MODEL-BASED DESIGN AND SPECIFICATION OF A HYBRID ELECTRIC CHEVROLET CAMARO FOR THE ECOCAR 3 COMPETITION

A Thesis
Presented to
The Academic Faculty

By

Jonathan Douglas Cox

In Partial Fulfillment
Of the Requirements for the Degree
Master of Science in Mechanical Engineering

Georgia Institute of Technology

May 2016

MODEL-BASED DESIGN AND SPECIFICATION OF A HYBRID ELECTRIC CHEVROLET
CAMARO FOR THE ECOCAR 3 COMPETITION

Approved by:

Dr. Michael Leamy, Advisor
School of Mechanical Engineering
Georgia Institute of Technology

Dr. David Taylor
School of Electrical Engineering
Georgia Institute of Technology

Dr. Aldo Ferri
School of Mechanical Engineering
Georgia Institute of Technology

Data Approved: April 2016

Dedicated to my wonderful wife Kelsay. You've sacrificed as much as I have for this, and I am truly grateful.

ACKNOWLEDGEMENTS

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF SYMBOLS

A-hr   Ampere-hours

AWD   All Wheel Drive

BCM   Battery Control Module

CAD   Computer-Aided Design

CAN   Controller Area Network

CD   Change Depleting

CS   Charge Sustaining

CSM   Current Sensing Module

EDM   Electronic Distribution Module

E&EC   Emissions and Energy Consumption

EM   Electric Machine

EMF   Electromotive Force

EPA   Environmental Protection Agency

ESS   Energy Storage System

EV   Electric Vehicle

FEA   Finite Element Analysis

FWD   Front Wheel Drive

HEV	Hybrid Electric Vehicle
HIL	Hardware in the Loop
HWFET	Highway Fuel Economy Test
I	Current
ICE	Internal Combustion Engine
IVM	Initial Vehicle Movement
MIL	Model in the Loop
MPG$_{ge}$	Miles per gallon gasoline equivalent
MTC	Motor and Transfer Case
P	Power
P1	Parallel architecture with motor on engine accessory side
P2	Parallel architecture with motor between engine and transmission
P3	Parallel architecture with post-transmission motor
P4	Parallel architecture with motor on different axle than that driven by engine
PHEV	Plugin Hybrid Electric Vehicle
R	Electrical resistance
RBE2	Rigid Body Entity
RWD	Rear Wheel Drive
SIL	Software in the Loop

SMS	Systems Modeling and Simulation
SOC	State of Charge
TCM	Transmission Control Module
UF	Utility Factor
V	Volts
VDP	Vehicle Development Process
V_{oc}	Open-circuit voltage
V_t	Terminal voltage
xIL	X in the Loop

SUMMARY

Georgia Tech has the privilege of competing in EcoCAR 3, a four-year competition in which 16 universities are given a stock 2016 Chevrolet Camaro and work to transform it into a hybrid electric sports car. In this thesis, an overview of the first two years of the author's work on the team as the Engineering Manager, the graduate student overseeing all vehicle engineering work, will be detailed. The competition will be introduced and described before a discussion on vehicle electrification and the various ways it has been achieved by manufacturers and competition teams. Next, the design of the Georgia Tech vehicle will be presented with a focus on powertrain and supporting component selection. The vehicle model underlying many of these decisions will then be discussed in detail, showing how the team used Simulink and Engineering Equation Solver to effectively predict vehicle performance, emissions, energy consumption, and cooling needs. Building on this, the controls design process known as model/software/hardware in the loop will be discussed in the context of the Georgia Tech team's use of this process. Finally, a progress update will be given, including photos of the team vehicle in current build state weeks before the Year 2 Competition.

INTRODUCTION

Advanced Vehicle Technology Competitions

As the use of fossil fuel-powered vehicles has grown in importance throughout the last century, so has the awareness that its improvements to the human condition are offset by growing concerns about the effect of these vehicles' emissions on human health and atmospheric chemistry. By now, it is of near unanimous agreement that steps must be taken to reduce the greenhouse gases and other emissions that are by-products of the combustion process, but as will be explored in detail, this is neither an easy nor cheap task.

In order to prepare themselves for the coming decades of greater vehicle hybridization and electrification in response to these concerns, members of the North American auto industry have since 1988 co-sponsored, with the United States Department of Energy, a series of university engineering competitions known as Advanced Vehicle Technology Competitions (AVTCs) [1]. Each of the 11 competitions so far has focused on a different technology or vehicle class, roughly following the trends of potential upcoming technological breakthroughs.

The AVTC series began in 1988 with the Methanol Marathon and was followed by the 1990-1993 Natural Gas Vehicle Challenge, each of these focusing on converting vehicles to run on alternate fossil fuels with generally more favorable emissions profiles than gasoline. The HEV Challenge came next from 1992-1995 on the eve of the introduction of the Toyota Prius as the world's first mass-produced hybrid electric vehicle (HEV) in 1997. 1995 through 1997 saw the Propane Vehicle Challenge, similar in theme to the first two, alongside FutureCar from 1995 to 1999, which featured teams working with various fuels, powertrains, and lightweighting techniques in partnership with the Clinton administration's Partnership for a New Generation of Vehicles initiative. After the Ethanol Vehicle Challenge of 1997-1999 came FutureTruck (1999-

2004), which was much like FutureCar in promoting a variety of advanced methods to reduce emissions and energy consumption, this time from an SUV. After these, the AVTC series took a more hybrid-focused approach, with Challenge X (2004-2008), EcoCAR (2008-2011), EcoCAR 2 (2011-2014), and the current competition, EcoCAR 3 (2014-2018).

The automaker sponsoring these competitions has changed over the years, including Ford in FutureTruck, Chrysler in FutureCar, and GM in all competitions since Challenge X, but many elements have remained the same. Student teams are tasked with significant modifications to stock vehicles for the purpose of enhancing environmental friendliness. The Department of Energy and automakers invest heavily in the teams' efforts so that their students are trained in the skills required to make advancements for actual production vehicles, thus seeding the future with upcoming generations of environmentally-conscious and well-trained engineers. And in the past few competitions, an emphasis has been placed on following a modified version of the General Motors Vehicle Development Process, which will be covered in detail, allowing students the opportunity to see their vehicles as not just a project but as a product throughout various development stages from ideation to consumer-readiness.

Figure 1: Members of the Georgia Tech EcoCAR 3 team at work

Georgia Tech has had the opportunity to compete in a number of these competitions, starting with FutureTruck in 1999, EcoCAR in 2008, and now EcoCAR 3. This thesis is focused on the first two years of the Georgia Tech team's work on EcoCAR 3, with a focus on the structural and thermal engineering, vehicle component packaging, and modelling and simulation efforts led by the author.

Overview of Vehicle Electrification

Electricity is one of many options available for energy storage, transportation, and usage, but its conversion from stored energy into rotation energy is by far one of the most efficient such conversions, and thus its implementation in vehicles is one way to greatly reduce the ineffiencies inherent to gasoline-powered movement.

The discovery of gasoline and subsequent improvements in its usage throughout the 19[th] century gave humankind an unprecedented ability to transport themselves quickly and over great distances. The energy density of gasoline, specifically, is responsible for the enormous changes that have occurred in society because of the invention of the internal combustion engine, which harnesses the exothermic combustion reaction of gasoline and oxygen in a cylindrical pressure vessel, aptly known as a "cylinder," to push a piston connected to a rotating shaft. However, the products of this reaction, vented to the atmosphere as exhaust gases, have numerous negative effects on both the environment and people, including functioning as atmospheric greenhouse gases and causing numerous health issues in individuals who live in car-heavy areas.

Engine efficiency, measured as energy output over the total available energy present in the the gasoline consumed in producing that output, is around 20% when averaged over a range of speeds and loads. This is primarily due to the constraint of Carnot efficiency, shown below, which states that efficiency of any heat engine η_{max} is limited by the temperature of its cold reservoir T_C (where T_H below is hot reservoir temperature).

$$\eta_{max} = 1 - \frac{T_C}{T_H}$$

Beyond this, the mechanics of engines make them overall wasteful machines. Friction in the moving parts, incomplete combustion due to inavailability of air, and wasted heat through the engine block and exhaust system all represent mostly inescapable losses. However, if an alternative source of energy could be added to somehow complement the functionality of an engine, bringing out the best characteristics of fuel-powered transportation (cost and energy density, for example) while minimizing its negative effects, it would be possible to begin to reduce society's overwhelming dependence on conventional vehicles. This is the philosophy behind the concept of the hybrid vehicle- keep the engine for now at least while it's still necessary, but reduce its usage (and only operate it when it would be at its most efficient operating regimes) by adding another form of energy, thus "hybridizing" the powertrain between the multiple energy sources. Numerous ways to do this exist, with options for the second source of energy including compressed air, flywheels, and pressurized hydraulic fluid, but it is the hybrid *electric* vehicle that has so far enjoyed the most success at coming to market and reducing fuel-based emissions and energy consumption.

While numerous attempts at hybrid and electric vehicles have been made throughout the years, the story of commercially successful vehicle electrification best begins with the introduction of the Toyota Prius in 1997. Since then, a number of hybrid vehicle architectures have been proposed to maximize efficiency without compromising performance. At the same time, numerous electric vehicles (EVs) are available on the market or coming down the pipeline as of this writing as well, including the Nissan Leaf; Tesla Model S, Model X, and the forthcoming Model 3; the Chevrolet Bolt; and the BMW i3. As well, a class of HEVs known as Plug-in Hybrid Electric Vehicles (PHEVs) has been developed that are somewhere in between EVs and HEVs. These vehicles feature large batteries that are meant to be charged when the vehicle is not driving, and then when the vehicle is driven, only electrical energy will be used for propulsion until the battery reaches a state of charge low enough to require normal hybrid operation.

Despite this progress, it is still unclear how the transition away from conventional vehicles will take place due to a lack of strong electrical infrastructure to support EVs and the high cost and low energy density of batteries and fuel cells, among other issues. The EVs just mentioned all have a low range or a high cost, making them impractical for many consumers. Thus HEVs represent a clear stepping stone, at least, by utilizing existing gasoline infrastructure and providing vehicles with long ranges at competitive prices to standard vehicles while moving society towards a more environmentally-friendly and electrified method of transportation.

When classifying hybrid electrical vehicles, it is common to refer to them as having a Series, Parallel, or Power-Split powertrain architecture based on their combined fuel-electricity power flow. Describing the layout and operating principles of these architectures is a helpful way to present what has thus far been treated at a surface level- how exactly HEVs work- so these architectures will be explored here in detail.

Series HEVs

The series HEV is the simplest of true hybrid architectures, featuring an engine powering a generator (which is a simply an electric motor turned against an electromotive force [EMF] instead of by it) which charges a battery. The current from the generator, as well as the battery as necessary, is then used to power another motor that drives the wheels. In this configuration, the engine does not in any way provide torque to the wheels, rather it is used to produce electricity.

The benefit of this architecture is in its simplicity- the engine can be run at its most efficient operating point as an on-board generator, and otherwise the vehicle is just an electric vehicle. The engine can also be smaller than would otherwise be required for a given vehicle, as its role is to provide a steady source of power rather than provide a wide range of torques and speeds. The downside to this architecture, however, is in the number of energy conversions, each of which carrying an associated efficiency factor. While the total system efficiency can be higher

than that of a gasoline-powered conventional vehicle even with the number of energy conversions, these losses are non-negligible and would be best avoided. As well, because the traction motor is the only source of power to the wheels, it must be large enough (and thus incurring a corresponding cost and packaging difficulty) to power the vehicle in a range of driving scenarios.

Because of these limitations, commercial series HEVs have been offered but in small numbers. Examples include the BMW i3 with Range Extender (the i3 is an EV otherwise) and the Fiskar Karma.

Parallel HEVs

A parallel HEV architecture takes the opposite approach to traction by, as the name suggests, providing parallel energy paths. Both the engine and any number of motors are connected to the wheels, and this availability of both sources of power enables the selection of either or both for different driving situations. For example, the engine can be used for steady state highway speeds where its operating characteristics will provide effective combustion, while the motor can be used during stop-and-go urban driving where its regenerative braking capabilities will allow for much of the energy spent during acceleration to be recouped during stopping. When the driver demands high acceleration, both sources can be used to propel the vehicle, and if the battery state of charge gets too low, the motor can not only provide regenerative braking during stopping events but can also act against the engine, thus acting as a generator and forcing the engine into a more efficient higher-throttle operating regime.

Beyond these control options, another benefit of this architecture is in component sizing. Because both power sources are available to provide power to the wheels, both can be employed for high torque demands, and so both can be downsized from what would be needed if they were to power a vehicle by themselves.

The use of gears, shafts, and belts allow for this architecture to be laid out in numerous different configurations. One construction used in the EcoCAR 3 competition to describe different options is P-notation, i.e. P1 and P2. The "P" refers to "parallel," and the number refers to the location of the motor in relation to the engine, with definitions listed in Table 1 below.

Table 1: P-Notation

Layout	Location of Motor
P1	On the accessory side of the engine
P2	In between the engine and transmission (on the crankshaft side of the engine)
P3	After the transmission
P4	On a different axle than that driven by the engine

These options will now be presented in detail.

P1

P1 Parallel HEVs generally have a small electric motor coupled to the engine by a belt on the engine's accessory side. This motor, running on a voltage higher than the standard automotive 12 volts but often at only 36 or 48 volts (most HEVs and EVs operate at an order of magnitude higher), acts as both an alternator and a starter motor. Because of its location in place of a standard alternator it is often referred to as a Belted Alternator-Starter or BAS motor. Essentially the least powertrain-altering way to turn a conventional vehicle into a hybrid, the usage of higher voltage allows the engine to be more quickly restarted than with a standard 12 volt starter motor. This in turn allows the engine to be turned off at short stops, such as at a red light or stop sign, in a way that doesn't significantly impede a driver's desire for torque upon pressing the gas pedal.

The only fuel savings in this form of hybrid come from turning the engine off when it would otherwise be idling, but because of its simplicity in implementation it has been commercially

available in numerous vehicles, including pre-2016 Chevrolet Malibu hybrids and other GM hybrids.

P2

A P2 parallel HEV places a motor between the engine and transmission, a strategy that immediately calls to mind the issue of packaging. Because transmission bell housings are generally not sized or shaped for the presence of a motor, and because there is usually very little room in a vehicle to push the engine forward or transmission backward (or side-to-side with a transverse transmission) to make room for a motor, packaging a P2 usually requires a very thin motor built for the purpose. If that can be arranged, however, the torque of the motor will be available over a wider range of vehicle speeds than otherwise because it will be multiplied through the same multi-speed transmission as the engine.

This architecture has been featured in most of Honda's hybrids, including the Insight, as well as the Porsche Panamera Hybrid SE.

P3

A P3 hybrid features a motor attached to the drivetrain after the transmission, removing the benefit of multiplying the torque by the transmission gear ratios but generally lightening the packaging issues. Options include a motor inline with the driveshaft or wheel axles, spinning at relatively lower speeds and needing to produce higher torque than a P2 motor for the same vehicle acceleration (at most low-to-medium speeds at least), or one coupled to the drivetrain with a torque ratio of some sort, such as by using a belt drive or a gearbox.

