. The idea behind the power controller is that limiting specific inserts' access to power at the right time can lead to a large reduction in peak power usage with only minor impacts on galley functionality. As such, it is important to think of the power controller as a delicate give-and-take between energy gains and original intended functionality. Below, we have outlined our implementation of a power controller, but we have left plenty of variables flexible in order to accommodate the needs and responsiveness of a specific galley.

2.1.1 Priority Based Control

The power controller prevents power peaks using a priority based control method. Every second the controller analyzes which inserts are requesting power. These inserts are then assigned a priority value based on information like if the insert is in its temperature bounds or if it has been delayed before. All inserts requesting power are then sorted according to their priority values and granted power sequentially until a maximum power allowance is met. This process is repeated throughout the flight with insert priorities efficiently fluctuating so that all inserts are serviced in a timely manner. In the case that powering an insert is of utmost importance, the insert can be assigned a maximum priority, and thus any power request made is granted regardless of the power allowance.

2.1.2 Standby Shutoff

According to the baseline model, standby power consumption accounts for around 9% of total power consumption in the galley. Standby mode is important if an insert is going to be used imminently as it ensures responsiveness to a user request; otherwise, it is an unnecessary use of power. As such, the team implemented a standby shutoff for all inserts excluding the ACU during takeoff and landing (first and last 45 minutes of the flight) as these inserts aren't being used anyways. Furthermore, the power controller turns an insert from standby to off if it has not been used for over 30 minutes.

2.1.3 Assumptions

The controller is set to allow a maximum power of 9500W. The priority of an insert relies on the basic assumption that the power controller's job is to limit power peaks while ensuring that intended insert functionality is maintained. For example, we want to make sure that the ovens are hot enough or that a passenger can receive coffee in a reasonable time frame.

This was translated into code by creating priority functions for each individual insert that outputs the priority based on the time delay since the insert requested power but has not yet had it granted, or the insert's deviation from its desired temperature. If an insert breaks any bounds that have been set in the priority functions, the insert receives maximum priority. Furthermore, in order to give an insert's priority meaning when compared to another type of insert, priority factors were assigned. This means that each insert's specific priority is multiplied by a constant to ensure that it scales relative to a different type of insert. For this, our assumptions were based on the insert's functionality's importance. For example, keeping the ACU cool is vastly more important than heating water. The table below outlines the priority calculations for each insert.

Table 2.1.3-1: Insert priorities

Insert	**Priority Function Explanation**	**Max Priority Setting**	**Priority Factor**
ACU	Priority is determined by the variation between current temperature and the proper temperature.	Maximum priority is triggered when the ACU breaks its upper temperature threshold	2
Oven	Priority is determined by the variation between current temperature and the proper temperature.	Maximum priority is triggered when the oven breaks its lower temperature threshold	1
Beverage Maker	Priority increases with the time since an initial ungranted power request was made.	Maximum priority is triggered after 5 minutes of delaying the beverage maker since an initial request.	0.5
Water Heater	Priority is determined by the variation between current temperature and the proper temperature.	No maximum priority included.	0.25

When it comes to estimating the added mass of this solution, we assume that we'll need a separate temperature sensor for each insert. If Honeywell T7022A temperature sensors are used

[10], and there is one for each insert, the added mass would be 7×68g, or 0.476kg. Additionally, a separate computer would be needed for the controller. Assuming an NVIDIA Jetson Xavier NX computer is used [11], this adds 0.762kg of mass. Finally, we hypothesize the need for centralized control using a tablet for the power controller in Section 6.2, and assuming an iPad mini is used [12], this adds 0.30kg of mass. Thus, the power controller solution is estimated to add 1.54kg of mass to the galley.

2.2 Aerogel

Aerogel is a silica gel in which the liquid has been removed to leave air pockets. This material has a density as low as 160g/m^3 and a thermal conductivity value as low as 0.004W/m-K, which stays the same over its lifetime. However, it is extremely rigid and fragile, and as such, efforts to decrease its fragility can increase the thermal conductivity value. The current insulation in the cart bay is a ¼”. thick high-density polyurethane foam which has a thermal conductivity of 0.022 - 0.024W/m-K but increases to 0.031W/m-K over 5-10 years.

2.2.1 Aerogel Insulation

Aerogel can be made into insulation sheets, but to create these sheets, the aerogel’s thermal conductivity value increases to 0.012W/m-K and can go upwards to 0.023W/m-K. Although the worst-case aerogel insulation would have the same thermal conductivity value as the current insulation, the lower density would also contribute to energy savings in the form of mass reduction.

2.2.2 Assumptions

In our model, we implemented the best-case scenario aerogel insulation with a thermal conductivity value of 0.012W/m-K and a thickness of ¼”. The aerogel insulation would replace the current insulation at the locations it is currently implemented.

The aerogel insulation has a density of 70kg/m^3 [13] compared to the current insulation which has a density of 320kg/m^3 [14]. With the insulation filling the same volume, the aerogel would have a mass of just under a quarter of the current insulation’s mass; the aerogel would have a mass of 0.413kg while the PU foam would have a mass of 1.89kg, so this solution results in a reduction of 1.48kg in the mass of the galley.

2.3 Variable Air Flow

The idea of this solution is to redirect air flow either from cart bays that do not actively contain carts to cart bays that do or from cart bays that do not require cooling to cart bays that do. Our hypothesis for this solution was that the cooling redirected from empty cart bays to the filled cart bays would reduce the amount of time it would take for the ACU to cool those carts due to a higher mass flow of cold air into those cart bays and a higher forced convective heat transfer coefficient from the increase in air speed.

2.3.1 Electrical Mechanism

The electrical mechanism, as opposed to the mechanical variant described later in Section 6.4, would require multiple return air sensors and motors. We did not design a functional mechanism for this solution, but it is not needed to understand the impacts of implementing this solution.

2.3.2 Assumptions

We assumed that there would be no recirculation in the ventilation system and that each cart bay would be distributed an equal amount of cold air.

Using the equation for mass flow rate: $m_{dot} = \rho V A = constant$, where ρ is the density of air (interpolated based on temperature), V is the air speed, A is the outlet area, and $V \times A$ is the volumetric flow rate (120L/s, which was given to us by the sponsors at Collins). The flow speed of cold air being distributed was calculated based on the number of carts in the bays, and then we scaled the forced convective heat transfer coefficient with this as well.

Assuming that the mechanism is built with eight HiTec Multiplex HS-311 Servos [15], one for each inlet and return air vent, the added mass would be 8 × 43g, or 0.344kg. The vents would require covers as well with a minimum surface area of 260mm × 43.5mm; this was taken from the CAD model. Assuming the material is made from Aluminum 6061 and a thickness of 1mm, the added mass from the cover would be 0.0305kg for each cover. There would need to be eight covers for a total added mass from the covers of 0.244kg. See Appendix A for density information. The total added mass for the entire mechanism would then be 0.588kg.