As of this writing, no commercially available P3 parallel HEVs have been offered to the market, though the Georgia Tech EcoCAR 3 team vehicle is itself a P3 parallel, and the design of such will be explored in great detail throughout this paper.

P4 "Through the Road"

A P4 parallel hybrid features one or more motors driving a separate axle from that driven by the engine, giving rise to the name "through the road hybrid" to describe how the speed coupling between the engine and motor is through their commonly traversed road. The benefits of this architecture include modular packaging and independent control of the two systems, with each physically independent of the other. Placing the motor on the front axle is a way to maximize regenerative braking, as the weight transfer to the front wheels during braking acts against skidding, but this configuration does give rise to the interesting effect of flexing the entire vehicle if the motor is braking and the engine is propelling, requiring attention during chassis and body structure design.

Commercial implementations of the P4 parallel architecture include the BMW i8, the Subaru VIZIV concept car, the Audi duo, and the Peugeot 3008 HYbrid4.

Combination Parallel Layouts

Because the preceding options place motors in different locations along the drivetrain, it is possible to combine them with multiple motors into, for example, a P1-P3 parallel, with a belted starter motor and a motor on the driveshaft, or a P2-P4-P4 parallel, with a front wheel drive (FWD) transverse transmission, a built-in P2 motor in the transmission's bell housing, and separate motors driving each of the rear wheels. While these layouts are more common in competition cars such as those produced in AVTCs than those produced by major automakers, commercial examples include Acura's RLX [2] and 2016 MDX Hybrid [3], which both feature the P2-P4-P4 orientation just described.

Some combination parallel layouts can provide operating characteristics similar to the Power-Split architecture described below. A P1-P3, for example, could allow the P3 motor to drive the wheels, the P1 motor to act as a generator, and the engine to provide torque to both the

generation and propulsion paths. This architecture is also known as a Series-Parallel hybrid since it combines elements of each of its namesakes.

Power-Split HEVs

The Power-Split is generally considered the most complex hybrid architecture option available. This architecture also combines the functionality of a series hybrid and a parallel hybrid by allowing the power of the engine to go fully to the wheels, fully to charging the battery, or anywhere in between. This is usually accomplished through the use of one or more planetary gearsets and it allows for constant optimization of fuel and electrical energy usage by directing both forms of energy in a virtually infinite combination of paths. The presence of these planetary gears and the continuously variable transmission they provide distinguish a Power-Split hybrid from, say, a P1-P3 series-parallel hybrid with discrete gear ratios.

Two commercial implementations of this powertrain have enjoyed considerable success on the market, GM's Voltec powertrain and Toyota's Hybrid Synergy Drive. The first is featured on a number of vehicles, including the Chevrolet Volt and 2016 Malibu hybrid, and the second is featured most notably on the Toyota Prius, as well as numerous other Toyota and Lexus hybrids and the Nissan Altima hybrid.

Figure 2: Georgia Tech EcoCAR 3 Shop

With the background of AVTC history and the current state of vehicle hybridization and electrification in mind, the EcoCAR 3 competition that is the focus of this paper will now be introduced. EcoCAR 3 gives 16 universities the opportunity to transform a 2016 Chevrolet Camaro into a hybrid electric performance vehicle. These universities are:

- Arizona State University

- California State University Los Angeles

- Colorado State University

- Embry-Riddle Aeronautical University

- Georgia Tech

- McMaster University

- Mississippi State University

- Ohio State University

- Penn State University

- University of Alabama

- University of Tennessee-Knoxville

- University of Washington

- University of Waterloo

- Virginia Tech

- Wayne State University

- West Virginia University

This competition, which began in 2014, comes in the market context of increasing excitement at the prospect of vehicle electrification, especially in the performance vehicle segment. The Tesla Model S, released in 2012, features impressive acceleration and safety features, and numerous higher-end automakers including Porsche (with the Panamera Hybrid) and BMW (with the i3 and i8) are now offering vehicles that take advantage of the high and available low-end torque provided by electric motors.

The Chevrolet Camaro, first introduced in 1966 as a competitor to the Ford Mustang, is a classic American sports car that prioritizes power and handling in its sell to consumers but is by no means generally regarded as terribly fuel-efficient, making it a fantastic target for hybridization at the hands of student teams. The 2016 model year is the first of a new generation of the vehicle and features roughly 200 pounds of weight savings over the previous model, giving the teams a cutting-edge product to work on.

Figure 3: Camaro delivery truck stopping at Georgia Tech. Pictured here, Alabama's Camaro.

Each of the schools participating in EcoCAR 3 were given a brand new 2016 Camaro in late 2015 with the goal of reducing emissions and maximizing energy efficiency through powertrain hybridization. The overall timeline of the four-year competition, as previously mentioned, follows a modified version of the GM VDP as follows.

In Year 1, teams received the CAD files of the entire Camaro and a MathWorks Simulink model of the stock vehicle powertrain and longitudinal dynamics. These resources were used to develop a plan for the team's transformed vehicles, with the year's activities and end-of-year Year 1 competition featuring an emphasis on packaging studies and modeling and simulation.

In Year 2, teams received their vehicles and after initial testing disassembled them, removing the stock powertrain, and began installation of their new hybrid powertrains. As of this writing

near the end of Year 2, teams are well into the work of integrating their motors, engines, transmissions, and all other components needed to make a fully-functioning HEV. Vehicles will be judged at the Year 2 Competition on their extent and quality of integration as well as initial functionality.

In Years 3 and 4, the focus of the competition will shift to producing consumer-ready vehicles with optimized controls and finessed driver experiences. Vehicles will be expected to start quickly and easily, drive over long distances, and perform as a Camaro-purchaser would expect their sports car to. Teams will be judged on a balance of energy consumption, emissions, performance, handling, consumer acceptability, and more.

Emissions and Energy Consumption Scoring

While the methods for evaluating team vehicle acceleration, handling, ride quality, and other characteristics in the EcoCAR 3 competition are fairly straightforward, the ways that environmental impact is measured are more involved. The EcoCAR 3 Emissions and Energy Consumption (E&EC) scoring philosophy is based on the methods used by the United States Environmental Protection Agency (EPA) to evaluate vehicles' overall performance in these areas when they have a blended source of energy, as do HEVs. The approach simply measures net energy consumption of both sources over a set driving schedule, known as a drive cycle, and sums them for overall energy consumption. By using a standardized value for the energy content of gasoline, this number can be expressed in units such as kWh/km (energy per distance) or converted into mile per gallon gasoline equivalent, or mpg$_{ge}$.

The advent of PHEVs, however, required new methods to be introduced to equitably rank these vehicles. Because the standard EPA drive cycles are no more than a few miles long, a plugin HEV with a long enough electric range could easily perform each of the drive cycles on electric-only power and thus avoid any vehicle-side emissions or fossil fuel consumption. In reality, many owners of these vehicles will frequently exceed their electric range, also known as Charge

Depleting (CD) range, requiring the gasoline-powered engine to in order to continue driving and thus introducing its emissions into the surroundings. This second mode of operation is often referred to as the Charge Sustaining (CS) mode, as it is a common strategy of PHEV makers to use engine power in combination with electric power in such a way that battery state of charge is held roughly constant over time.

To judge PHEVs alongside HEVs and conventional vehicles, a number known as the Utility Factor (UF) was introduced. The UF, a number between 0 and 1, corresponds to a given distance and represents the fraction of Americans who drive that distance or less in a single day. For example, if 40% of American drive 20 miles a day or less (this being an example, not necessarily factual), the UF associated with the distance of 20 miles is 0.4. A vehicle's CD range is used to determine its UF, and then this UF is used to calculate a weighted average between CD driving characteristics and CS driving characteristics. In equation form, with EC representing Energy Consumption,

$$EC_{total} = (UF)(EC_{CD}) + (1 - UF)(EC_{CS})$$

Thus it is clear that a vehicle featuring a longer CD distance, which is provided by larger battery capacity, will be rewarded with (and in this competition, judged favorably for) a lower energy consumption rating and lower emissions, as the average of its CS operation results with its more efficient electrical CD results will be more heavily weighted towards the CD driving.

Of note, the actual correlation between CD distance and UF, as defined in the competition rules [4], is

$$UF = 1 - e^{-\left[C_1\left(\frac{x}{D_{norm}}\right) + C_2\left(\frac{x}{D_{norm}}\right)^2 \dots + C_6\left(\frac{x}{D_{norm}}\right)^6\right]}$$

With the following variable definitions.

Table 2: UF Equation Variables

x	CD Distance
D_{norm}	399.9
C_1	10.52
C_2	-7.282
C_3	-26.37
C_4	79.08
C_5	-77.36
C_6	26.07

With this framework in place, it was desired to not only transform the team's Chevrolet Camaro into a fuel-efficient hybrid but also to increase its competitive scoring through an increased CD distance, provided for by a larger capacity battery pack. This will be explored in a later section.

Vehicle Performance Scoring

While Emissions and Energy Consumption comprises a large portion of available points in the year-end competitions in EcoCAR 3, numerous other categories exist and should be mentioned. A discussion of every single category and their point values in different years would get rather arcane, but a few items are worth mentioning.

Vehicle acceleration is measured in two ways, IVM-60 (as opposed to the more standard 0-60) and 50-70 (passing acceleration). IVM, which stands for Initial Vehicle Movement, is measured as the point in time at which the vehicle has travelled one foot after the driver has, with the vehicle at rest, fully depressed the accelerator pedal. This test is used in place of the more recognizable 0-60 test to allow for fair comparison between the teams' vehicles with different

drivers. For the 50-70 mph acceleration test, team vehicles are brought to a steady-state 50 mph and then accelerated to 70 mph.

These two tests, in combination with a number of vehicle handling evaluations, are included to emphasize the goal in this competition of producing a vehicle which can be enjoyed by the average consumer of a stock Camaro. It is incumbent upon the teams to meet all of these various goals in order to have a well-balanced vehicle.

A Note on CAD images

One of the great privileges of working directly with General Motors in this competition was their sharing of CAD files for the entire Camaro and numerous other components, including their donated engines and transmissions. While CAD has great potential to clearly convey design decisions and where components are located, and was thus extremely useful to the team's efforts, unfortunately very little CAD can be included in this paper in order to comply with non-disclosure agreements signed at the beginning of the competition. Photographs of the actual subjects of discussion will be included where possible, but as of this writing much of the vehicle is still not fully integrated, and there are even major subsystems that haven't been constructed yet. Thus words will have to suffice for some of these topics.

CHAPTER 1: DESIGN OF A P3 HYBRID ELECTRIC CAMARO

The EcoCAR 3 competition begins with a simple-sounding premise: take a stock 2016 Camaro and turn it into an electric or hybrid electric vehicle. The process of doing so, however, involved an enormous amount of planning, analysis, and fabrication work. In this section, the overall design process followed by the Georgia Tech EcoCAR 3 team will be covered through the initial architecture selection process and the specifics of building a hybrid Camaro.

Throughout this section, a table will be used to show the progress and results of the team's design and component selection at each stage of the decision-making process. Starting here is the stock 2016 Chevrolet Camaro. Note that the transmission option included below and in many following iterations was based on best knowledge available at the time but was later proved incorrect, as will be described.

Table 3: Architecture Selection Process 1: Stock Vehicle

Engine	3.6 liter V6 (LGX)
Transmission	6-speed longitudinal automatic (6L45)
Architecture	RWD gasoline-powered

Architecture Selection

Initial Decisions

As previously elaborated, hybridizing a vehicle can take numerous forms. Because "hybrid" means a blend of any two (or more) things, a hybrid Camaro doesn't necessarily have to be a hybrid electric Camaro, or electric at all. Fuel cell hybrids (electric but without batteries), hydraulic hybrids, compressed air hybrids, not to mention the various types of engines and fuels, such as propane, natural gas, ethanol, and diesel- all of these were options to be

considered. However, the interests of the team and the overwhelming emphasis of the competition on electrification prompted a decision to pursue the design of an HEV, and a few guidelines and rules that furthered simplified the universe of options were set at the beginning of the competition. The most important were:

- Remove the Camaro's stock engine, the 3.6 liter LGX V6
- Remove the stock fuel tank and replace it with one that has a capacity of 10 gallons or smaller (largely to enable ease of handling during competitions)
- The vehicle must be able to travel around 150 miles (exact distance TBD)
- Do not use hydrogen fuel cells
- Torque balance between front and rear axles must be at least 51% biased rearward
- Maximum vehicle and axle weights and distributions could not exceed the values in Table 4 and Table 5 below.

Table 4: Weight and Distribution Limits Without Occupants

Front Axle Weight Distribution (minimum)	43%
Rear Axle Weight Distribution (maximum)	57%
Total Weight	1938 kg

Table 5: Weight Limits With Occupants

Front Axle	1000 kg
Rear Axle	1170 kg
Total Weight	2120 kg

Figure 4: Stock LGX engine, to be removed

The next most important factor was the availability of sponsored components. GM not only

offered a selection of engines and transmissions for donation, as did A123 and Bosch for

batteries, but these manufacturers also offered continuing support for their products through the

duration of the competition. It would be difficult for a new team like Georgia Tech to

independently secure component sponsorships that could offer competitive alternatives to these

products, so the team decided to rely on these donated options as much as possible.

Table 6: Architecture Selection Process 2: Initial Decisions

Engine	General Motors-made
Battery	A123 or Bosch
Architecture	BEV or HEV

Architecture Selection

A major variable in component and architecture selection is packageability. Any engineered system will face geometric space constraints at some point it its design- it's not an infinite world that we live in. The packaging considerations present in the EcoCAR 3 competition were particularly difficult and presented a significant challenge to the team's early efforts to plan the selection and layout of its components.

When an automaker designs a modern vehicle, it is standard practice to base that vehicle on a platform, a vehicle structure that can be used for various models intended for different applications. For example, a truck and an SUV, or an entry-level sedan and a more upscale version, can share much of the same underlying engineering work, physical structure, powertrain components, and more while having completely different body paneling, front and rear fascias, engines, and interior qualities. While automakers strive to reuse their engineering efforts through these shared platforms as much as possible, they have the freedom to choose whichever platform they wish or design a new one entirely if packaging or structural strength concerns dictate that no existing platform is sufficient for an application.

In the EcoCAR 3 competition, however, teams are tasked with integrating significant battery capacity into a pre-built, small performance vehicle with no space allocated for it. They must then integrate an electric motor that is large enough to offset the power loss caused by removing the stock engine (while one team, Virginia Tech, decided upon an architecture including a bigger engine, all other teams in the competition opted for smaller engines that are likely to be more fuel efficient than the stock V6). Contending with the engine replacement itself is a non-trivial task, as the Camaro's chassis was designed with this stock LGX engine in mind and a smaller displacement does not ensure fit in every dimension. Finally, it was necessary to find some place to put all the other somewhat smaller components, such as the inverter, battery

charger, supervisory controller, high voltage junction box, electronic air conditioning compressor, high voltage-to-low voltage DC/DC converter, and more.

While these packaging concerns played an outsized role in the Georgia Tech team's architecture selection, model and simulation was also used as a primary driver for decision-making. Teams worked in parallel on these two design aspects to converge on a possible and optimal set of options.

With these two major drivers discussed, the process by which a P3 parallel PHEV was selected as the team vehicle will now be reviewed by working through the decision-making process architecture by architecture.

Battery Electric Vehicle

It was quickly realized through modeling work that a potential best-case scenario for this competition was to convert the Camaro to a full BEV. After all, electrical energy consumption is many times more efficient and generally less pollutant-emitting than is fossil fuel consumption. However, as previously mentioned, batteries are nowhere near as energy dense as gasoline is by volume or mass, meaning that a BEV Camaro would require an enormous battery pack with far more capacity than those offered by the competition sponsors, which would thus require a major sponsorship effort, significantly increased weight, and decreased cargo and/or packaging space. Given the maximum weight regulations and packaging constraints for this vehicle, it was decided to pursue a hybrid electric architecture instead of a BEV.

Table 7: Architecture Selection Process 3: HEV

Engine	General Motors-made
Battery	A123 or Bosch
Architecture	HEV

With that decision made, the team pursued an exhaustive search strategy through the various HEV architecture possibilities to determine which options would be feasible.

Initial Ruled-out Options

A simple P1 hybrid and a Power-Split hybrid were quickly ruled out as being too simple and too complex, respectively, for the Georgia Tech team's capabilities. Future improvements to the vehicle could include doing both, interestingly- adding a P1 BAS motor to the powertrain would actually enable series-parallel operation by using the P1 as a generator during driving. Whether or not this strategy will be pursued is a decision left to the team in future years, but for the first two years it was ruled out to focus efforts on a single and substantial electrical source of traction.

P2 Hybrid

One way to utilize an electric motor's torque is to place it between the engine and transmission in a P2 configuration, which, as described in the Introduction, is used by Honda for many of their HEV models. A multi-speed transmission allows for a torque source to operate in a narrower speed range than otherwise, which is how engines are able to operate over such a wide range of vehicle speeds while generally staying within a few thousand RPM of idle speed. The same benefit applies to a motor in a P2 configuration, allowing it to be selected without its maximum allowable speed being a major limiting factor (which, as will be seen, plays a major role in P3 motor selection). This does, however, mean that while the motor may stay in or near its peak torque speed range for all practical speeds experienced by the vehicle, the motor torque to the wheels is effectively decreased at increasing vehicle speeds by the decreasing gear ratios.

With these tradeoffs in mind, it was determined that no commercially available motors would fit between the engine and transmission without causing major interference between the engine and the structure of the vehicle. While modifications to the vehicle structure are allowed in the

EcoCAR 3 competition, they require approval from GM based on FEA simulation of various load cases to determine that the requested changes will not harm the structural integrity of the vehicle. This analysis would be very time-consuming, as would the modifications themselves, so it was decided that pursuing a P2 architecture would be too costly for a time- and resource-limited team. It is possible that this was a missed opportunity, as it appears that other teams were able to select and integrate P2 motors into their Camaros, but it is unclear if this was accomplished without needing structural modifications.

P4 "Through-the-Road" Hybrid

Next, the P4 architecture was explored. There are two options for a P4 HEV, the engine driving the front wheels and the motor driving the rear wheels, or vice versa (one might add that AWD transaxles and transfer cases can be used to broaden these possibilities, such as engine AWD with motor RWD or FWD, but the packaging constraints discussed remain essentially the same regardless). The benefits of this architecture include the fact that the two powertrains are completely modular- the engine/transmission and motor/transmission combinations can be built, integrated, and tested independently, and either could drive the car by itself without needing the other functioning. As mentioned, a potential downside would be invoked by placing the motor at the front axle with the engine driving the rear, in which case regenerative braking during engine propulsion would cause vehicle structure flexing, but this would also enable higher regenerative capabilities versus placing the motor on the rear axle.

It was quickly discovered that fitting a transverse transmission in the chassis of a Camaro was not feasible with any of the donated engine options (ruling out engine-front/motor-rear), and no motor of decent size (as determined by vehicle modeling) would share the engine bay with a full-sized engine (ruling out the inverse). Due to these packaging constraints, this architecture was removed from consideration.

While the Camaro is a fairly compact car without much wiggle room for added components, an opportunity presented itself in the required removal of the stock fuel tank. As shown in Figure 5 below, the stock fuel tank is in the configuration known as a saddle tank, spanning both sides of the vehicle and straddling the driveshaft and exhaust.

Figure 5: Camaro stock underbody with saddle tank spanning exhaust routing

Being required to remove this tank and replace it with one of 10 gallons or smaller makes it a logical decision to build a new tank that only occupies the tank cavity on one side of the driveshaft. As will be explored in a later section, the final fuel tank design for the team's vehicle has ample capacity while only using the passenger side of the vehicle, where the stock fuel filler neck is located. Thus the remaining cavity on the driver's side allows for the placement of an electric motor that can be coupled to the driveshaft.

Various other locations were explored for the P3 architecture, including on the transmission output and inline on the driveshaft, but a couple of concerns precluded these placement options. The driveshaft tunnel itself was designed for just that, a driveshaft, leaving no room for the addition of a motor, and placing a motor inline further down the driveshaft was troublesome for

two reasons. For starters, few manufacturers of compact electric motors make products that can connect to a shaft on both of their axial ends, as this is a non-standard application and the rotor shaft would be required to withstand the same torques that are applied to a vehicle driveshaft in such a configuration, namely, the torque of the engine multiplied by the largest transmission gear ratio, in addition to its own electromotive torque. While this is of course possible, it appears to generally not be a design requirement for these manufacturers. More importantly for our team, if a motor is placed on the driveshaft, its torque to the wheels is only multiplied by the final drive ratio. Motors that produced enough torque to the wheels to meet our competition acceleration goals in this location (as predicted by our vehicle simulation) were all too large to fit in this location. Thus the team focused its efforts on the space left by the fuel tank and coupling the motor located there to the driveshaft with some form of gearing.

In summary, because the various options mentioned were ruled out in turn and because of the packaging opportunity granted by the fuel tank, the P3 architecture was selected for the team vehicle.

Table 8: Architecture Selection Process 4: P3 Parallel HEV

Engine	General Motors-made
Battery	A123 or Bosch
Architecture	P3 Parallel HEV

Specification of P3 Components

With the overall architecture and placement of components decided, a few specifics needed to be addressed such as the engine, batteries, motor, and other related components.

Engine Selection

Because the vehicle to be modified was a General Motors product, with a GM CAN bus, and because the offer was already on the table, it was decided that best engine selection would be one of the donated engine options, which are listed below in Table 9. It would not only be given to the team free of charge but would also include ongoing technical support, and the ability to naturally interface with the vehicle's CAN bus would be extremely helpful for getting it to work in the first place.

Table 9: GM-donated Engine Options

Engine Code	LEA	LVL
Fuel	E10 or E85	B20 (20% biodiesel)
Displacement (L)	2.4	1.6 Turbo
Maximum Torque (Nm)	233	321
Maximum Power (kW)	136	100

CAD packaging studies showed that the two engines, despite the 50% difference in displacement, occupy roughly the same footprint in all three dimensions and would each fit in the engine bay of the Camaro, so either would work from an initial packaging study perspective. However, while the diesel LVL offered impressive torque, modelling results showed that the negative emissions characteristics associated with diesel combustion would offset any benefits it provided. Meanwhile, the LEA was E85 compatible, or "flex-fuel". E85 is 85% composed of ethanol that, in the United States, is largely produced from corn. Because of this former life as a carbon dioxide consumer, even with its combustion taken into account it still has a net negative effect on carbon dioxide emissions. Thus its usage is favorably weighted in the EcoCAR 3 competition, and for these reasons, the LEA was selected as the team's vehicle's engine.
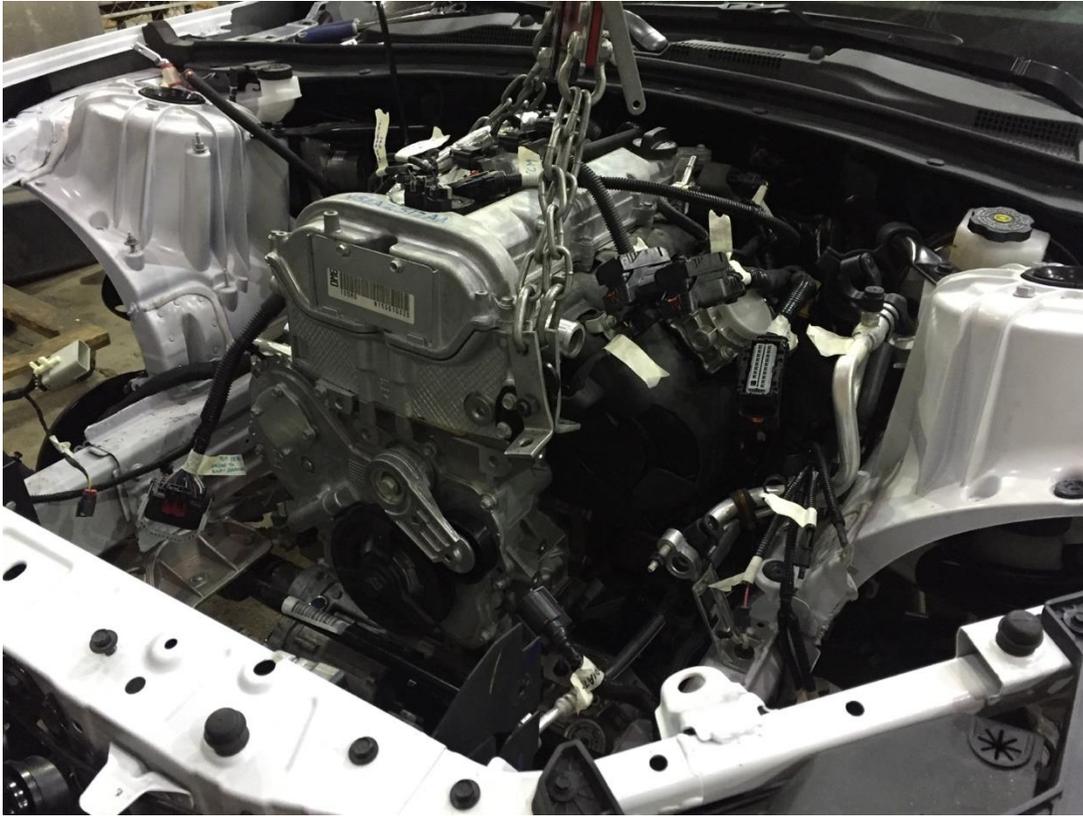
Figure 6: Installation of the LEA engine into the Camaro

Table 10: Architecture Selection Process 6: Engine Selection

Engine	GM LEA
Battery	A123 or Bosch
Architecture	P3 HEV

Battery Pack Selection

Battery selection was a fairly straightforward tradeoff between the benefits of high capacity and the drawbacks of the corresponding higher weight and volume associated with that capacity. Selecting between the sponsor-donated options, however, required full simulation of the proposed vehicle with each option as well as packaging studies to determine feasibility.

As previously discussed, the trunk was found to be the only feasible place to install the batteries. A meaningful point penalty is associated with removing the back seats of the Camaro in this competition, and it wouldn't be feasible to remove the engine from the front of the vehicle to put the batteries in the engine bay, so the trunk was the only remaining option.

The batteries offered as donations to each of the competition teams took the form of either packs or kits. Packs were pre-built, including all wiring, sensors, control modules, and structures. Kits required teams to design and build the structures and wiring needed to house and connect the individual battery modules, and teams would be required to go through a multi-stage design review process under the guidance of engineering consultancy group Ricardo in order to have approval to construct and use their team-built packs. The options offered by competition sponsors A123 and Bosch are listed below.

Table 11: Available Battery Kits

Kit Name	Energy
A123 6x15s2p	10.8 kWh
A123 6x15s3p	16.2 kWh
A123 7x15s2p	12.6 kWh
A123 7x15s3p	18.9 kWh

Table 12: Available Battery Packs

Pack Name	Energy
A123 100s1p	4.5 kWh
A123 116s3p	21.6 kWh
Bosch 8x104s2p	10.7 kWh

The pack and kit names listed above contain important information in the standardized format of [number of modules in series] "x" [number of cells in series per module] "s" [number of cell

31

groupings in parallel per module] "p". So the first kit, the A123 6x15s2p, would contain 6 modules, each module containing two parallel rows of 15 cells in series.

Going into battery selection, the initial desire of the team was to select a pre-built pack in order to best budget limited time and resources. As well, while simulations of the various packs showed that there were negative emissions and energy consumption consequences to choosing a larger, heavier option, these effects were well offset by the increased all-electric distance they made possible. These two factors made the A123 116s3p pack look the optimal selection.

However, simultaneous packaging studies revealed that neither the A123 pre-built pack nor the Bosch would fit in the limited Camaro trunk space, effectively ruling each of them out. Looking at the remaining kits, it was discovered that each of them could fit in the Camaro trunk but the 3p modules, which were larger than their 2p cousins, would significantly decrease available clearance. These lines of reasoning left three A123 options: 100s1p pack, 6x15s2p kit, and 7x15s2p kit. Since the 100s1p pack had such a comparatively small capacity it was ruled out, and the larger of the two kits, the 7x15s2p, was chosen because the seventh module did not introduce a significant marginal packaging difficulty vs the six-module kit.

Note that while in this section a difference between "packs" and "kits" has been introduced, the team's batteries will be referred to frequently outside of this section as a "pack" to employ a commonly-used term that describes a packaged grouping of batteries, as well as to describe their aggregate nature and function in the context of the vehicle.

Table 13: Architecture Selection Process 7: Battery Pack Selection

Engine	GM LEA
Battery	A123 7x15s2p
Architecture	P3 HEV

Motor Selection

A number of motors were considered for the team's vehicle, with a few initial requirements. The motor had to be a 3-phase synchronous AC machine, the most efficient type of motor widely available for vehicular applications. It had to be able to operate on a DC bus voltage of around 340 volts in order to match the nominal voltage of the A123 battery pack selected above. And it needed to offset the loss of power incurred by using a smaller engine than the stock V6. With these requirements in mind, motors produced by Brusa, Parker, and TM4 were considered.

The Brusa motors immediately stood out for their compact construction and high torque capabilities. However, there was a premium on these benefits – the Brusa motors cost tens of thousands dollars more than the Parker and TM4 motors, putting them out of reach for the Georgia Tech team.

TM4 produces two motors for use in EVs and HEVs, the MOTIVE and the SUMO, the latter so named for its intended use case in busses. The SUMO provided enough torque to not require a gear ratio to the driveshaft at all, but unfortunately could not fit anywhere without conflicting with the vehicle body structure. The MOTIVE, however, was an enticing option that produced 255 Nm of torque and would fit in the desired location.

Parker, on the other hand, publishes a catalog with literally hundreds of motor configuration options in the GVM-142 and -210 series, where GVM stands for "Global Vehicle Motor." This impressive selection contains numerous variations in axial and radial dimensions, windings, and permanent magnet strength. While TM4 had one potential product, it was only a matter of searching through the catalog to find a Parker model with the appropriate torque and speed characteristics before a suitable match was found in the GVM210-150P. This model was selected over the TM4 by the team's electrical graduate students and faculty advisor.