3. Code Overview and Explanation

The code is structured according to the following diagram:

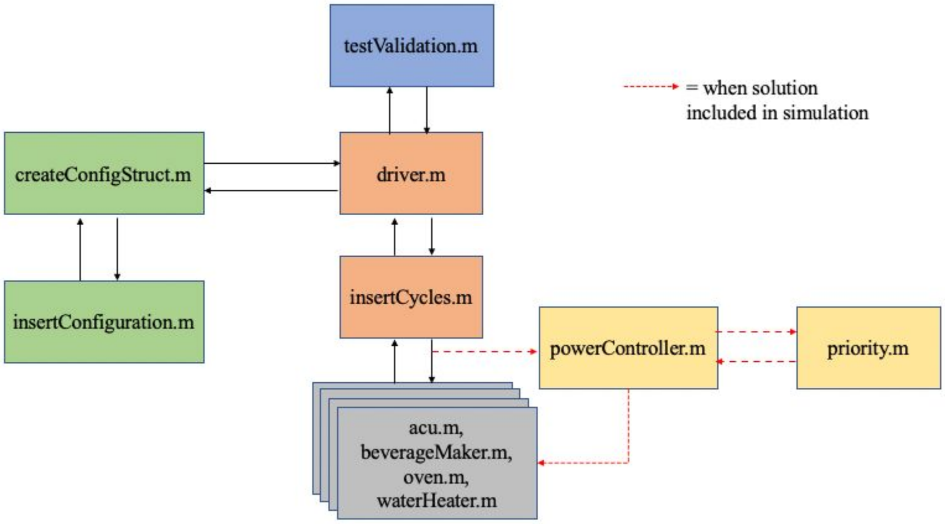

Figure 3-1: Code structure

The execution of the model is performed from the *testValidation.m* script. It defines the number of simulations that are run, which plots are generated, and the desired solution configurations included (power controller, aerogel insulations, etc.). It will then randomly generate top-level parameters with the *createConfigStruct.m* file and run the *driver.m* program on the baseline galley configuration and the selected solution for the specified number of runs. It then calculates the percentage difference in average and peak power consumption, displaying the results in the command window.

The *createConfigStruct.m* file is used to control the highest-level parameters, such as the total flight time and the number of onboard passengers (those outlined in Section 1.2.1). Using this data, it then creates a randomized configuration for each insert that defines its usage over the course of the flight, using the assumptions laid out in Sections 1.2.2-1.2.5 above. This collection of parameters defines a unique simulation configuration, and this information gets passed on to the driver.

The driver program (*driver.m*) iterates through the flight duration second-by-second calling insert functions (*acu.m*, *beverageMaker.m*, *oven.m*, and *waterHeater.m*) that model the thermal interactions and their effect on the state of an insert at that moment in time. Additionally, if the power controller solution is included in the simulation, its function (*powerController.m*) runs at the beginning of each loop, assessing the current requests for power and distributing/restricting it accordingly. This function implements power controller functionality by organizing power requests according to priority (calculated in *priority.m*) before fulfilling a subset of the requests. It will also write data to a file called *restricted.txt* when it restricts power to an insert, and to a file called *thresholdViolation.txt* if the controller's threshold is ever violated. If the simulation is set to generate plots, once the entire flight has been simulated it will display a series of graphs showing the state of various variables over time.

Finally, there is a lone file called *acuFlow.m*, where work was started on a potential future solution involving multiple ACUs. This next step is outlined in Section 6.5, and since it's not part of the current solution system it doesn't play a role in the current model.

4. Requirements and Validation Methods

Our team worked together with our project sponsor at Collins to develop a list of five critical requirements for the success of this project, and we developed validation methods for each requirement which were all approved.

4.1 Requirements

The project's five most critical requirements are laid out in Table 4.1-1 below.

Table 4.1-1: Top five critical requirements

	Subsystem Group	User Requirements, Target & Units	Origin of Validation Methodology	Team / Individual Responsible for Completion
1	Galley power consumption	Reduce average galley power consumption by 10%	Student Developed	Power and thermal subteams
2		Reduce maximum galley peak power consumption by 25%	Student Developed	Power subteam
3	Aircraft power consumption	No increase in aircraft power consumption required to maintain steady level flight due to added mass; galley power savings must offset power increase from added mass	Student Developed	Power subteam
4	FAA requirements	Compliance with simplified, relevant FAA fire regulations derived from *Title 14, Chapter 1, Subchapter C, Part §25.851 through §25.869*	Student Developed	Systems Integration
5		Compliance with simplified, FAA emergency conditions regulations derived from *Title 14, Chapter 1, Subchapter C, Part §25.581, §25.1362*	Student Developed	Systems Integration

4.2 General Validation Process

The origin of all the validation methods are student developed since there isn't a standard industry methodology to follow. These validation methods were developed after consultation with our sponsor at Collins. We have five critical requirements which had to be validated, which are detailed in the subsequent sections.

4.2.1 Decrease in Average Power Consumption Test

The requirement validated with this test is a 10% reduction in average power consumption with respect to the baseline model. In order to pass, the average power consumed by the galley solution over 100 flights must be 10% less than the average power consumed by the baseline galley over 100 flights. To validate this requirement we ran the MATLAB test script *testValidation.m*, which records the average power consumption of the baseline and solution galleys and automatically calculates the percentage change, outputting whether or not the requirement is satisfied to the user in the command window.

4.2.2 Decrease in Max Power Peak Intensity Test

The requirement validated with this test is a 25% reduction in peak power consumption with respect to the baseline mode. In order to pass, the maximum power peak in the solution galley model over 100 randomly generated flights must be 25% less than the maximum power peak in the baseline model over 100 randomly generated flights. The methodology we used to validate this requirement is nearly identical to the one laid out in Section 4.2.1 above: the *testValidation.m* script also records the power peaks of the baseline and solution galleys, calculates the percentage difference, and outputs to the user whether or not the requirement was satisfied.

4.2.3 Total Aircraft Power Consumption Test

The requirement validated with this test is that there is no increase in overall aircraft power consumption required to maintain steady level flight due to the added mass from our solutions. Essentially, this ensures that our galley power savings must offset the power increase from added mass of the solution implementation. In order to pass this test, the average decrease in galley power consumption over 100 flights must be equal to or greater than the increase in total aircraft power needed to maintain steady level flight from any added mass to the galley system. To validate this requirement we ran the *testValidation.m* script, which not only outputs the percentage difference in average power consumption between the baseline and solution galleys, but also the solution galley's actual average power consumption. This reduction in power can then be compared to the increase in aircraft power needed to maintain steady level flight, which is calculated through the following equation:

$$\Delta P = 2.6130 \times 10^{-6} \times \Delta m^2 + 11.3928 \times \Delta m$$

where Δm is the change in mass due to our solutions. The mass change assumptions are outlined in Sections 2.1.3, 2.2.2, and 2.3.2. This equation is derived from the known equation for power consumption in steady level flight: $P = \frac{1}{2}\rho V^3 S C_{D_0} + \frac{W^2}{\frac{1}{2}\rho V S}\left(\frac{1}{\pi e AR}\right)$, and since weight ($W$) is the only variable that changes, the first term drops out when calculating ΔP. The assumptions that were made in simplifying this to the equation above are the following: 80% efficiency in the engine's generators when converting to electrical power, $\rho = 0.4135\text{kg/m}^3$ (air density at cruising altitude of 10km), $e = 0.7$ (rectangular wing assumption) [16], $AR = 9.49$ [2], $S = 442\text{m}^2$ [3], $V = 250.8\text{m/s}$ [2], and the baseline weight $W = 490{,}000\text{lbs}$. This weight comes from an operating empty weight of 314,000lbs [2], a half tank of fuel weighing 121,831lbs (since the maximum flight time is 16.6 hours and we are modeling an 8-hour flight) [2], and 325 passengers with an average weight of 137lbs and 30lbs of baggage [17]. It sums to slightly larger than 490,000lbs, which was rounded down since these assumptions are rough, and this weight corresponds to about 80% of the maximum takeoff weight [2].

4.2.4 Compliance with FAA Fire Regulations Inspection

This validation method is an inspection to ensure compliance with relevant (but simplified) FAA fire regulations derived from FAA Regulations Title 14, Chapter 1, Subchapter C, Part §25 [18]:

Table 4.2.4-1: Fire protection regulations

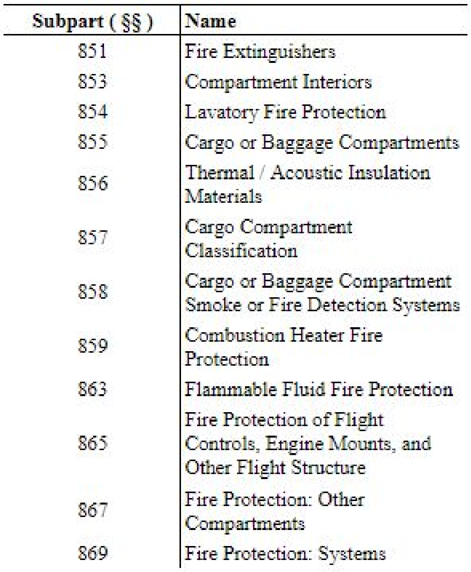

Subpart (§§)	Name
851	Fire Extinguishers
853	Compartment Interiors
854	Lavatory Fire Protection
855	Cargo or Baggage Compartments
856	Thermal / Acoustic Insulation Materials
857	Cargo Compartment Classification
858	Cargo or Baggage Compartment Smoke or Fire Detection Systems
859	Combustion Heater Fire Protection
863	Flammable Fluid Fire Protection
865	Fire Protection of Flight Controls, Engine Mounts, and Other Flight Structure
867	Fire Protection: Other Compartments
869	Fire Protection: Systems

These regulations were simplified into the following checklist:

1. Ensure no material is highly flammable in any galley modification
2. Ensure no material emits toxic fumes in the presence of heat in any galley modification

4.2.5 Compliance with FAA Emergency Conditions Regulations Inspection

This validation method is an inspection to ensure compliance with relevant (but simplified) FAA fire regulations derived from FAA Regulations Title 14, Chapter 1, Subchapter C, Part §25 [18]:

Table 4.2.5-1: Emergency conditions regulations

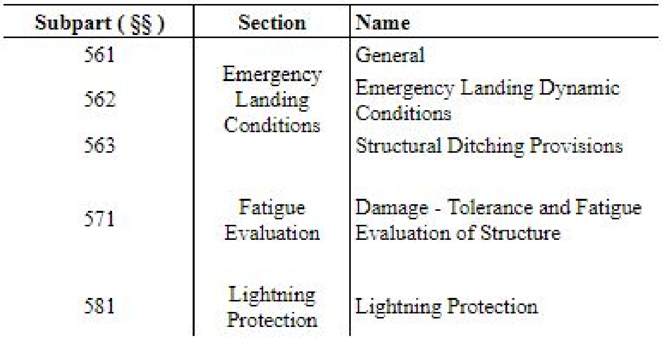

Subpart (§§)	Section	Name
561	Emergency Landing Conditions	General
562		Emergency Landing Dynamic Conditions
563		Structural Ditching Provisions
571	Fatigue Evaluation	Damage - Tolerance and Fatigue Evaluation of Structure
581	Lightning Protection	Lightning Protection

These regulations were simplified into the following checklist:

1. Ensure that no physical galley modification will alter the aircraft's protection against catastrophic effects from lightning or endanger the aircraft
2. Ensure that no physical galley modifications is likely to break loose
3. Ensure that no physical galley modification is likely to injure a passenger in the case that it does break loose
4. Ensure that no physical galley modification is likely to nullify any of the escape facilities provided for use after an emergency landing.

5. Results

The solutions we employ ensure that we realize our top five critical requirements. The solutions work together to reduce the average power consumption by 12.8%, hence meeting the 10% goal laid out in the requirement. In addition, the solutions reduced the maximum power consumption by 27.4%, satisfying the requirement of 25%. Based on the assumptions outlined in Sections 2.1.3, 2.2.2, and 2.3.2, the solutions combined increase the galley mass by 0.918kg, which constitutes a power increase of 10.5W according to the equation outlined in Section 4.2.3. Since the average power reduction is approximately 130.4W, it easily offsets this power increase, satisfying our third critical requirement.

Both sets of FAA requirements outlined in Sections 4.2.4 and 4.2.5 are confirmed as "Compliant" by the team. No materials that are included in any of the three solutions are flammable, nor do they emit toxic fumes in the presence of heat, which satisfies the FAA fire protection regulations. In addition, our modifications to the galley are all fairly small and lightweight, and can be fastened to the galley in the same manner that existing components are fastened, so they are not likely to break loose, or injure anyone even if they do. Since they are contained in the galley monument, they are not likely to nullify any escape facilities, and also do not alter the aircraft's protection against catastrophic effects from lightning. Therefore, the solutions satisfy the FAA emergency conditions regulations as well. The satisfaction of the five most critical requirements are outlined in the table below.

Table 5-1: Top five critical requirements and results

User Requirements, Target & Units	Achieved Result
Reduce average galley power consumption by 10%	12.8%
Reduce maximum galley peak power consumption by 25%	27.4%
No increase in aircraft power consumption required to maintain steady level flight due to added mass; galley power savings must offset power increase from added mass	-119.9W
Compliance with simplified, relevant FAA fire regulations derived from Title 14, Chapter 1, Subchapter C, Part §25.851 through §25.869	Compliant
Compliance with simplified, FAA emergency conditions regulations derived from Title 14, Chapter 1, Subchapter C, Part §25.581, §25.1362	Compliant

The following sections discuss the breakdown of how each solution contributes to the results.

5.1 Power Controller

The power controller on its own shows very promising results: over 100 simulations, it is responsible for a 7.54% decrease in average power consumption and a 27.4% decrease in peak power consumption. When analyzing the plots generated by a random simulation, we can observe some of the controller's behavior over the course of a flight. The first set of plots to analyze is the total power consumption over the flight between the baseline and power controller configurations.

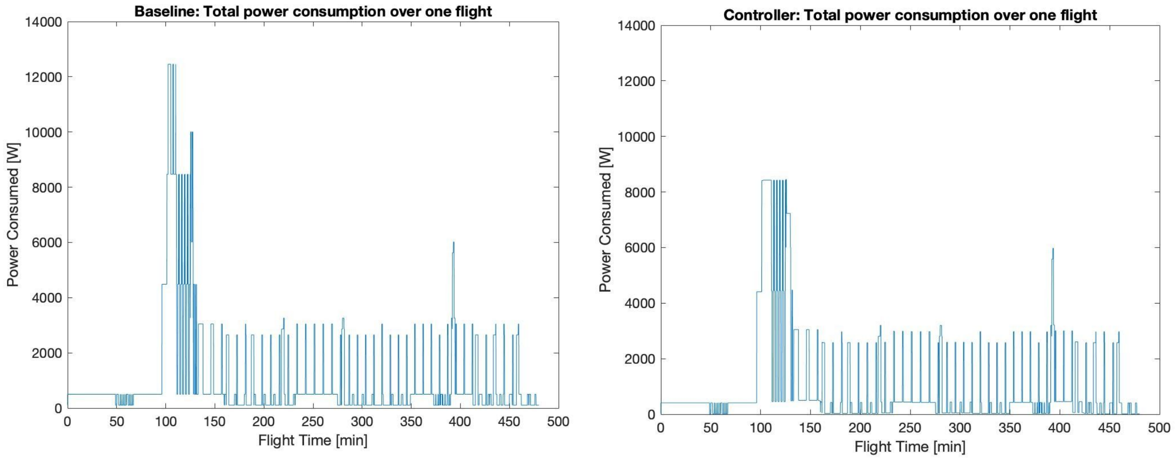

Figures 5.1-1 and 5.1-2: PC of a galley over one flight
(Left: Baseline, Right: Power controller)

We can see by comparing Figure 5.1-1 to Figure 5.1-2 that the power controller is effective with reducing the peak power consumption. The controller successfully reduces a peak of over 12kW to less than 9.5kW. This is because in its current configuration, the controller is set to allow a maximum consumption of 9.5kW, and it's unlikely to ever violate that as it would require multiple inserts to have maximum priority. As can be seen from the figures, the controller successfully held the power consumption below that level, without drastically affecting the power consumption over the rest of the flight. In fact, with the clustering of insert usage around the first meal period (in the 90-150 minute region), the controller does not have to restrict power too often during the flight, but demonstrates that it's capable of doing so when needed. Since the power controller is solely responsible for managing peak power loads, it alone achieves the 27.4% reduction in peak power consumption needed to meet the critical requirement. Additionally, it's easier to see the power controller shutting off inserts that have been in standby mode for more than thirty minutes by zooming further into the figures above.