Table 14: Architecture Selection Process 8: Parker Motor

Engine	GM LEA
Battery	A123 7x15s2p
Architecture	P3 HEV
Motor	Parker GVM210-150P

Transfer Case

Team-built driveline connections are said to be a frequent cause of vehicle failure in AVTCs, so it was desired to find a pre-built solution to mitigate that risk. The Borg Warner 4477 transfer case, found in the 2008-2011 Cadillac CTS and shown in Figure 7 and Figure 8 below, is a chain-driven transfer case that was selected to transfer torque from the Parker motor to the driveshaft due to its compact construction and the removability of its sprockets. After analysis of the effect of various sprocket ratios on the safe operation of the motor, it was decided to use a 31 teeth to 12 teeth, or 2.583:1, sprocket ratio between the motor sprocket and the driveshaft sprocket. The custom sprocket will be designed and fabricated in future years, with the team retaining the stock 1:1 ratio for the Year 2 competition.

Figure 7: Transfer Case, showing front driveshaft connection

Figure 8: Transfer Case, showing rear drivershaft and motor connections

To convert the DC supply voltage from the battery to AC voltage to drive the motor, a suitable inverter was required. The Sevcon GEN4S8 was selected for its high efficiency, configurability to virtually any AC electric machine, and voltage/current limitations that matched the vehicle's electrical system. The inverter, shown below in Figure 9, is the dark grey component on top of the two power supplies stacked in the center.
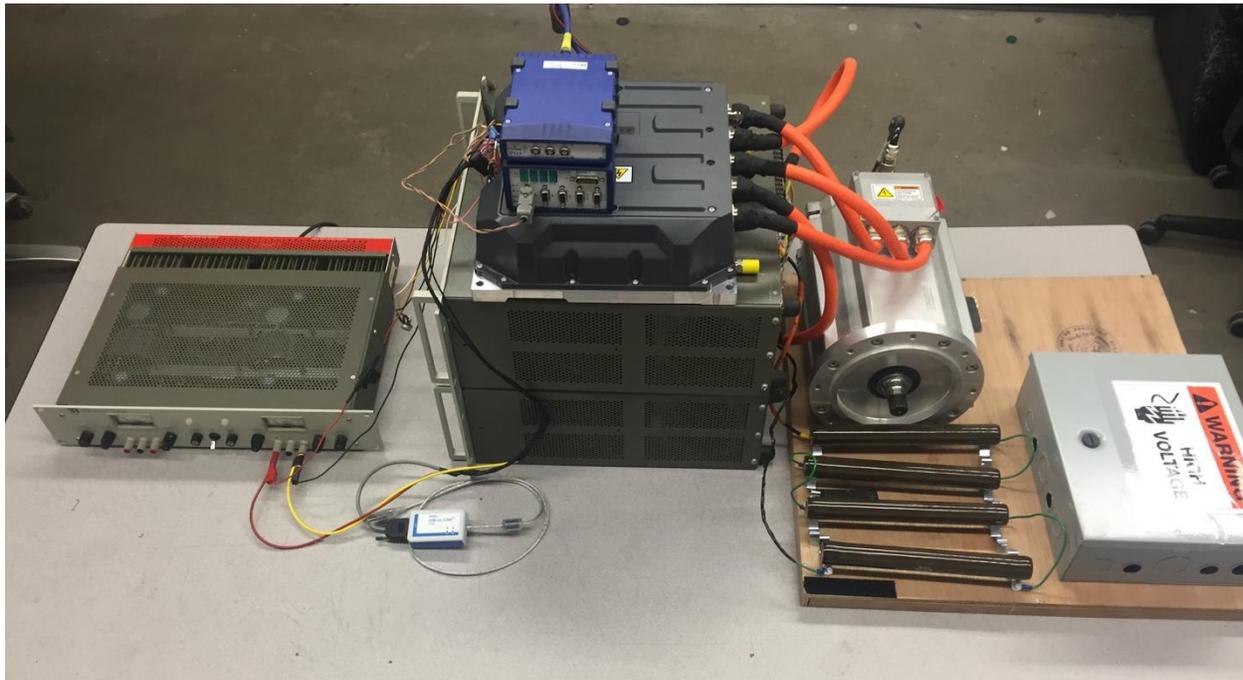
Figure 9: Motor and Inverter in bench test setup

In conventional vehicles, many auxiliary tasks such as air conditioning and 12 volt battery charging are accomplished by driving components with a belt connected to the engine's crankshaft. While this is an effective way of providing and transferring power for these functions when no high voltage energy source is present, it is by no means the most efficient way to do so, as the power provided is coming from wasteful fuel combustion. In an electric or hybrid

electric vehicle it is common to use a DC/DC converter in place of the alternator and an electric compressor in place of the belt-driven AC compressor.

The DC/DC converter receives high voltage on one side and steps it down to a level on the order of 12 volts on the other side, allowing the 12 volt battery to be recharged from the high voltage battery. In a similar way, an electric compressor uses high voltage to drive a refrigerant compressor that operates the same way a belt-driven compressor does. In the EcoCAR 3 competition, DENSO, a Japanese auto components supplier, sponsored each of these components for use by teams and so they were selected for use in the Georgia Tech team vehicle. A DENSO electric compressor is shown in Figure 10 below.
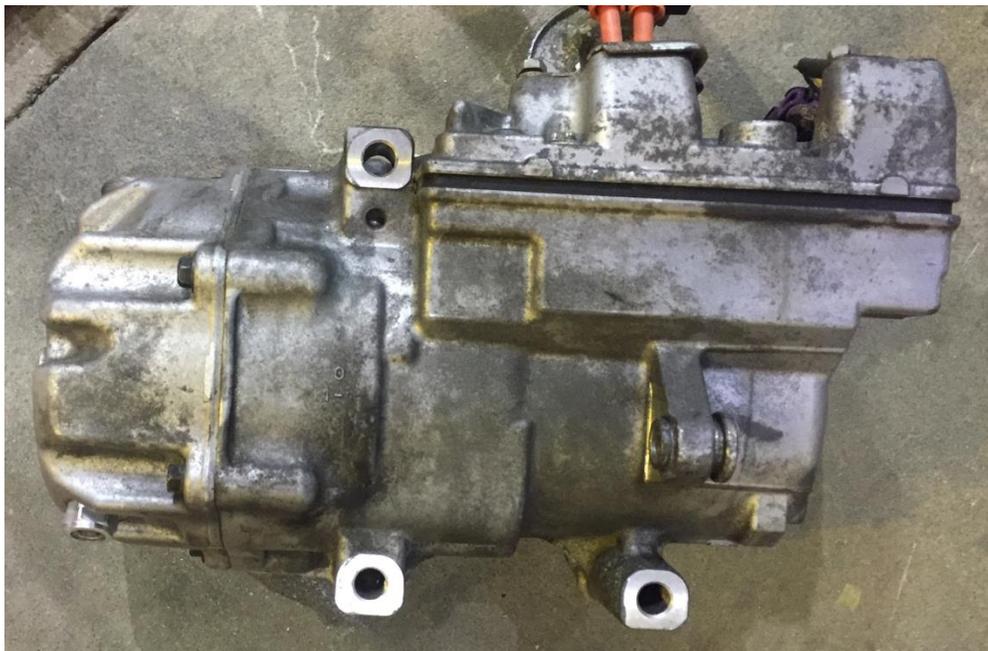
Figure 10: DENSO Electric AC Compressor

Finally, the high voltage batteries would be charged through regenerative braking and engine-powered generation through the P3 motor in CS mode, but to prepare the vehicle for CD mode

it would need to be plugged into a wall outlet. A Brusa battery charger, shown in Figure 11 below, was selected to convert the lower wall AC voltage into higher DC pack voltage.

Figure 11: Brusa Battery Charger

Table 15: Architecture Selection Process

Engine	GM LEA
Battery	A123 7x15s2p
Architecture	P3 HEV
Motor	Parker GVM210-150P
Inverter	Sevcon GEN4S8
AC Compressor	DENSO Electric Compressor
DC/DC Converter	DENSO DC/DC Converter
Battery Charger	Brusa NLG513

Transmission Selection

Notably absent from discussion thus far has been the selection of a transmission, which is of course crucial for efficiently conveying the power of an engine to the wheels. The timeline for selecting the aforementioned components was fairly compressed, requiring prioritization of time spent on each component's selection, and GM had not at the time released the transmission options available for donation. Thus for the architecture selection process it was assumed that the stock (as was thought at the time) 6L45 transmission, a 6-speed longitudinal, would be used in the vehicle with the hope that GM would eventually make known other freely available options.

Given that it will fit in a vehicle, some overall trends are evident for any transmission. First, of course, the lighter a transmission is, the less it will increase the vehicle's overall weight. Second, a GM transmission will connect more seamlessly to the Camaro's CAN bus than will another manufacturer's transmission, which is helpful from a controls perspective for getting the vehicle running. Finally, gear ratios have some predictable trends- the lower the ratios available, for example, the better fuel economy will be obtained at higher speeds, and the higher the ratios available, the quicker the possible acceleration from a stop. Combining these two trends, the greater the spread of ratios, the better the overall performance.

Two developments that were beneficial to overall vehicle planning occurred after the architecture selection process was finalized. First, GM offered the following two transmissions for donation:

Table 16: 6L45 vs 8L45 Gear Ratios

Transmission	Ratios
8L45	4.62, 3.04, 2.07, 1.66, 1.26, 1.00, 0.85, 0.66 [5]
8L90	4.56, 2.97, 2.08, 1.69, 1.27, 1.00, 0.85, 0.65 [6]

Each are 8-speed longitudinal transmissions, with the 8L90 being designed for higher torque inputs. Given the gear ratio trends mentioned above, it was immediately obvious that the 8L45, which was lighter than the 8L90 and fully capable of withstanding our engine's torque, would likely be the better option. The second development, of note, was the announcement that the 2016 Camaro would include the 8L45 as its stock transmission option, meaning that the 8L45 would fit perfectly in the Camaro. As it later turned out, a transmission mounting adapter was needed to for the LEA engine to fit in the engine bay that wasn't designed for it, but this was a minor design and fabrication exercise. Thus the 8L45 appeared to be the best selection available.

Before finalizing this decision, however, the 8L45 and 6L45 were compared by simulating each in the team vehicle model. A shift strategy was developed for each option that shared a common principle, minimizing fuel consumption at any given RPM and throttle position combination up to 100% driver throttle input, at which point maximum torque to the wheels would be prioritized over fuel economy. This was done by iterating through the speed and throttle combinations and calculating engine efficiency for each gear, then selecting the most gear providing the most efficient operation as the selection for that operating point. It should be noted that this would give the vehicle a very unsatisfying and sometimes unpredictable drivability were it to be implemented this simplistically in real life- as the driver further depressed the accelerator pedal, torque would build very slowly until a sudden jump near the floor. But this strategy enabled direct comparison of the two transmissions on the basis of their potential for fuel savings, which was the topic of interest. The shift maps developed are shown in Figure 12 and Figure 13 below.
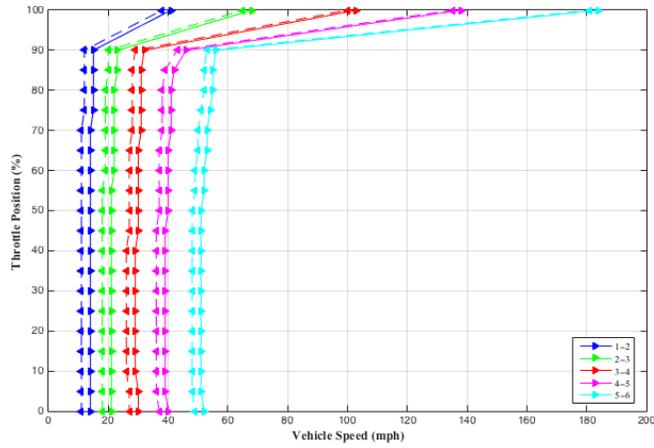
Figure 12: Optimized 6L45 Shift Map
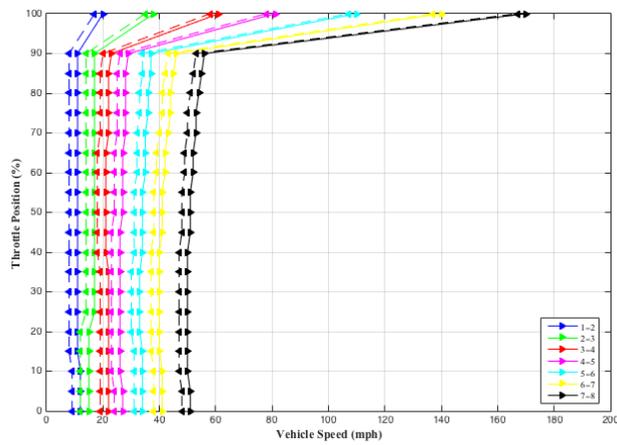
Figure 13: Optimized 8L45 Shift Map

The results of this simulation are shown in Table 17 below.

Table 17: Modeling Results for Transmission Comparison

	6L45	8L45
Energy Consumption (Wh/km)	585.3	557.7
Fuel Economy (mpgge)	35.8	37.5
Projected Range (miles)	264	277

With a decrease in fuel economy and a corresponding increase in vehicle range, the 8L45 was selected over the 6L45 as the team vehicle's transmission. Figure 14 below shows the coupling of the 8L45 and the LEA, with the adapters required for such, prior to installation in the Camaro.
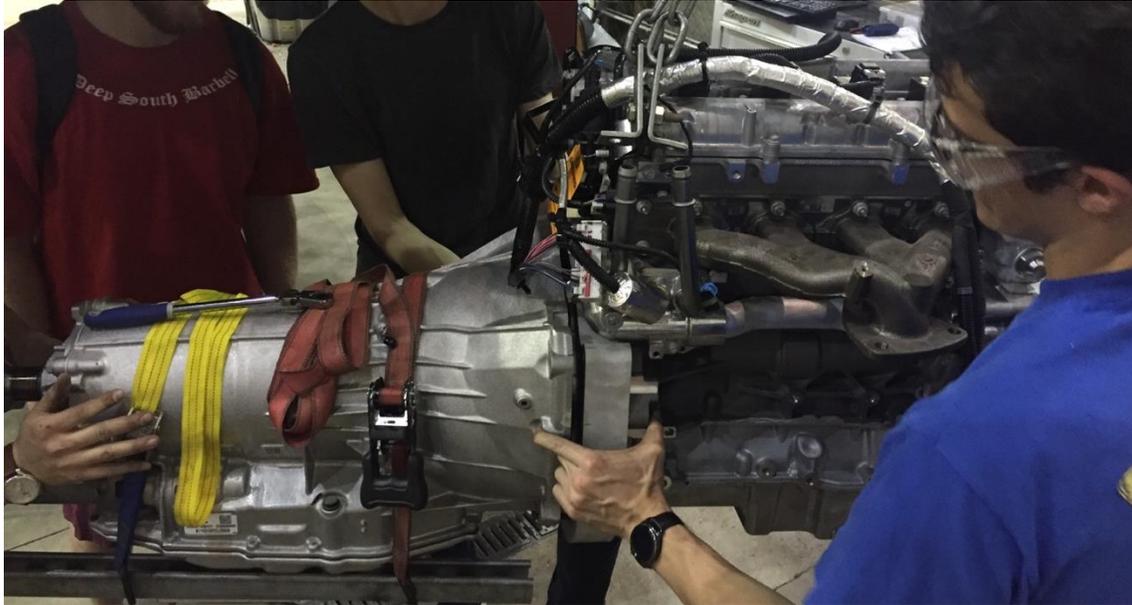
Figure 14: Mating the 8L45 and the LEA

Non-Powertrain System Engineering and Design

Finite Element Analysis

Finite Element Analysis (FEA) was used extensively to predict the strength of team components, so some general notes that apply to multiple non-powertrain components will be mentioned before discussing the design of the battery pack, motor and transfer case mounting, fuel systems, and more. Siemens NX 9.0 was used for all FEA.