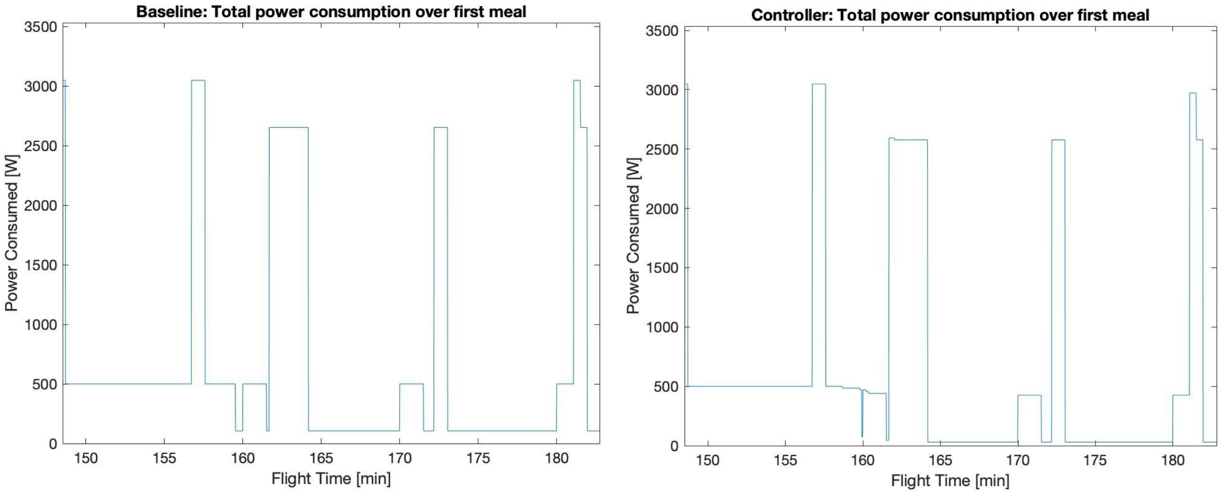

Figures 5.1-3 and 5.1-4: PC of a galley over the first meal period
(Left: Baseline, Right: Power Controller)

This shutoff, which is apparent in the lower power consumption seen in Figure 5.1-4 when compared to Figure 5.1-3, combined with the insert shutdown (excluding the ACU) during the first and last 45 minutes of the flight (takeoff and landing) is what results in the 7.54% decrease in average power consumption when the power controller is used on its own. Finally, it's important to ensure that power controller restrictions do not have too drastic of an effect on insert functionality so that flight attendants do not have to wait for extended periods of time during food/beverage preparation. As was mentioned earlier, the power controller restrictions are clustered around the first meal period, in particular because of the oven usage during that time, so they are the most important inserts to analyze. The plots below show the temperatures of the three ovens over their use during the first meal period.

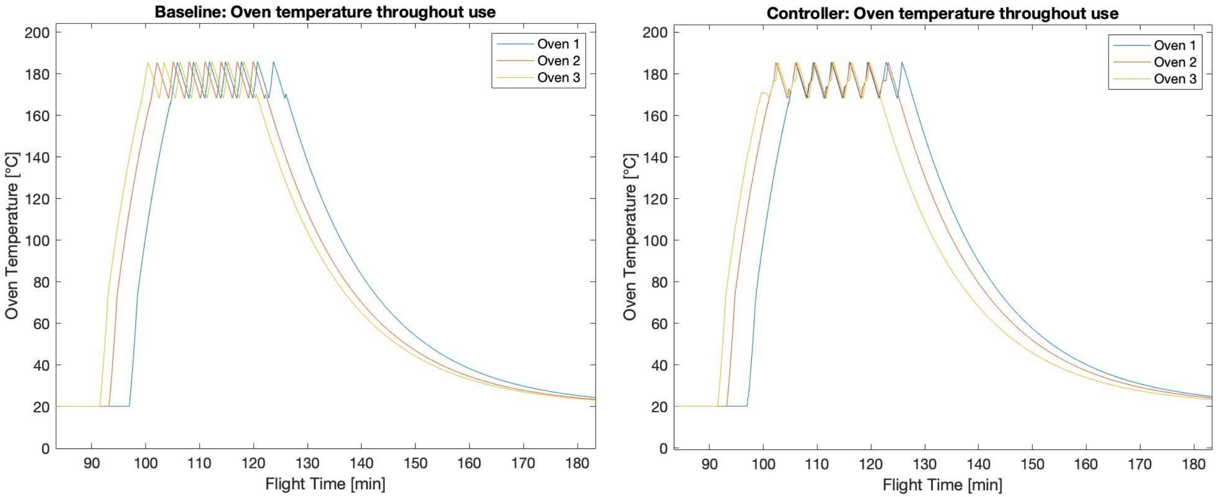

Figures 5.1-5 and 5.1-6: Temperature of ovens over the first meal period
(Left: Baseline, Right: Power Controller)

As can be seen in Figures 5.1-5 and 5.1-6 above, the ovens finish cooking the meals less than five minutes later when the power controller is used as opposed to the baseline galley, which in our opinion is an acceptable delay for the decrease in peak power consumption. This behavior is common across many simulations: delays do not often exceed five minutes, which is a positive outcome and reinforces the notion that a power controller is a feasible solution for galley power management problems. We can see in Figure 5.1-7 below that this is achieved without the ovens entering maximum priority (a value of -1) very often.

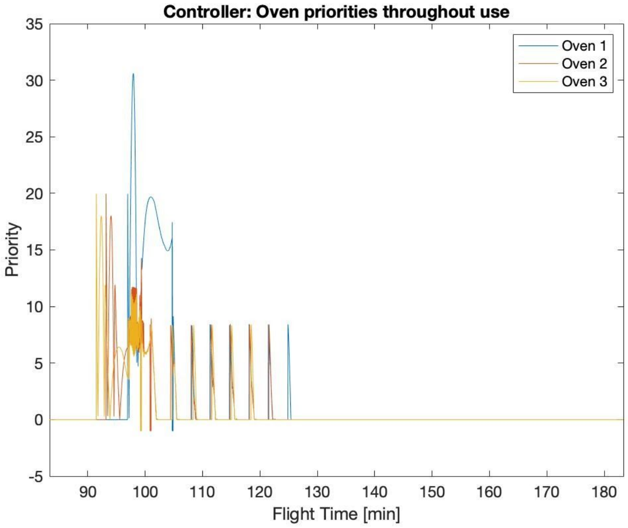

Figure 5.1-7: Oven priorities over the first meal period

The temperature curves for other inserts (the water heater, beverage makers, and ACU) show similar properties, but it's more interesting to look at the power consumption graphs to observe another common phenomenon with the controller.

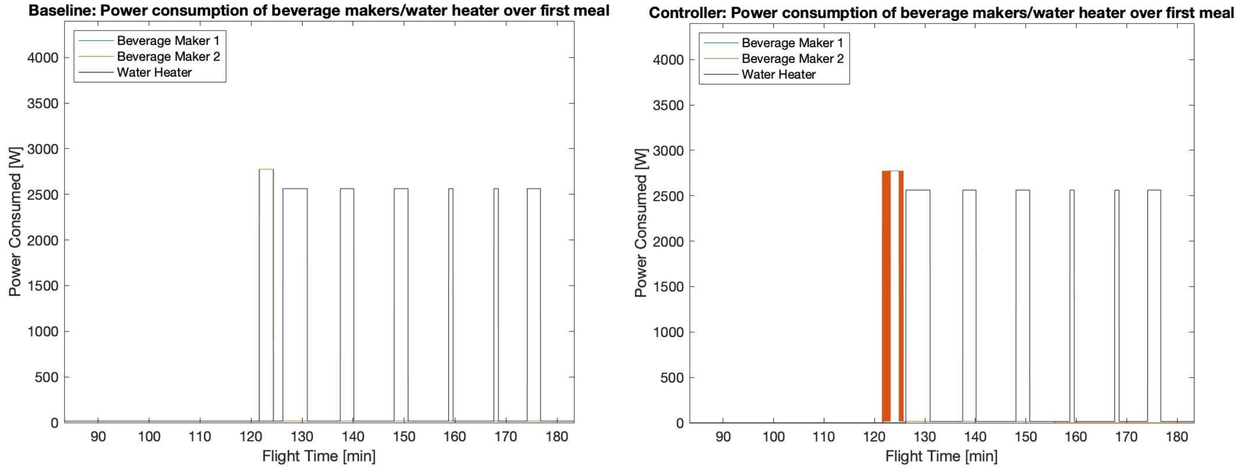

Figures 5.1-8 and 5.1-9: PC of beverage makers and water heater over first meal
(Left: Baseline, Right: Power Controller)

In Figure 5.1-9, the "block" of red is the beverage maker turning on and off approximately every second due to rapid fluctuations in its priority relative to other inserts (see Figure 5.1-10 below). The power controller is not limited in how quickly it can cycle an insert on and off: our sponsors at Collins confirmed that the inserts can handle second-to-second fluctuations in power with minimal degradation, but altering the controller so that cycling happens over a longer period of time may be an interesting point of exploration, and will be discussed further in Section 6.2 later. While the example in Figure 5.1-9 shows rapid cycling for a beverage maker, this same cycling has been observed in the ovens during certain simulations as well.

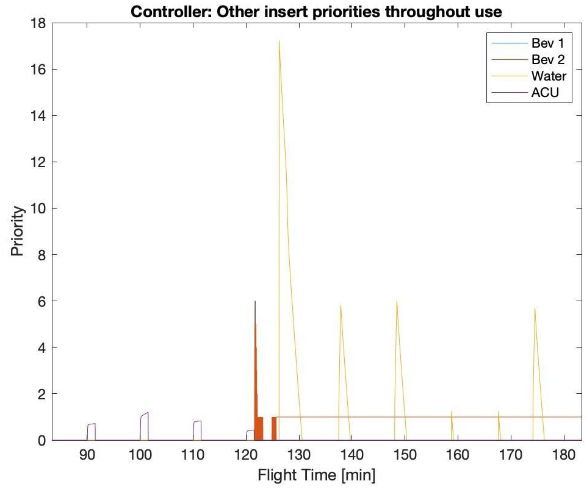

Figure 5.1-10: Priorities of beverage makers, water heater, and ACU over first meal

5.2 Aerogel

Aerogel insulation, when implemented as the sole solution, reduces the average power consumption by 5.19%. This reduction in power consumption can be attributed to creating a better thermal boundary between the cart bays and the ovens. When comparing Figure 5.2-1 to Figure 5.2-2, and Figure 5.2-3 to 5.2-4, the largest difference occurs in the middle of flight. The ACU in the baseline model must undergo multiple cycles to maintain the temperature in the cart bay; the solution model is more effective at keeping heat out of the cart bay and as a consequence, the ACU is not drawing power as often.

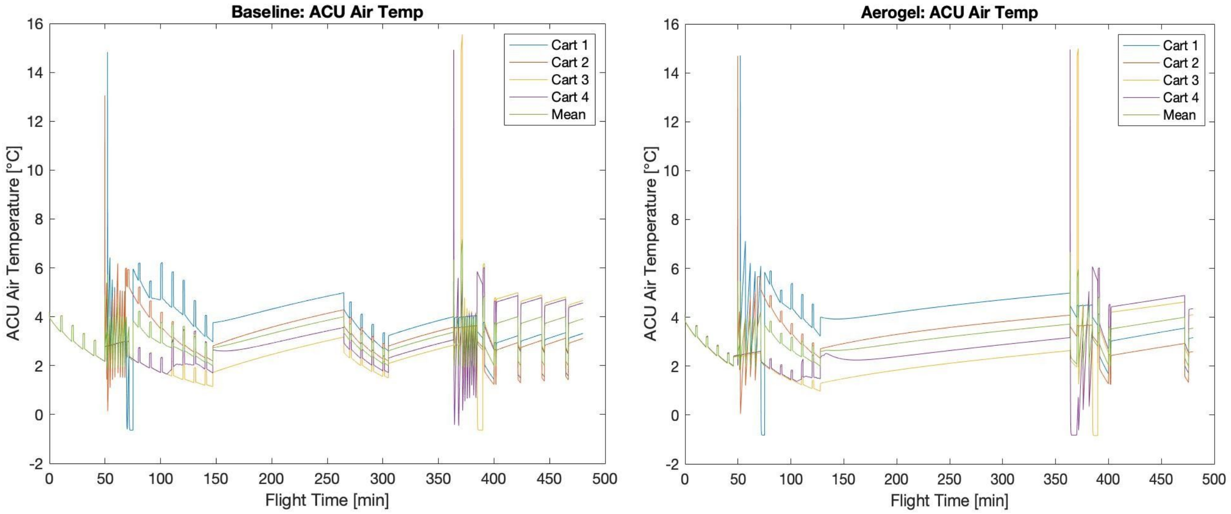

Figures 5.2-1 and 5.2-2: ACU return air temperature over flight
(Left: Baseline, Right: Aerogel)

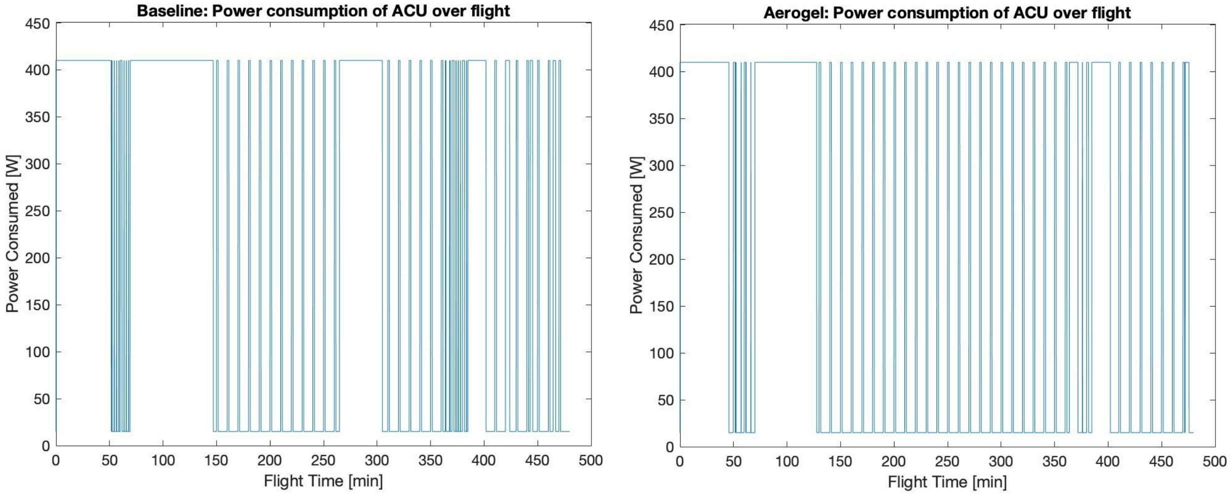

Figures 5.2-3 and 5.2-4: ACU PC over flight
(Left: Baseline, Right: Aerogel)

5.3 Variable Air Flow

Our final version of the variable air flow solution decreased the average power consumption by 0.40% when implemented on its own. When comparing Figure 5.3-1 and Figure 5.3-2 below, the solution allows for more temperature fluctuation in the empty cart bays. This can be seen especially after Cart 2 is taken out of the cart bay at around 50 minutes where the temperature spikes higher in Figure 5.3-2 than Figure 5.3-1. The spike is due to cold air being redirected from the empty cart bays to the filled cart bays. However, upon analysis of the graphs during this time, the occupied cart bays are only cooled to a temperature only 0.1°C lower than the baseline at the same time.