Competition regulations [4] require the following load cases to be considered for all major mounting structures:

- 20G longitudinal (representing a frontal impact)

- 20G lateral (representing a side impact)

- 8G vertical (representing a large bump load)

Competition regulations also require that all critical parts be designed to a safety factor of at least 1.5 for these load cases.

The simulation of bolted connections in FEA is a well-studied question, and there is no one way that will be both accurate and computationally efficient for every situation. To illustrate this point, of course the most accurate way to simulate such a connection would be to include a 3D model of the bolt, nut, and washers (including threads) and apply the correct preload, friction, and contact elements to them and the components they are fastening. This method, however, would greatly increase computational time because of the nonlinearities in the contact elements and the small elements required for the threads. Numerous alternative options have been developed, and analysis of bolted joints took two forms for the analysis presented in this paper.

When numerous bolts were to be used for a single component, an example being the 28 bolts attaching the battery modules to their structure, modeling was simplified by applying the necessary forces, constraints, or connections to a point in the center of each hole that was then connected to the side walls of the bolt holes by a Rigid Body Element (RBE, or specifically RBE2 in NX). This, of course, is a location where no forces are directly applied in an actual bolted joint, as all forces from a bolt are applied via the frictional interface between the washer and member being bolted, not through bearing loads from the bolt shaft. However, the value in this approach lies in the ease with which it can be applied to a large number of bolt holes, something that is especially valuable early in the design cycle when frequent changes are made in response to FEA results. An illustration of this method from a motor/transfer case mounting plate is shown in Figure 15 below.
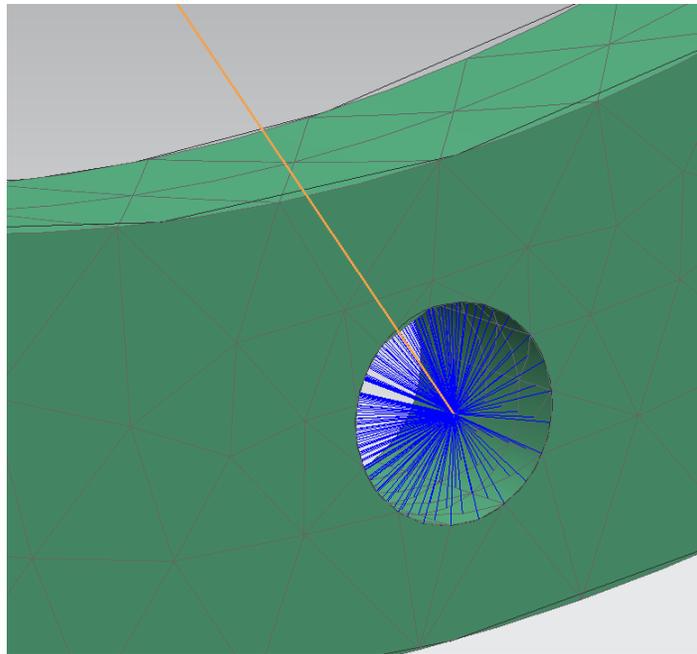
Figure 15: First Method for Bolted Connections

The second method is more accurate and was generally utilized when fewer bolted joints were present to carry a given load. The faces of the members being bolted were split by two circles concentric to each bolt hole to represent the washer footprint. The forces, constraints, and connections were then applied to the washer footprint faces via a central point, as in the first method. This method is favorable for being more accurate to the interfaces carrying loads in real-life bolted joints but requires more setup time for each hole individually. An illustration of this method from the front motor/transfer case mount is shown in Figure 16 below.
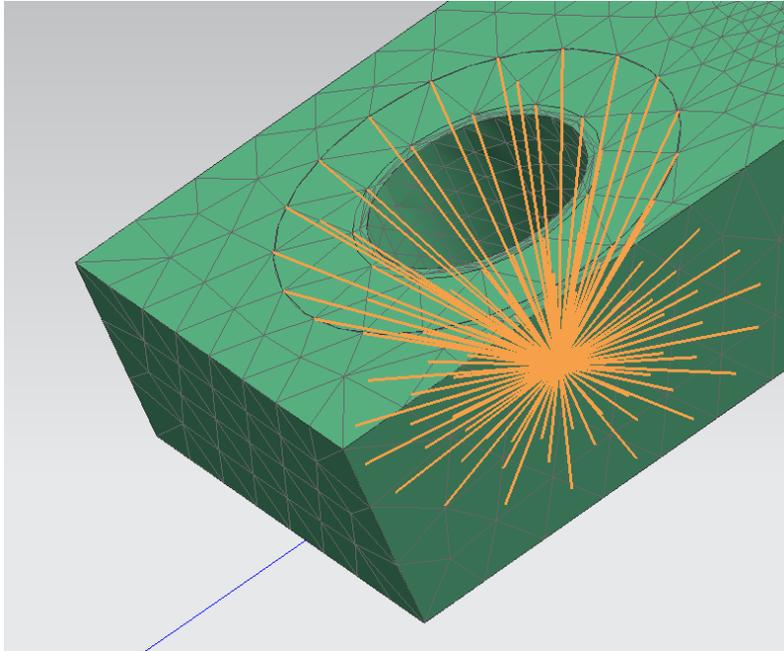
Figure 16: Second Method for Bolted Connections

For each of these methods, all connections used in the analysis contained in this section were RBE's. This is a conservative method that neglects the way that various components represented by these connections (including bolts and, as will be explained, the motor and transfer case themselves) take some of the loading in a given scenario, as RBE2's are defined to be infinitely stiff.

Battery Pack

As mentioned, packaging the 7 modules of the team's A123 battery pack proved impossible outside of the trunk. Thus the battery pack was designed to efficiently package not only the modules but the Battery Control Module (BCM), Electronic Distribution Module (EDM), Current Sensing Module (CSM), and associated structural and cooling components without unduly filling up the entire trunk and thus taking away all luggage space.

Figure 17: BCM (grey), EDM (orange), and CSM (orange) from left to right

First, the three modules just mentioned, each shown in Figure 17 above, should be described.

The BCM is fairly straightforward- this control module connects to the seven battery modules

and measures each of their cells' voltages and temperatures, thus sensing whether or not the

battery pack is in a state that allows for safe operation and taking action if it isn't. For example, if

cells within a battery module have differing enough voltages (which will happen to some level

due to manufacturing imperfections), the BCM will activate cell balancing within a module, and if

cell temperature is observed to be too high, the BCM will decrease allowable battery current to

diminish the corresponding heat generation.

The CSM consists of a small resistor through which battery current flows and across which

voltage drop is measured, providing an accurate way of measuring total pack current. The EDM,

finally, consists of contactors, which are fairly large mechanical switches that can complete or

break an electrical circuit. These contactors are closed, thus completing the circuit and allowing

current to flow from the battery to other components in the vehicle, only when the BCM senses

that all conditions allow safe operation and when a CAN message is received from the hybrid supervisory controller that commands these contactors to close.

There are numerous ways to successfully lay out the modules in the trunk of a Camaro (for reference, each of the modules is roughly the size of a standard automotive 12 volt battery). The team's design goal was to arrange and package them in as simple a box as possible and to make the construction of this box as simple a task as possible, goals driven by the compressed timeline of the competition and the numerous other demands on the team to build out and package the components of a hybrid electric Camaro within two years.

Figure 18: Small Camaro trunk opening

Another factor relating to module arrangement was available trunk openings. Because the opening of the Camaro's trunk, shown in Figure 18 above, is far too small to insert the entire pack at once, the pack was designed to be modularly constructed outside of the vehicle and

then inserted in pieces. Thus the modules are divided into two groupings, a 2x2 grouping of four modules and a separate grouping of three.

Each of these groupings was given a 1/16" thin steel plate (referred to as the 4-module and 3-module plates) to rest on and to assist in transfer from the team's electrical work room to the vehicle.

The thermal design of the battery pack will be detailed later in a dedicated section, but will also be considered here from the perspective of mechanical design. Because of limited space it was decided to place a 5/8" aluminum cooling plate (aluminum for good thermal conductivity and weight savings) underneath the 3- and 4-module plates. Channels were milled through this cooling plate and copper pipes routed through them, as shown in Figure 19 below.
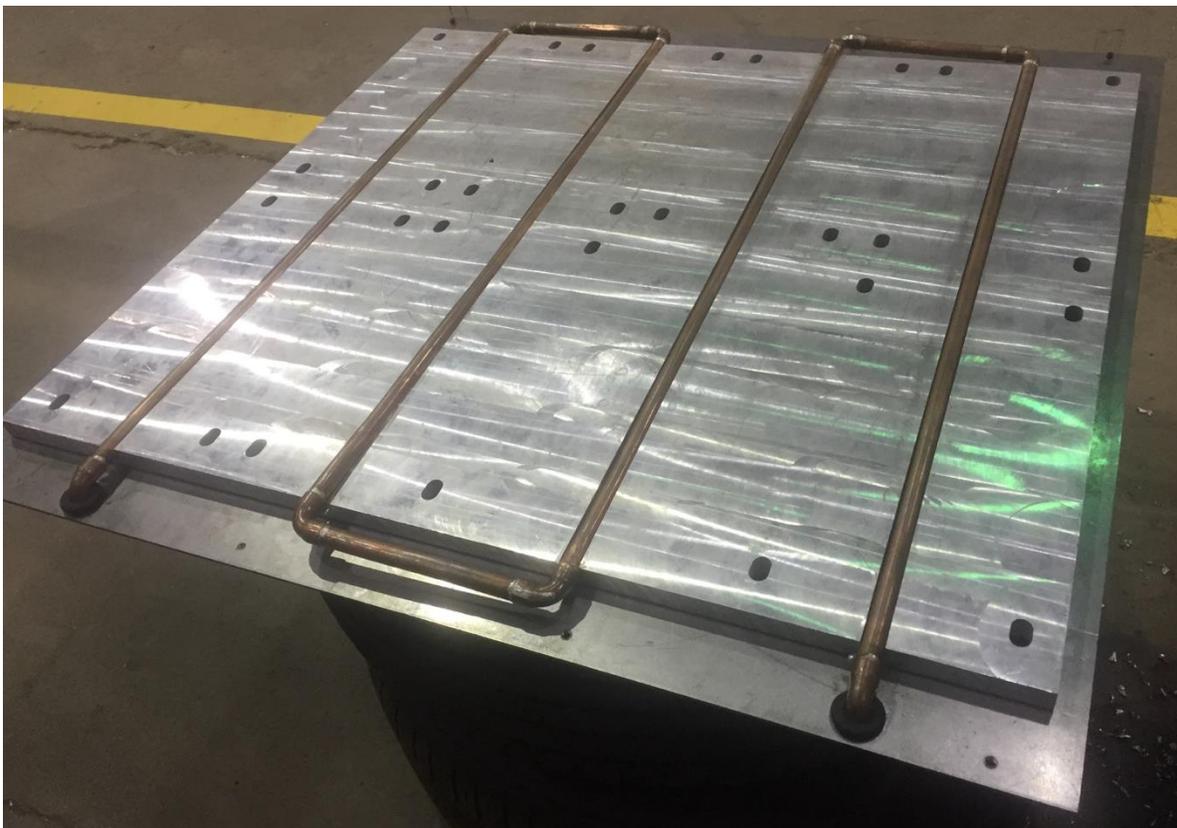
Figure 19: Cooling plate mounted on steel baseplate

Finally, this all rests on a 1/8" thick steel baseplate that provides structural support and a common connection point for the aluminum cooling plate, acrylic battery box structure (to be discussed), and battery pack mounting straps. These straps, fashioned from ¼" thick steel bars, run across the trunk frame rails as shown in Figure 20 below; note that the rear portion of the trunk floor is cut out in this image.

Figure 20: Straps across trunk frame rails

Smaller steel blocks, which will be welded to the vehicle frame towards the side of the battery pack facing the front of the vehicle, will provide additional support. A finished block is shown below in Figure 21.

Figure 21: Battery Pack Mounting Block

Per competition rules, the battery pack must be sealed to the environment so that unintentional contact with any electrically conducting elements may be prevented for any tool or fingers in the batteries' vicinity. This also prevents dangerous gases from entering any area of the vehicle in the event of battery failure. An acrylic box was thus designed to surround the modules and acrylic L-brackets were used to hold the sides together and to the steel baseplate, as shown in Figure 22 below.
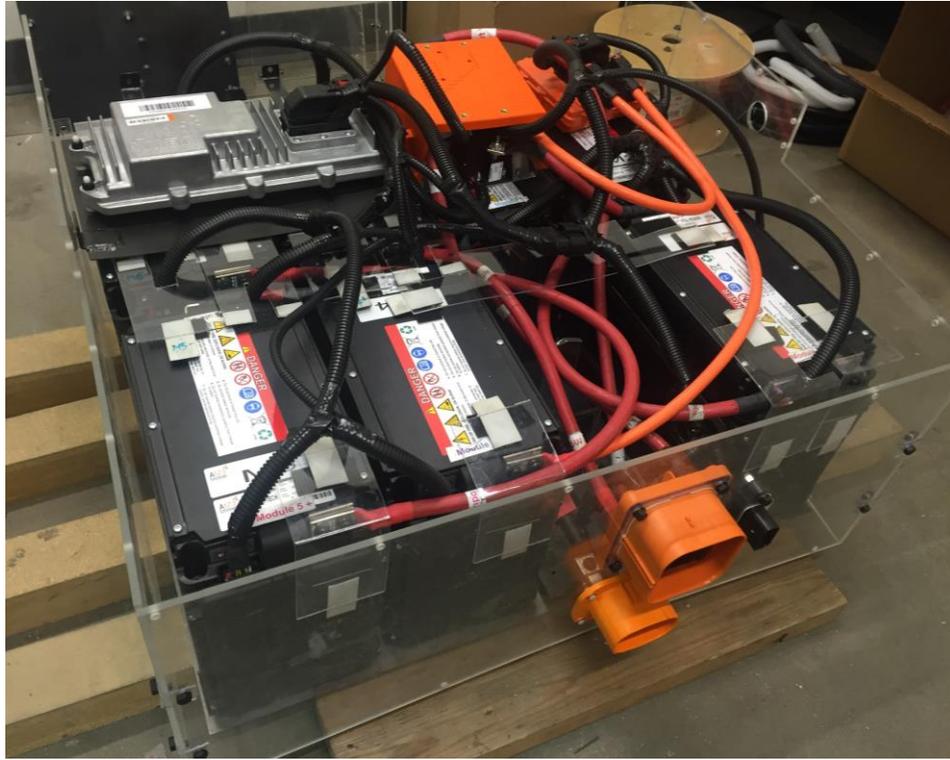
Figure 22: Battery pack partially assembled
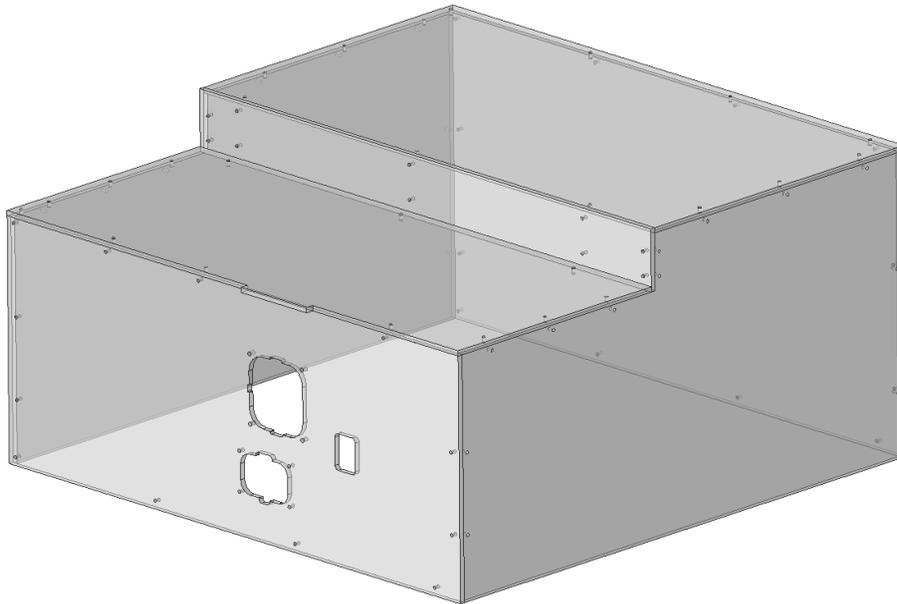
Figure 23: Acrylic Box CAD

The shelf-like construction of the battery box's top section, shown in the CAD image in Figure 23, is designed to allow the placement of the high voltage junction box on top of it, a location convenient for wire routing. The junction box acts as a sort of combined power strip and fuse box for high voltage components, giving the inverter, battery charger, AC compressor, and DC/DC converter a common and protected place to access the battery's voltage and be protected by fuses sized appropriately for each individual device.

The seams of the acrylic structure are sealed with vacuum bagging tape, or "tacky tape", a sticky tape-like product used in composites manufacturing to provide a seal during vacuum bagging.

The BCM, EDM, and CSM each require mounting within the battery pack. The most convenient placement for them is across the four modules closest to the front of the vehicle, so thin steel plates were designed to hold them. Steel L-brackets were fashioned and welded to the sides of these plates so that they could attach to the acrylic structure with plastic fasteners, which are used so that there are no metallic connections between the interior of the battery pack and the outside world.

FEA

For this simulation, forces on the battery modules acted upon the 28 bolt holes through the aluminum cooling plate and the steel base plate. These forces were applied using the simplified method mentioned earlier, onto the faces of these holes' sidewalls, and to keep the aligned holes in the two plates together, the forces were actually applied to a node in the center of each hole that was connected to these faces with RBE2's.

Two types of constraints were applied to the mounting straps and blocks. Fixed constraints were applied to the faces of the mounting blocks and the front strap which are welded to the vehicle. The rear strap, however, is bolted to the vehicle, and because only one hole is present on each

side, the more complex method of constraint application was employed. A washer footprint was defined on each of the tops of the faces and a force corresponding to the maximum preload of a M12 Class 10.9 bolt [7] applied to each. Next, a "Simply Supported" constraint was applied to the bottom faces of the straps, those that rest on the frame rails. This constraint fixes travel in the up-and-down direction while leaving the remaining five degrees of freedom free. Finally, a constraint was added to the bolt hole sidewalls to fix radial movement (where "radial" references the holes' local cylindrical coordinate systems) to prevent longitudinal or lateral slipping.

This simulation setup is shown in Figure 24 below, with results following.
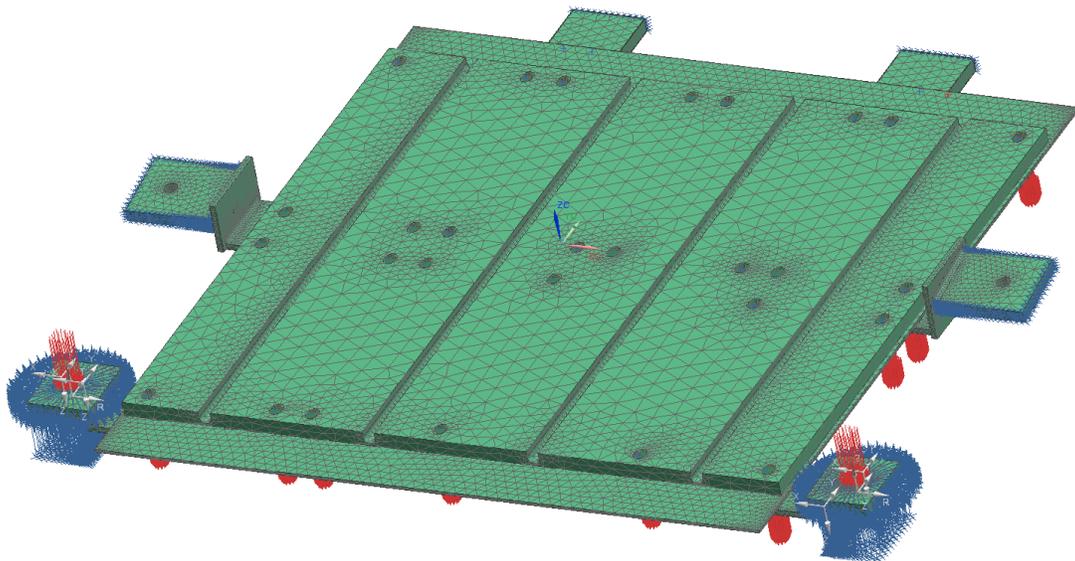
Figure 24: Battery Mounting FEA

8G Vertical
Stress - Elemental, Von-Mises

247.80
220.27
192.74
165.21
137.68
110.16
82.63
55.10
27.57
0.04
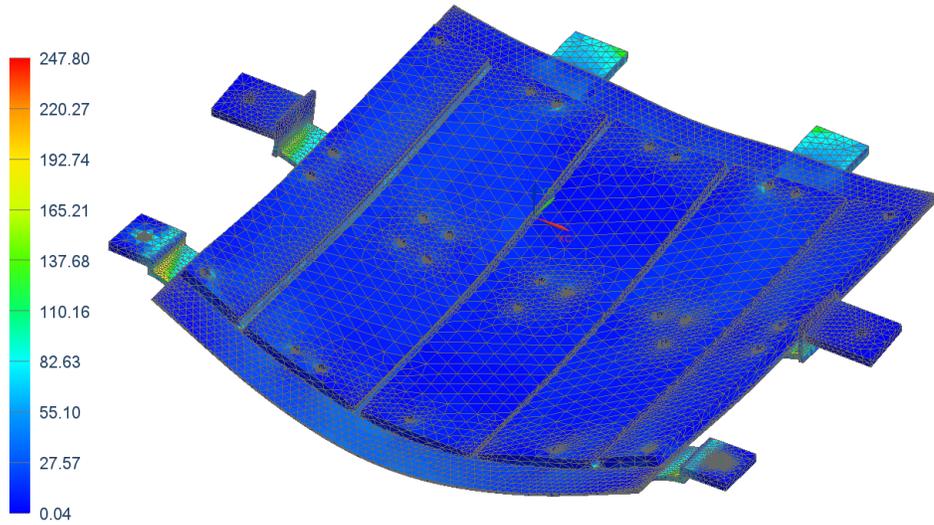
Units = N/mm^2(MPa)

Figure 25: Battery Structure Vertical Loading

20G Lateral
Stress - Elemental, Von-Mises

150.40
133.69
116.98
100.27
83.56
66.85
50.14
33.43
16.71
0.00
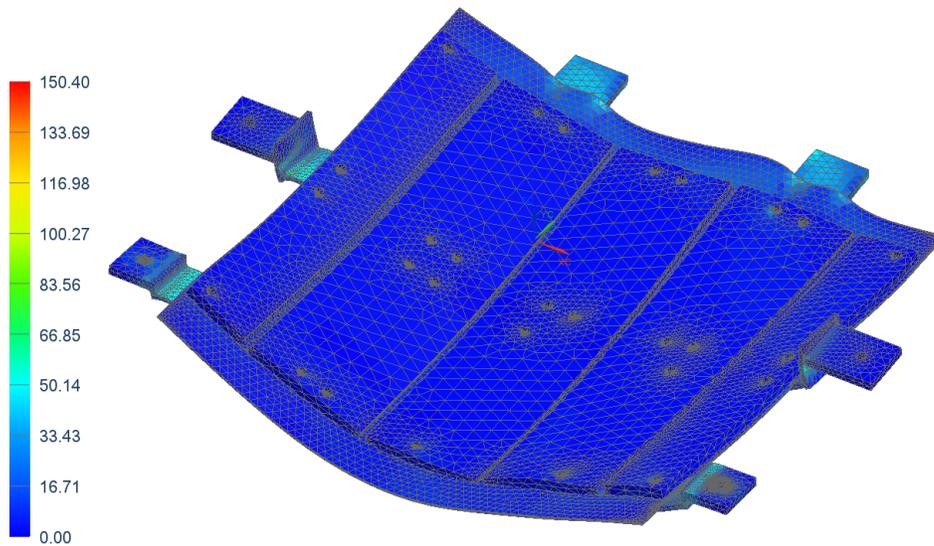
Units = N/mm^2(MPa)

Figure 26: Battery Structure Lateral Loading

54

20G Longitudinal
Stress - Elemental, Von-Mises

118.21
105.07
91.94
78.81
65.67
52.54
39.40
26.27
13.14
0.00

Units = N/mm^2(MPa)
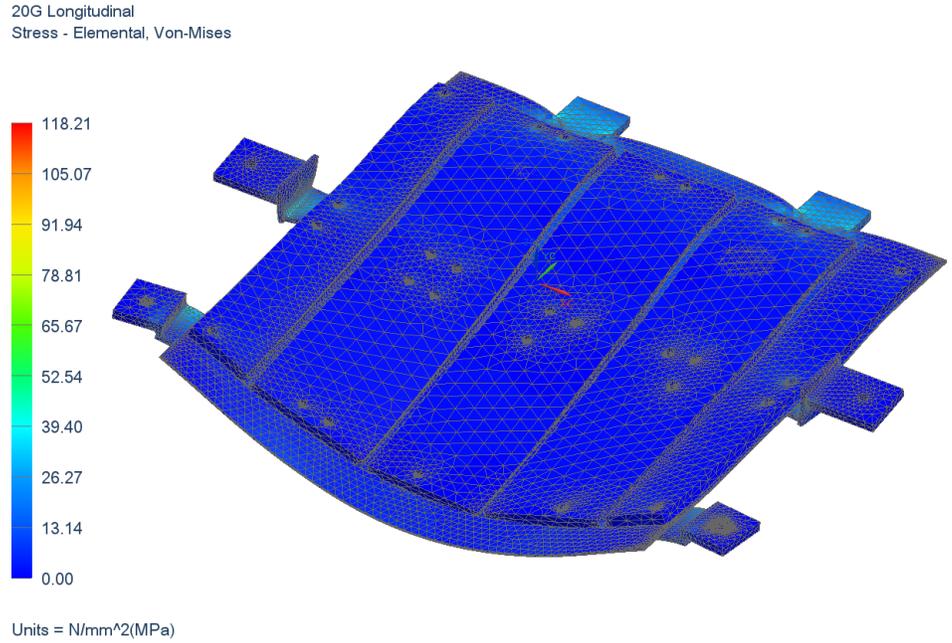
Figure 27: Battery Structure Longitudinal Loading

The maximum stress in each of these three load cases was found in the straps that span the frame rails, which are made from 1018 steel. Safety factors below are based on 1018 steel's yield strength of 370 MPa [8].

Load	Max Stress (MPa)	Safety Factor
Vertical	247.80	1.5
Lateral	150.40	2.5
Longitudinal	118.21	3.1

Lessons Learned

The design and construction of the battery pack presented possibly one of the greatest challenges to the team throughout the first two years of EcoCAR 3, largely because of lack of

experience with the design of battery packs. At the end of the process during the time of this writing, some reflection gives a few lessons learned for any future battery pack design.

First, while previous experience with powertrain selection had taught the overwhelming importance of packaging in the design process, in this case it would have been wiser to start with the goal of simplifying wire routing rather than packaging. Once the placement of the modules was set, untold hours were spent designing a wire routing scheme that would satisfy manufacturer-specified minimum bend radii and competition regulations related to safety and best practices. If the modules could have avoided wiring altogether, such as through the use of copper bus bars, or the modules could have been oriented in such a way as to reduce wire lengths and bends, the overall design of the pack could have been greatly simplified and its construction made both quicker and safer.

A related issue discovered during pack build was that the way the high voltage cables bolted to the modules in some locations blocked both the low voltage harness and the module mounting bolts from being inserted. While not directly a design flaw of the team per se, as it is the manufacturer who designed it that way, noticing and mitigating that issue would have prevented numerous days' labor and allowed for much more timely completion of pack integration.

Finally, while the thermal analysis to be presented makes a case for effective battery pack cooling even in very warm-weather situations, the team was unaware that the sides of the modules were designed for contact with cooling plates. The bottom surfaces of the modules present less surface area and a higher thermal resistance than the sides, but were chosen as the cooling surface for packaging reasons. Mounting the modules on their sides would have yielded more effective cooling.

These lessons in mind, a newly-designed pack would be lighter, more compact, easier to assemble, and more efficiently cooled.

Motor and Transfer Case Mounting

Design

The Georgia Tech team emphasized active involvement from the undergraduate team members who comprise the majority of the team, and powertrain mount design was one of the tasks devolved almost entirely to them. However, the author was responsible for advising and approving the designs, ensuring that they were able to fabricated, and conducting structural analysis on them. The most interesting work done in this regard concerned the FEA of the motor and transfer case mounting structures.

These mounts, referred to from here on as the MTC (motor and transfer case) mounts, experience an unusual amount of loading in this vehicle because they not only support the weight of their mounted components (both their static weight and in the considered load cases) but also resist the propelling torque of the motor multiplied through the gearing of the transfer case. Their analysis was also complicated by the fact that three separate pieces were required to mount the MTC. These three pieces will be described before beginning their structural analysis.

From the front of the vehicle, the first MTC mount is a simple rectangular channel bar spanning the driveshaft tunnel, welded to its sides, and connecting to a threaded hole at the front of the transfer case. This part is shown in Figure 28 below.

Figure 28: Front MTC Mount

Because the connection of this bar to the component acting upon it is towards the middle of and far away from its constrained ends, this structure is placed in bending in a variety of loading cases. While the design of this segment previously consisted of a flat bar in the same location, the rectangular channel was selected to greatly increase its second moment of area to resist this bending while at the same time decreasing its weight.