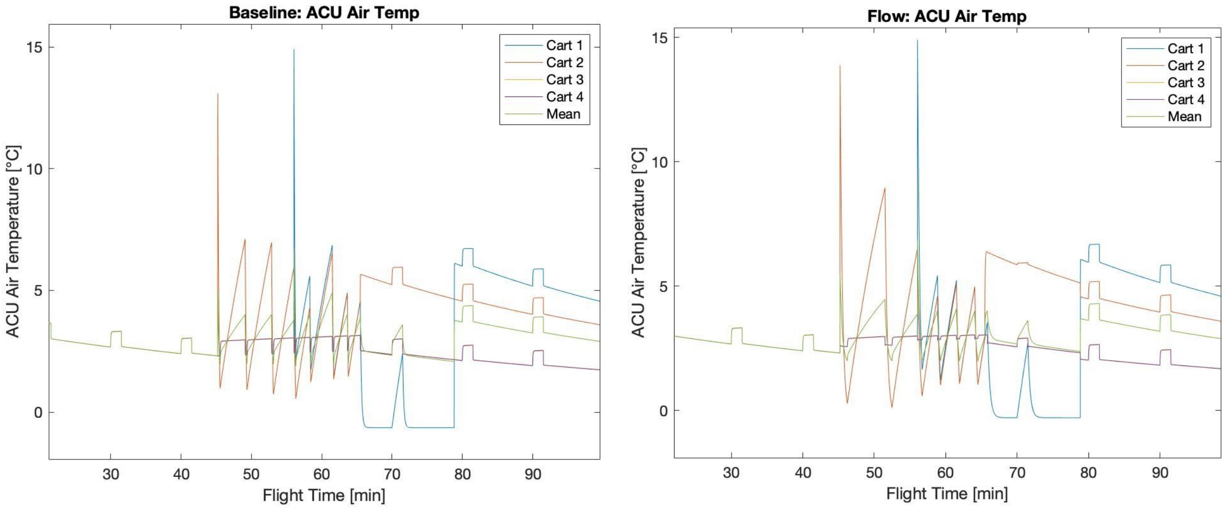
Baseline: ACU Air Temp

Flow: ACU Air Temp

Figure 5.3-1 and 5.3-2: Return air temperature during first cart service
(Left: Baseline, Right: Variable air flow)

In general, results for the variable air flow solution varied from -0.5% to +0.5% in average power consumption depending on the temperature threshold used for empty carts or the amount of air leakage allowed into the cart bays through “closed” vents. This is likely because the amount of power saved by the increased cooling to the filled cart bays is reduced by the amount of power required to cool the warm air once carts are returned from service (since unoccupied bays are cooled less). The expected result was that the occupied cart bays would cool much quicker with the variable airflow and the empty cart bays would not experience a noticeable impact in cooling, resulting in large power savings. This, however, did not occur in practice, and the ACU was turned on for a similar amount of time regardless of air distribution method. As a result, the ACU power consumption was very similar between the baseline and variable air flow model, which is outlined in Figures 5.3-3 and 5.3-4 below.

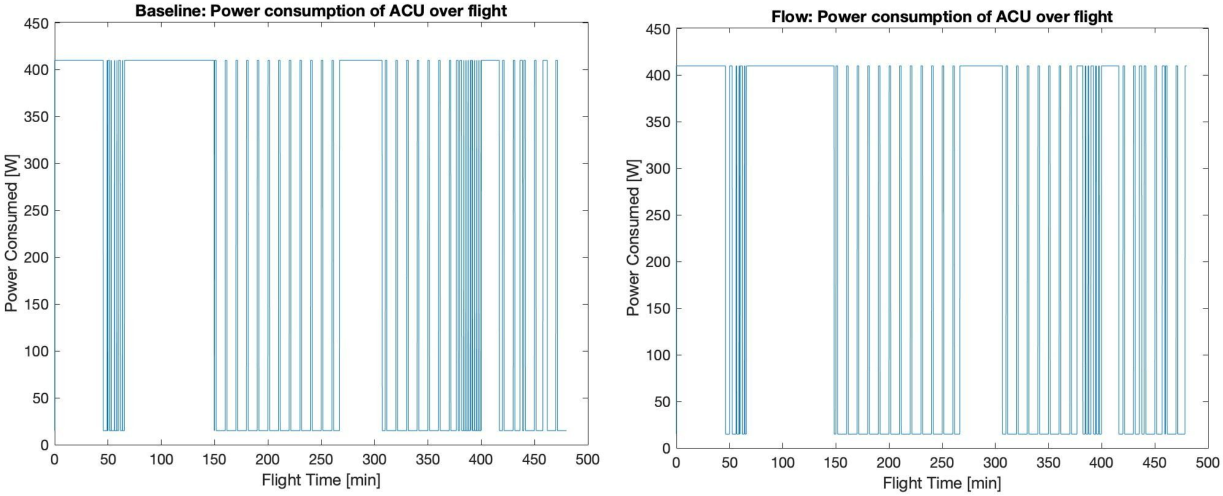

Figure 5.3-3 and 5.3-4: ACU PC over flight
(Left: Baseline, Right: Variable air flow)

6. Implications and Next Steps

This section will highlight the implications of and next steps for our baseline model, our solutions, and all of their assumptions, including areas where errors or false assumptions could have been made.

6.1 Baseline Model

In continuation of this project, a notable improvement could be to implement a Graphical User Interface (GUI) or User Interface (UI) for ease of use when creating galleys. In addition, by using the Partial Differential Equation (PDE) toolbox in MATLAB, a finite element analysis (FEA) can be performed to either make the model more robust by calling the FEA each iteration or by finding general heat transfer equations to optimize the model for speed.

The assumptions we made were based on the best information we had at the time. In hindsight, a few of the assumptions or ways the baseline is modeled would have to change. The following list details the top three changes (in no particular order) to be implemented to better model the galley.

1. The assumptions made in Section 1.2.1 regarding the number of passengers doesn't create a lot of variation in the capacity of the flight. We believe that including a higher standard deviation in the probability distribution and bounding the output to a maximum of 350 passengers would better model the varying flight conditions.
2. The assumptions made in Sections 1.2.3 and Sections 1.2.4 on the percentage of passengers ordering coffee/tea may differ significantly from airline data. These assumptions were built on statistics about the normal day-to-day consumptions of these beverages, which may differ from the behavior of passengers on an airplane.

3. We made a few assumptions with respect to modeling of the beverage maker which may not accurately represent the real world scenario, mostly from a lack of information. Without information about the heating mechanism, we hardcoded the temperature increase every second to reflect the brew cycle time of 2 minutes 45 seconds. Another piece of information which was lacking was whether there is a tank of water that maintains hot water or if the beverage maker heats up water as it dispenses a beverage. Altering the beverage maker model to make it more realistic would be beneficial to the model, and may impact results.

6.2 Power Controller

Based on the power controller results in Section 5.1, the power controller appears to be a very promising solution for the galley, and we definitely recommend iterating on its design and further exploring its capabilities. However, there are a few things that we believe should be addressed regarding this solution in the "next steps" for the project. Firstly, the oven priority uses a third-order polynomial to model the temperature curve when an oven is preheating, and uses this curve to estimate how much time remains for the oven to heat up. This method allowed the team to successfully prevent all three ovens from being turned on while allowing a maximum preheat delay of 4 minutes. This polynomial curve was fitted to data collected from our model, and hence is subject to inaccuracies from issues with the way the ovens are modeled. If the oven model is altered, or better yet, if data is recorded from a real insert, this function (in *priority.m*) should be updated. Based on the success of this method, it is our recommendation that the company explore this method for other inserts, but a refinement of the function is likely necessary.

Another aspect of the controller (which was discussed briefly in the power controller results), is the fact that it can fluctuate an insert's power every second. While our sponsors at Collins have confirmed that this does not noticeably degrade the inserts, we still think it would be worth exploring putting a limit on how quickly it can be cycled. For example, if an insert is given power at some point, it should continue to be granted power for 30 seconds before it can be shut off again. By observing Figure 5.1-6, we can see that inserts like the ovens can fall into sync with one another. This is because the controller is giving power to one or two of the ovens at each second, and changes which ones are granted power as well. Since the heating element doesn't cool off in one second, the ovens can still increase in temperature together whilst only one or two are actually consuming power. This may be okay, but we believe that limiting how quickly power can be cycled may cause the power controller to result in more natural cycling, where the ovens being granted power are more visibly staggered, with one cooling down whilst the two others heat up.

Moreover, using the power controller to turn inserts that would normally be in standby mode (15W) to off (0W) is an effective method to reduce galley power consumption. In fact, all of the power controller’s contributions to energy savings (7.54%) come from this. While this method is promising, it has important limitations that must be considered. When an insert is switched from standby to off, its customer-facing functionality is disabled as it no longer has the ability to forward a power request to the power controller. In practical terms, this means the insert will be unable to respond to user input. In order to avoid this, we recommend incorporating a unified smart screen and control base to the galley, similar to an iPad or tablet. This interface would enable flight attendants to have direct control over the power controller’s decisions as an override capability could be incorporated. Furthermore, by replacing each insert’s individual buttons and dials with buttons on a single unified panel, the power controller could likely switch inserts into off mode at all times that they are inactive as there would be no need for standby mode. The potential energy conservation from this would be substantial as standby mode power consumption even with the current standby shutoff still represents approximately 3% of total galley power consumption.

While the power controller is effective at reducing power peaks by 25%, if the company is interested in reducing peaks by more than this, then considerable effort should be invested in developing more robust insert priority functions. For this project, the relative priority given to different types of inserts (e.g. oven vs. beverage maker) was somewhat arbitrary. This was okay as the key to reducing the peaks by 25% was to ensure that not all three ovens were on simultaneously (which consumes 12,000W), and thus it was most important that insert priority was determined effectively for inserts of the same type. Furthermore, in order to achieve a 25% decrease in peak power consumption, many minor inserts like the water heater and beverage maker are hardly delayed. This would not be the case if the company was interested in further diminishing power peaks and it’s likely that the power controller in its current state would begin to drastically degrade the functionality of certain inserts whose priorities relative to other inserts have not been adequately determined.

6.3 Aerogel

The implementation of aerogel insulation sheets decreased the average power consumption by a significant amount, over 5%. The heat transfer from the ovens to the cart bays is a significant inefficiency due to thermal interactions between the inserts. With regards to the scope of our project, the tolerances for the oven and carts did not allow for implementation of aerogel synthesis in other areas. Further research outside the scope could be conducted to allocate space for aerogel insulation in the honeycomb paneling around the oven or inside of the oven.

In addition, the best case scenario was used in terms of the aerogel insulation’s thermal conductivity value; depending on the aerogel insulation used, this decrease in power consumption can change drastically.

Alternatively, another option to consider is to keep the aerogel in the pure form and support the material using a vibration-dampening material. The thermal conductivity value would increase based on the purity of the aerogel, the environmental conditions, and the composite material.

Another alternative is an experimental synthesis of aerogel. This mixes aerogel with a foam to decrease the fragility of the material while keeping a lower thermal conductivity value. If the insulating material were to break during its use, the foam can be reapplied rather than replaced to save costs. However, as this option is experimental, it is not currently available for commercial use. See Ref. [19] for more information.

6.4 Variable Air Flow

The variable airflow solution produced very little difference in power consumption. Furthermore, had we not accounted for air leakage in the model or had not set an upper temperature threshold (to prevent unoccupied cart bays from completely warming to cabin temperature), then the differences between the baseline and variable air flow solution would have been entirely nonexistent.

While energy consumption was not reduced, this solution does however lead to a more efficient cooling system where carts that remained inside their respective cart bays during service were cooled lower than their counterparts. This suggests that variable airflow could prove promising and that an error in our solution was poor distribution of air since we continued using the mean air temperature from a single intake to dictate airflow when the ACU turns on. Had four air intakes been used, or more simply, had we included a temperature sensor for each cart bay, the variable air flow may have been dramatically more effective in reducing the power consumption of the ACU. The ACU would have stopped cooling once an occupied cart bay cooled to the correct temperature rather than continuing to cool because the mean temperature was too high from warm unoccupied bays. As such, we recommend that Collins investigate variable airflow with the addition of temperature sensors for each individual cart bay.

The implementation suggested in Section 2.3 was to use servos to control the vents. An alternative to this for variable airflow would be to implement a mechanical mechanism using a spring loaded plate that would depress when carts are returned in order to open the inlet and return air vents.

6.5 Multiple ACUs

The use of multiple ACUs was a solution that only began to be considered too late into the project, and thus couldn't be carried out to completion. This solution incorporates a total of four ACUs, one for each cart bay, rather than one ACU distributing cold air to all four cart bays. This solution also separates the cart bay that holds two carts into two separate cart bays. From the ACU spec sheets, we know that the coolant system is separate from the ACU and the ACUs can be daisy-chained. An implementation was started for this potential solution in the *acuFlow.m* file, but there wasn't sufficient time to explore it further.

The current implementation began to conflict with how the baseline ACU is currently modeled, and would have taken too long to fix given the amount of time left in the project. However, initial results looked good despite the bugs and limited functionality. In addition to significant power savings, the cart contents are cooled more efficiently. Because each cart bay has its own temperature sensor, the cart bays will not individually violate the temperature threshold; with the baseline model, individual cart bays could violate the threshold as long as the mean return air temperature did not. The next steps would include finding a way to model multiple ACUs in the code rather than one. Rather than hard coding four ACUs, another option would be to give options for the number of ACUs. The difficulty here is modeling multiple ACUs while also maintaining full functionality of the baseline model. However, if only the impact of implementing multiple ACUs is desired, the process should be easy. Our team highly recommends a further study into this solution.

6.5.1 Assumptions

We made a few assumptions that are not entirely correct, but good enough for an idea of the solution's impact; these assumptions were approved by our sponsor at Collins. However, this solution could not be fully implemented by the end of the project as it was started later than the other solutions. We assumed that the amount of mass flow to each cart bay through the inlet air vents remains the same as before whether this is obtained by the same size fan with a lower fan speed or a smaller fan with a higher fan speed. The amount of power drawn by each ACU is exactly a fourth of the power drawn by the baseline ACU, and the CoP remains the same as the baseline ACU. The standby for each individual ACU is 15W, which would quadruple the total amount of standby power drawn if all four ACUs were off.

6.5.2 Preliminary Results

The model for multiple ACUs is not fully functional as other files would have to have been changed significantly; because of this, no graphs are shown in this section. Initial results, although skewed due to bugs and limited functionality, show a ~5% decrease in average power consumption. Additionally, with the assumption that all ACU components are scaled to a quarter of their size, there is no additional mass, and as such the total aircraft power consumption is reduced by the reduction in average power consumption.

7. Conclusion

The results of this project are promising and with more time and expertise, this project has the potential to be an invaluable asset in designing and improving future galleys. As this project has demonstrated, a modular galley model can easily compare multiple galley designs in a quick and efficient manner, as well as provide useful insight into the potential impact of a solution without the cost of physically implementing it.

Appendix A: Material Properties

Below, we list the material properties that are used in our model. All properties were gathered from Engineering Toolbox [20].