The second MTC mount is a ½" thick plate waterjetted in a smooth curve that surrounds the opposing mating faces of the motor and transfer case and features holes through which their mounting bolts can be passed. This plate is connected to the vehicle by a trussing structure, made from square channel tubing, that is welded to a flat plate. This plate, finally, bolts into an attachment plate welded to the vehicle.

Figure 29: Middle MTC Mount

The third MTC mount acts as a C-clamp around the rear of the motor, interfacing with four of the motor's eight mounting holes available along the axial length of the motor (the other four are unused). This structure connects to the vehicle in a similar way to the middle MTC mount, bolting to a plate that is welded to the vehicle structure.
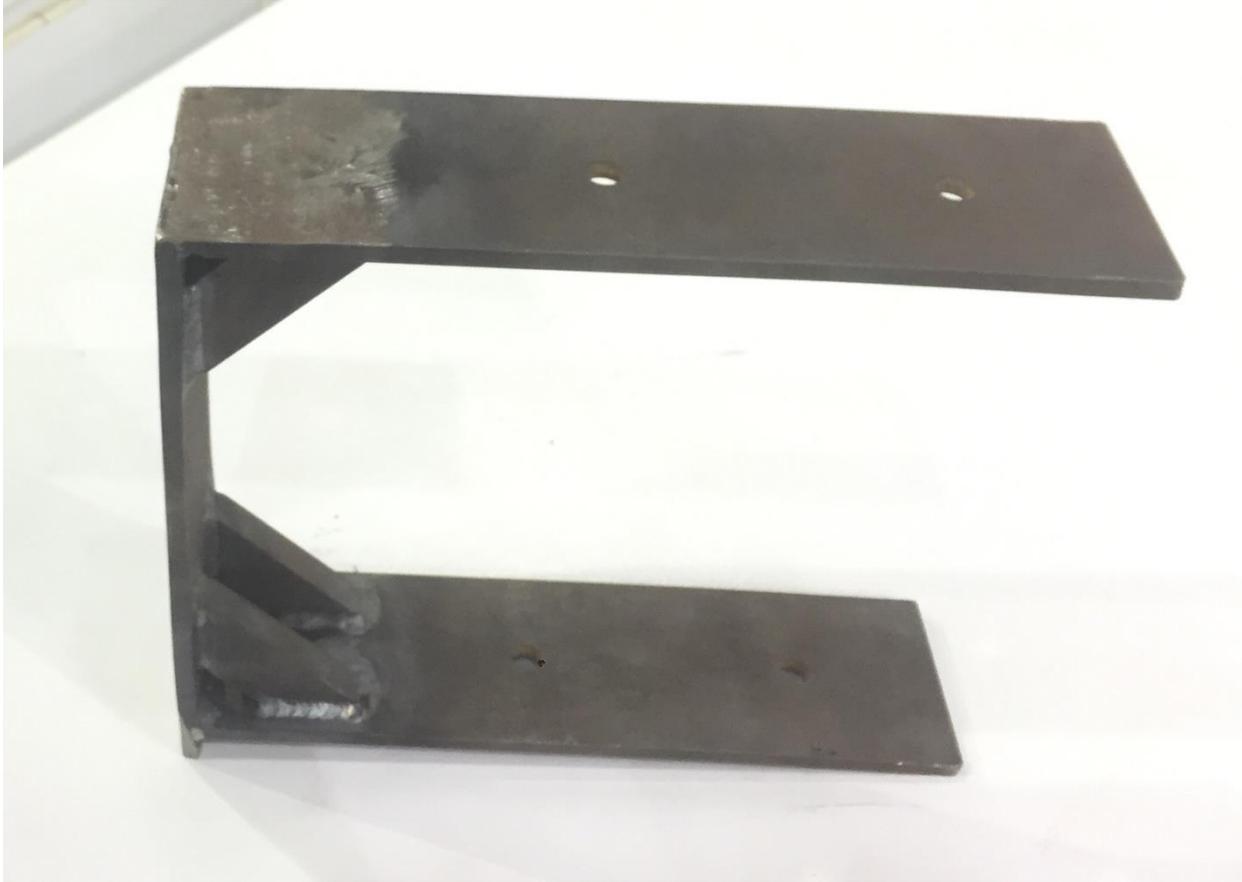
Figure 30: Rear MTC Mount

The seemingly complex strategy behind mounting these components, using plates bolted to plates then welded to the vehicle, is fairly simple upon explanation. Welding to the vehicle structure is difficult due to the thinness of its sheet metal, an issue inherent to modern unibody vehicles. Welding is also risky because it bonds these mounts to the vehicle and thus makes removing the motor and transfer case more difficult (and mistakes more costly). However, a quality and weld-distributed weld will enable forces to be better distributed on this sheet metal than would most bolting layouts. So plates are welded across a wide area of the sheet metal, then bolts connect the rest of the mounting structure to these welded plates. This enables the motor and transfer case to be easily removed in case they need to be tested off of the vehicle.

The FEA of the MTC mounting structures required a number of assumptions to correctly connect the three pieces while avoiding undue computational requirements or unknowledgeable assumptions about the strengths of the motor and transfer case themselves.

First, to avoid the latter complication, the motor and transfer case were modeled as nodes connected to the appropriate bolt holes by RBE2's. As previously discussed, this is a conservative approach since the infinitely stiff RBE2's cause more of a stress concentration than would the finitely stiff materials of the motor and transfer case.

The forces on the system were calculated as follows. First, the competition-required loads of 20G lateral acceleration, 20G longitudinal acceleration, and 8G vertical acceleration were applied to each of the component-representing nodes, with force calculated by multiplying these accelerations by the components' respective masses. The static weight of the two components was similarly applied for the first two load cases (it is already included in the third).

Next, the reaction forces of the motor torque were added to each of these scenarios to simulate the worst-case scenario of the rapid decelerations occurring simultaneously with peak torque production. The color-coded free body diagram in Figure 31 below shows the forces on the transfer case's chain and sprockets.
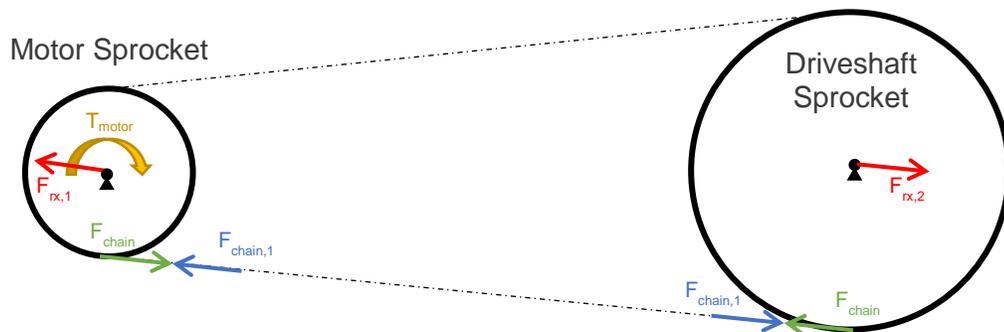
Figure 31: Transfer Case Sprocket Free Body Diagram

The torque of the motor, T_{motor} (yellow), pulls on the chain and puts it in tension (blue) with the corresponding reaction forces on both sprockets (green). These reaction forces then cause bearing reactions in each of the sprockets. To calculate these forces, first the radii of the two sprockets are defined as r_1 and r_2. Then,

$$F_{chain} = \frac{T_{motor}}{r_1}$$

$$F_{rx,1} = F_{rx,2} = F_{chain}$$

The driveshaft sprocket reaction force was applied to the transfer case node, as it is resolved through the MTC mount via the transfer case that holds it. The motor sprocket reaction force, on the other hand, was applied to the motor node, as the side of the transfer case where this force acts happens to not be bolted to the MTC mounts, while the motor that plugs into this sprocket is mounted in numerous points around it. Thus the load path for this reactive force, besides compressing the transfer case itself, an effect outside of this analysis, is conservatively placed entirely through the motor mounting points. These loads in the 8G vertical loading case are shown in the NX FEA environment (or FEM, as it is called in NX) in Figure 32 below.
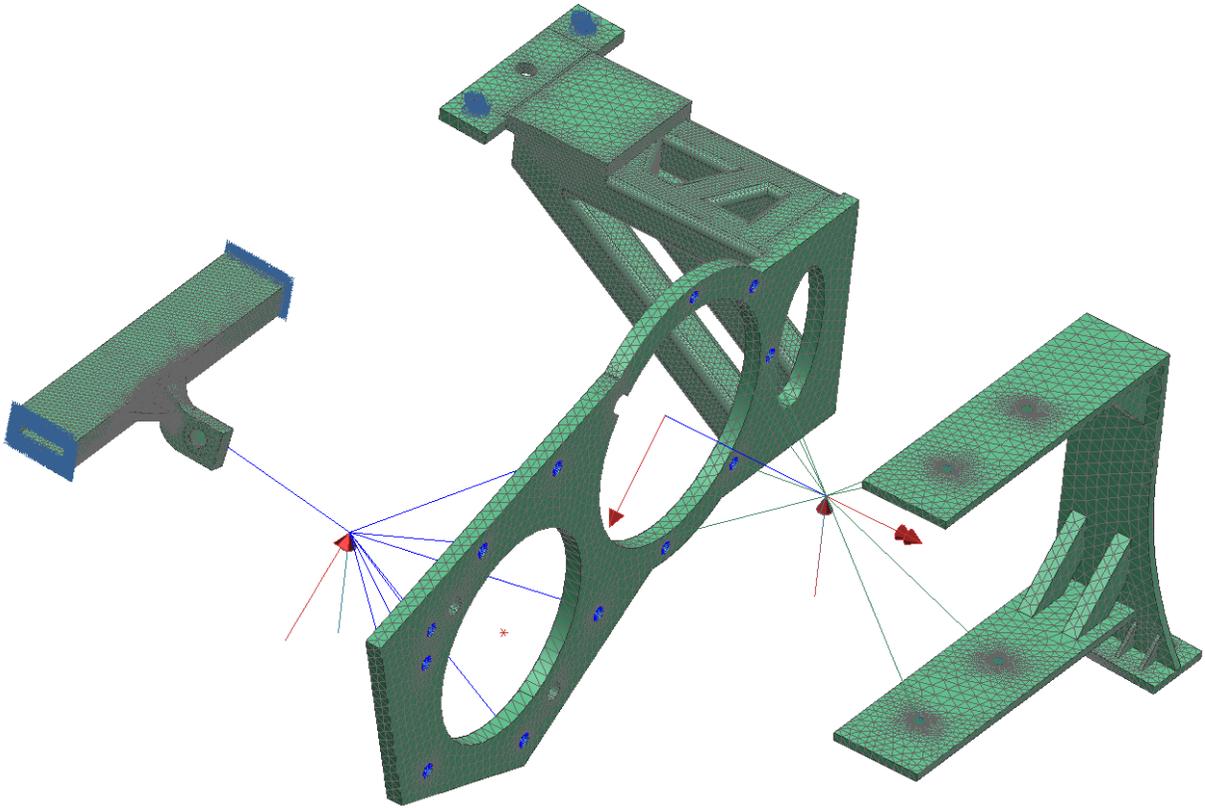
Figure 32: Motor and Transfer Case Mount FEA

Note the upward (towards the top of the page) forces, representing the vertical acceleration, and the angled forces, which represent the reactions just calculated. The barbed arrow pointing generally towards the right of the page is the torque of the motor.

Blue lines on the left represent the connections from the transfer case node to its mounting holes, and the green lines on the right represent the same for the motor. Finally, the blue clusters visible on the front MTC mount and the top of the middle MTC show the fixed constraints applied to the portions of the mounts that are firmly attached to the vehicle (the fixed constraints on the rear MTC mount are hidden in this view).

The results of this FEA simulation are shown in Figure 33, Figure 34, and Figure 35 below, and summarized in The maximum stress in each of these three load cases was found in parts which

are made from A514 steel. Safety factors below are based on A514 steel's yield strength of 690 MPa .

Table 18.
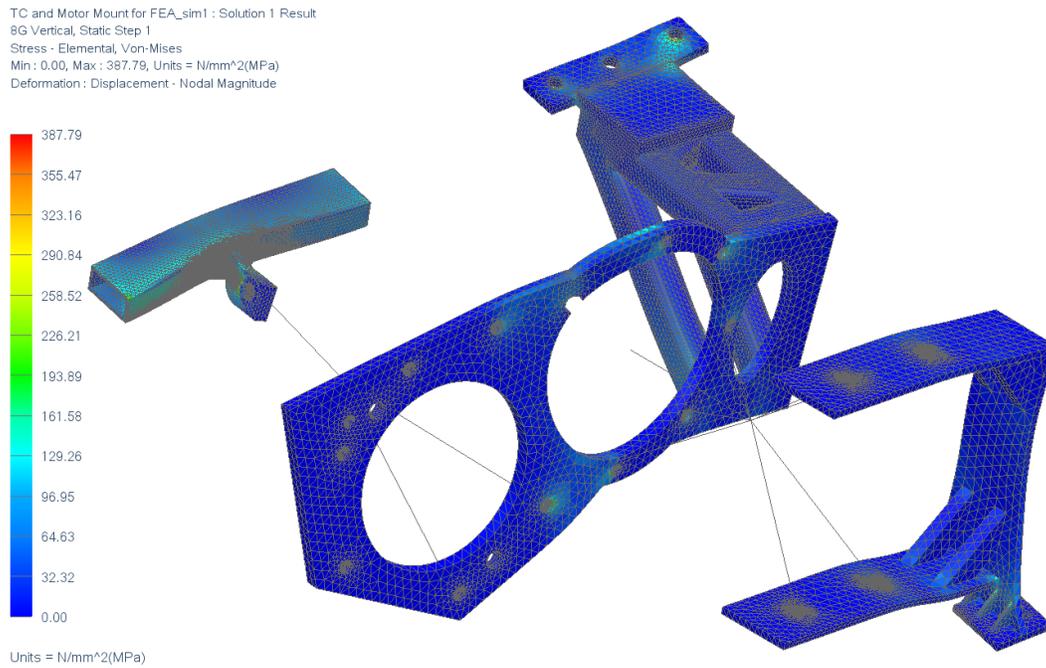
Figure 33: MTC Mount Vertical Results

TC and Motor Mount for FEA_sim1 : Solution 1 Result
20G Longitudinal, Static Step 1
Stress - Elemental, Von-Mises
Min : 0.00, Max : 375.16, Units = N/mm^2(MPa)
Deformation : Displacement - Nodal Magnitude