Table A-1: Thermal conductivity values

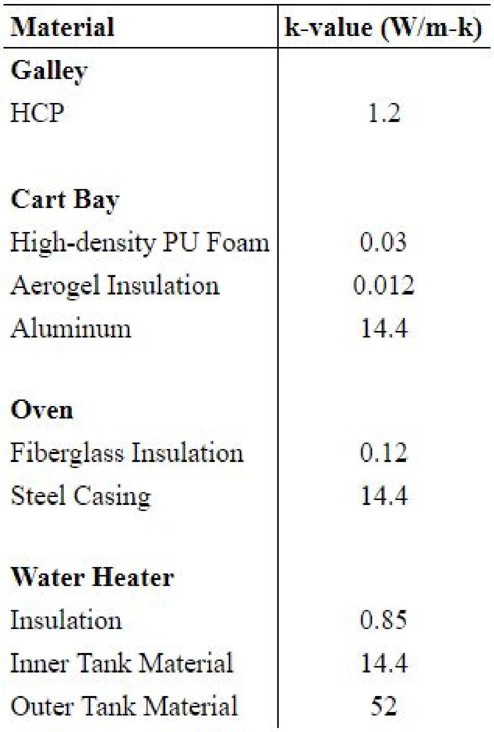

Material	k-value (W/m-k)
Galley	
HCP	1.2
Cart Bay	
High-density PU Foam	0.03
Aerogel Insulation	0.012
Aluminum	14.4
Oven	
Fiberglass Insulation	0.12
Steel Casing	14.4
Water Heater	
Insulation	0.85
Inner Tank Material	14.4
Outer Tank Material	52

Table A-2: Heat capacity values

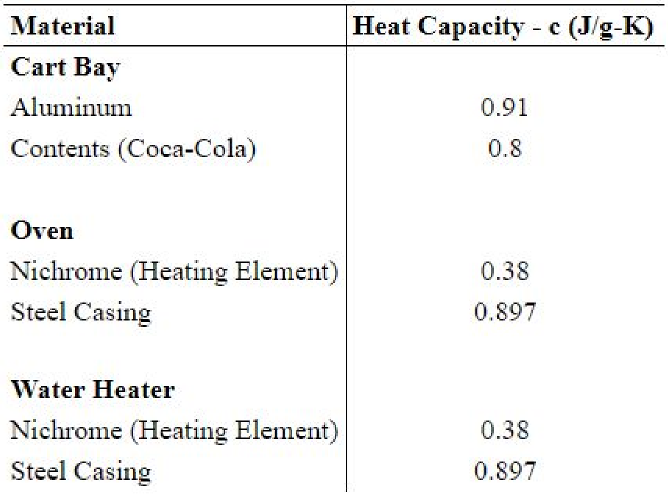

Material	Heat Capacity - c (J/g-K)
Cart Bay	
Aluminum	0.91
Contents (Coca-Cola)	0.8
Oven	
Nichrome (Heating Element)	0.38
Steel Casing	0.897
Water Heater	
Nichrome (Heating Element)	0.38
Steel Casing	0.897

*Heat Capacity of Air, Water, and Steam are interpolated in the model.

**These are the only materials that we track temperatures for.

Table A-3: Density values

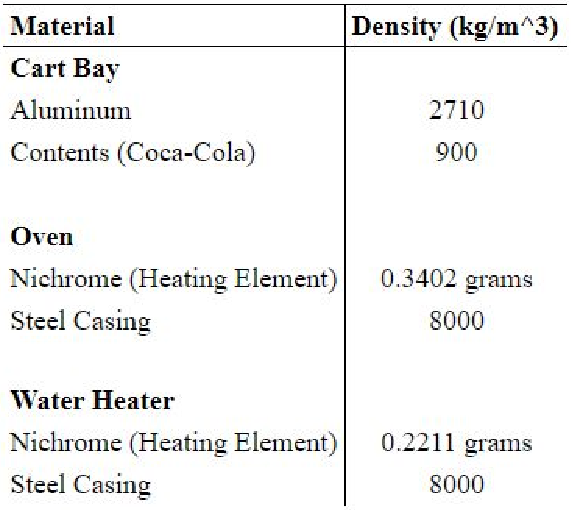

Material	Density (kg/m^3)
Cart Bay	
Aluminum	2710
Contents (Coca-Cola)	900
Oven	
Nichrome (Heating Element)	0.3402 grams
Steel Casing	8000
Water Heater	
Nichrome (Heating Element)	0.2211 grams
Steel Casing	8000

*Density of Air, Water, and Steam are interpolated in the model.

**These are the only materials that we track temperatures for.

***Heating Elements are given in grams as information on them were given in surface area and mass.

Appendix B: References

1. Alton, Robin. "How Many Hours Are In An Average International Flight?". Quora, 2019, https://www.quora.com/How-many-hours-are-in-an-average-international-flight.
2. "Airbus A350 XWB". Wikipedia, 2020, https://en.wikipedia.org/wiki/Airbus_A350_XWB.
3. "A350-900". Airbus, 2020, https://www.airbus.com/aircraft/passenger-aircraft/a350xwb-family/a350-900.html.
4. "Food & Drink". LEVEL, 2019, https://www.flylevel.com/en/why-level/food-drink.
5. "Oven Liners: Porcelain Vs. Stainless Steel". Katom, 2020, https://www.katom.com/cat/commercial-oven/oven-liners-porcelain-vs-stainless-steel.html.
6. "Preheat In A Convection Steam Oven". Sub-Zero, Wolf, And Cove, 2020, https://www.subzero-wolf.com/assistance/answers/preheat-in-a-convection-steam-oven.
7. Kunst, Alexander. "Americans' Daily Coffee Consumption At Home During The Week 2019". Statista, 2020, https://www.statista.com/statistics/250230/americans-daily-coffee-consumption/.
8. "Tea Fact Sheet - 2019-2020". Tea Association Of The USA, 2020, http://www.teausa.com/14655/tea-fact-sheet.
9. Maloney, Thomas et al. "Fire Behavior Of E-Tablets Stored In Aircraft Galley Carts". Federal Aviation Administration, 2015, https://www.fire.tc.faa.gov/pdf/TC-TN-14-40.pdf
10. “Honeywell T7022A Return Air Temperature Sensor Specifications.” https://www.amazon.com/Honeywell-T7022A1010-T7022A-Return-Temperature/dp/B008HOZV7I
11. “NVIDIA Jetson Xavier NX for Embedded & Edge Systems.”, *NVIDIA,* www.nvidia.com/en-us/autonomous-machines/embedded-systems/jetson-xavier-nx/
12. “IPad Mini - Technical Specifications.” *Apple*, www.apple.com/ipad-mini/specs/
13. “Thermal Wrap 6-mm Blanket Data Sheet.” http://www.buyaerogel.com/product/thermal-wrap/
14. “LAST-A-FOAM FR-3800 FST High Density Polyurethane Foam Sheet Specifications.” https://www.generalplastics.com/products/fst-3800

15. “HiTec Multiplex - HS-311 Standard Economy Servo Specifications.” https://hitecrcd.com/products/servos/sport-servos/analog-sport-servos/hs-311/product
16. “Induced Drag Coefficient.” *NASA*, NASA, www.grc.nasa.gov/www/k-12/airplane/induced.html
17. “Human Body Weight.” Wikipedia, Wikimedia Foundation, 14 Dec. 2020, https://en.wikipedia.org/wiki/Human_body_weight
18. “Airworthiness Standards: Transport Category Planes.” https://www.law.cornell.edu/cfr/text/14/part-25
19. Nazaren, N. and Moghaddas, J., “Synthesis and Characterization of Silica Aerogel Reinforced Rigid Polyurethane Foam for Thermal Insulation Application,” *Journal of Non-Crystalline Solids,* Vol. 461, April 2017, pp. 1-11. https://doi-org.proxy.lib.umich.edu/10.1016/j.jnoncrysol.2017.01.037
20. “Tools and Basic Information for Design, Engineering, and Construction of Technical Applications.” https://www.engineeringtoolbox.com/