375.16
343.90
312.64
281.37
250.11
218.84
187.58
156.32
125.05
93.79
62.53
31.26
0.00
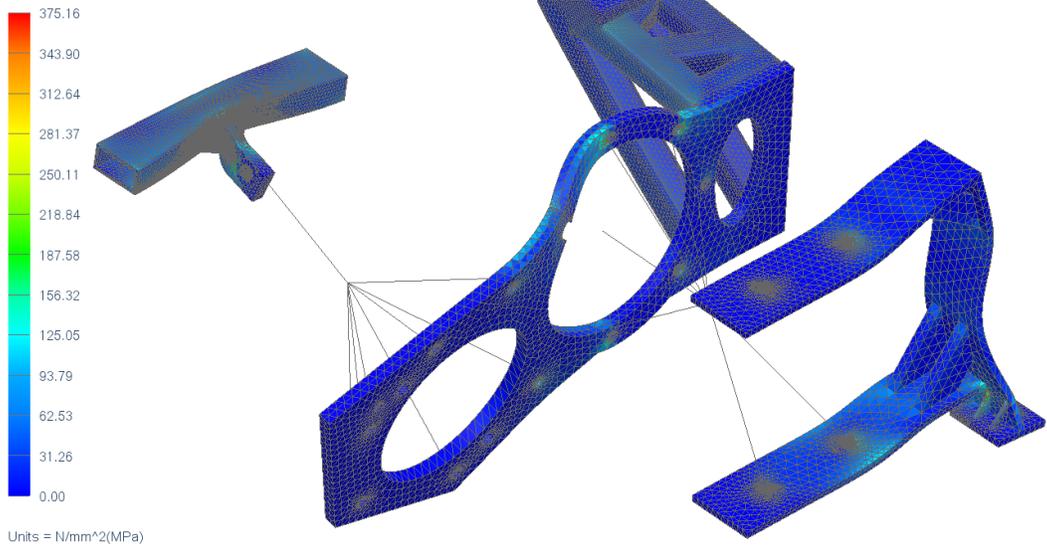
Units = N/mm^2(MPa)

Figure 34: MTC Mount Longitudinal Results

TC and Motor Mount for FEA_sim1 : Solution 1 Result
20G Lateral, Static Step 1
Stress - Elemental, Von-Mises
Min : 0.00, Max : 504.99, Units = N/mm^2(MPa)
Deformation : Displacement - Nodal Magnitude

504.99
462.91
420.83
378.75
336.66
294.58
252.50
210.41
168.33
126.25
84.17
42.08
0.00

Units = N/mm^2(MPa)
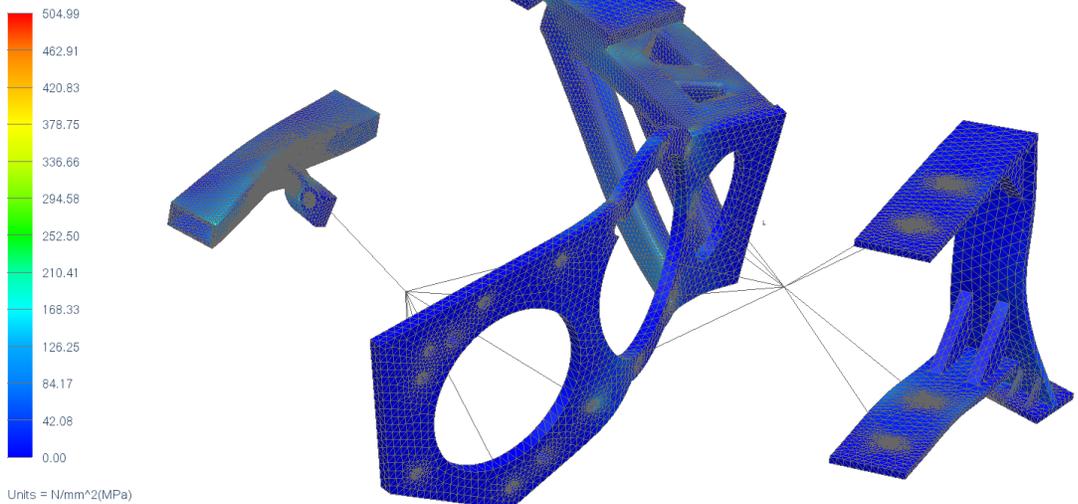
Figure 35: MTC Mount Lateral Results

65

The maximum stress in each of these three load cases was found in parts which are made from

A514 steel. Safety factors below are based on A514 steel's yield strength of 690 MPa [9].

Table 18: MTC Mount FEA Results

Load	Max Stress (MPa)	Safety Factor
Vertical	387.79	1.78
Longitudinal	375.16	1.84
Lateral	504.99	1.37

Note that the lateral load case has a safety factor under the competition requirement of 1.5.

While the value is still above 1, showing that this part will not fail, bracing will be added in the

requisite location to diminish the bending stresses and meet the competition requirement.

Fuel Systems

Competition rules state that the stock fuel tank of the vehicle must be removed and replaced with a smaller one, with the note that this requirement is mainly in effect to ease the process of tank removal and reinstallation at competitions during fuel weigh-ins. As well, the tank must be made of steel or aluminum (as opposed to plastic) to keep the tank's mounting points from fatiguing with repeated removals and reinstallations.

Fuel tank design for this competition must keep a few things in mind. First, unless a manufacturing partner is found, the design of the tank must be simple enough to be able to be constructed by students without unnecessary and thus time-wasting complexity. Second, it must have a large enough capacity to allow the vehicle to travel roughly 150 miles, the distance specified by the competition organizers as the likely distance for the emissions and energy consumption test in the Years 3 and 4 competitions. Finally, the tank requires the following interfaces:

- Fuel filler neck inlet (for refilling the tank)
- Fuel pump, which includes the following ports:
    - Fuel send line (the LEA engine is returnless)
    - Evaporative emissions to the engine and evaporative emissions canister
    - Electrical harnesses
- Refilling bung for out-of-vehicle refilling
- Venting bung to be opened during out-of-vehicle refilling
- Grounding point

While the location of one half of the stock fuel tank is occupied by the motor and transfer case, the space on the opposite side of the driveshaft is free and thus a perfect place for the new fuel tank. A simple shape, shown in Figure 36 below, was designed to fit between a chassis brace

piece on one side and exhaust routing on the other, and features a partially slanted top surface

to fit around the bottom of the rear passenger seats located directly above it.
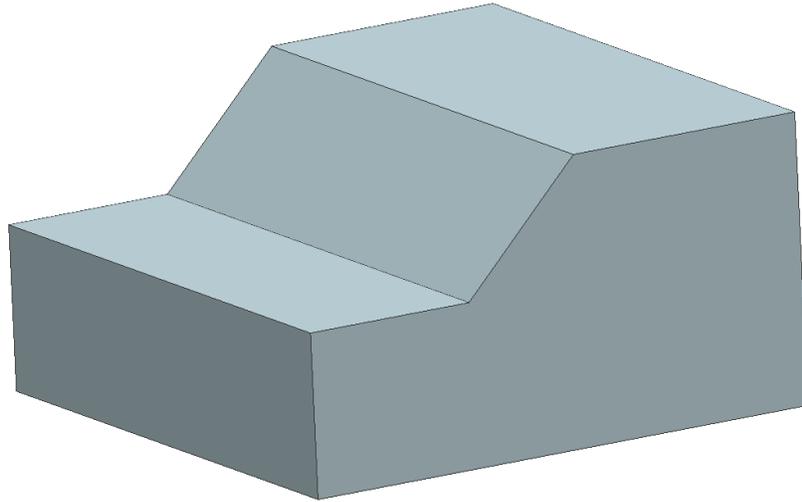
Figure 36: Fuel Tank Shape

While the interfaces mentioned above as part of the fuel pump could technically be integrated

separately, a number of reasons exist to use a General Motors fuel pump with these packaged

together. Beyond design simplicity, because this engine operates a return-less fuel system, the

engine is in constant communication with the fuel pump located in the tank about how much

pressure it needs to provide to the higher-pressure pump that is located on the engine itself. If a

3rd-party pump were to be used, this communication would be difficult at best to provide and

validate without over- or under-supplying the high pressure fuel pump. The fuel pump used in

the Buick Regal, the most common vehicle in which the LEA engine is found, is validated for

use with E85 and could be sourced through GM donation. It features an integrated evaporative

emissions solenoid and connections to the lines running to the evaporative emissions canister in

the rear of the vehicle and to the engine in the front. It is also small enough to fit in the team's

fairly short tank, unlike the fuel pump from the Chevrolet Equinox (another vehicle that features

the LEA) which was initially selected but later abandoned due to its size.

Figure 37: Buick Regal Fuel Pump

Figure 38: Chevrolet Equinox Fuel Pump

To simplify design and integration, the stock hard fuel line on the Camaro was utilized as a connection point for a short section of flexible tubing from the pump outlet. These hard lines run to the engine bay, and another soft line section was then added to connect them to the LEA's fuel inlet port.

Because of the added weight of the battery pack in the trunk, the weight distribution of the vehicle is heavily skewed rearwards. Competition regulations state that the vehicle may have no more than 57% of its total weight on the rear axles, so it is desirable to mitigate the batteries' effect by moving things forward and placing as few components as possible in the trunk. However, there are a few items which greatly simplify wire routing by being closer to the battery pack, including the junction box, the battery charger, and the inverter.

The battery charger, as previously described, is a rectifier and voltage transformer that receives AC wall electricity and converts it to the correct DC voltage to charge the battery. It contains CAN-connected electronics that allow for it to only charge when allowed by the supervisory controller and battery pack and also provide voltage and current as desired by the BCM for optimal pack health.

The inverter has the opposite effect on electrical energy, converting DC voltage from the battery into 3-phase AC voltage provided to the motor's three phase windings. The inverter acts as a "motor controller module," receiving a torque command from the supervisory controller and providing the pulse width modulated (PWM) voltage that will cause the stator current that will make the requested torque.

Because the battery charger and inverter each connect to the junction box that is located on top of the battery pack, placing them in the trunk makes sense to shorten the heavy and bending-resistant high voltage wiring needed in the vehicle. A location for them was found underneath the battery pack on the "shelf" of the trunk which, in most similar vehicles, would contain a spare tire (the 2016 Camaro, incidentally, does not include a spare tire but rather comes with a tire hole filling kit and air pump in this location).

Engine Bay

Eight types of routing are required for an engine to work and for all stock functionality to be provided. They are:

- Electric harness

- Air intake

- Engine exhaust

- Engine coolant

- Air conditioning refrigerant lines

- Vacuum lines

- Fuel lines

- Evaporative emissions lines

Providing these connections was a substantial effort requiring creative routing and full knowledge of each of their functionalities. They will be described here individually.

Electric Harness

When the LEA engine was delivered, it included the harness used with it in the Chevrolet Equinox. This harness connects various sensors and actuators on the engine to the engine control module (ECM), transmission control module (TCM), and the vehicle CAN bus.

Air Intake

The air intake requires a non-tortuous path for air from the front of the vehicle to the throttle body, some form of air filter, a mass airflow (MAF) sensor, and a positive crankcase ventilation (PCV) inlet. The team vehicle will use the intake system from the 2015 Buick Regal, another vehicle powered by the LEA (which almost guarantees compatibility with its MAF sensor and its location) with a cone filter at the front of the vehicle.

Engine Exhaust

The stock exhaust routing for the Camaro's V6 LGX engine, shown in Figure 39 below, consists of two lines, mirrored about the driveshaft axis, extending from both the driver and passenger sides of the engine to the 4-tip muffler in the rear.
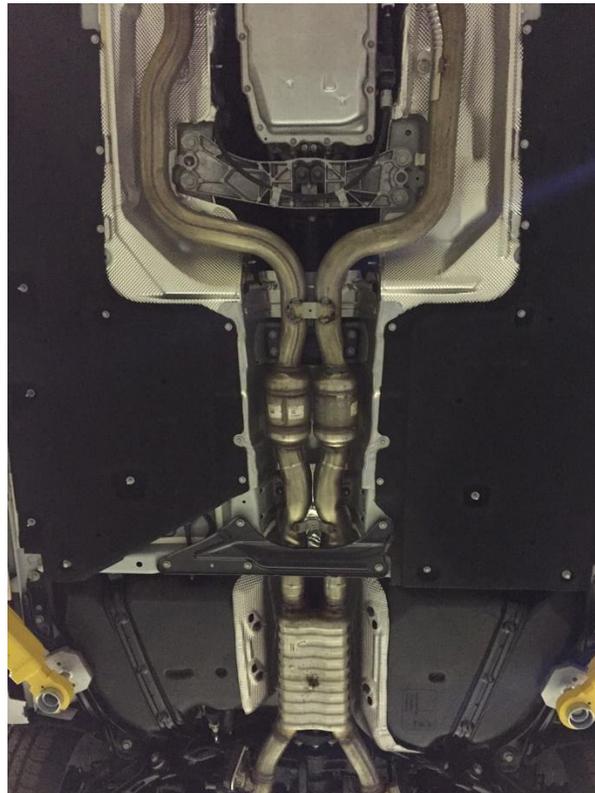
Figure 39: Stock Exhaust Routing

It was desired to use as much of the stock routing as possible to simplify design and installation, but the location of the LEA's exhaust manifold is far from the stock exhaust. Thus a custom section of exhaust tubing needed to be built to extend from the exhaust manifold (taken from the Saturn Sky, and shown in Figure 40) to the stock routing below and behind it. The stock routing is then retained until the first catalytic converter, as shown in Figure 41. After this, custom

routing will take the exhaust around the new fuel tank and to a single-tip muffler in the rear that is much smaller than the stock 4-tip muffler.
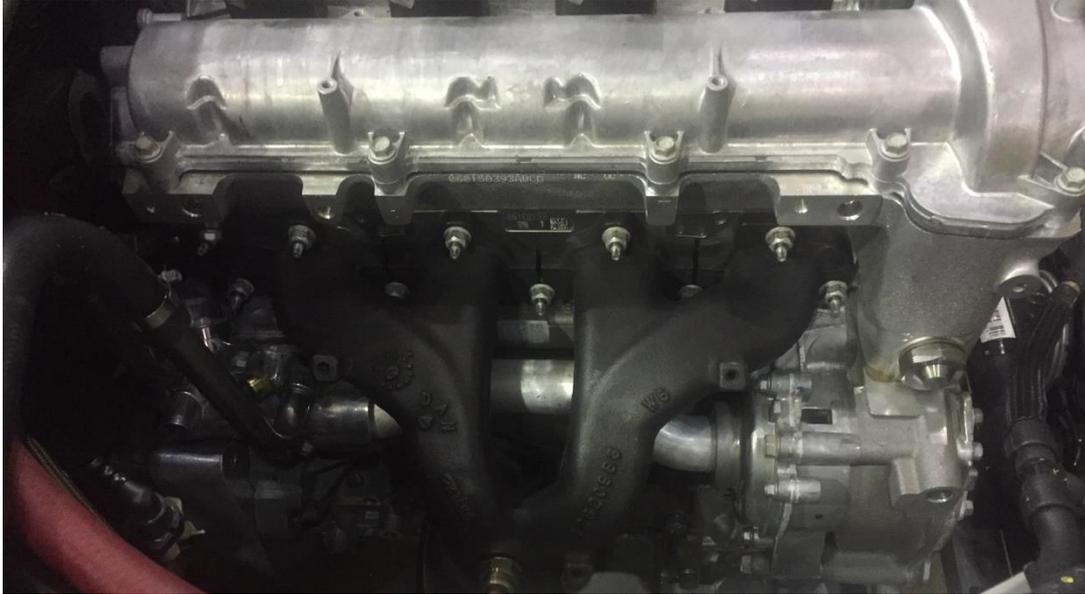
Figure 40: Saturn Sky exhaust manifold
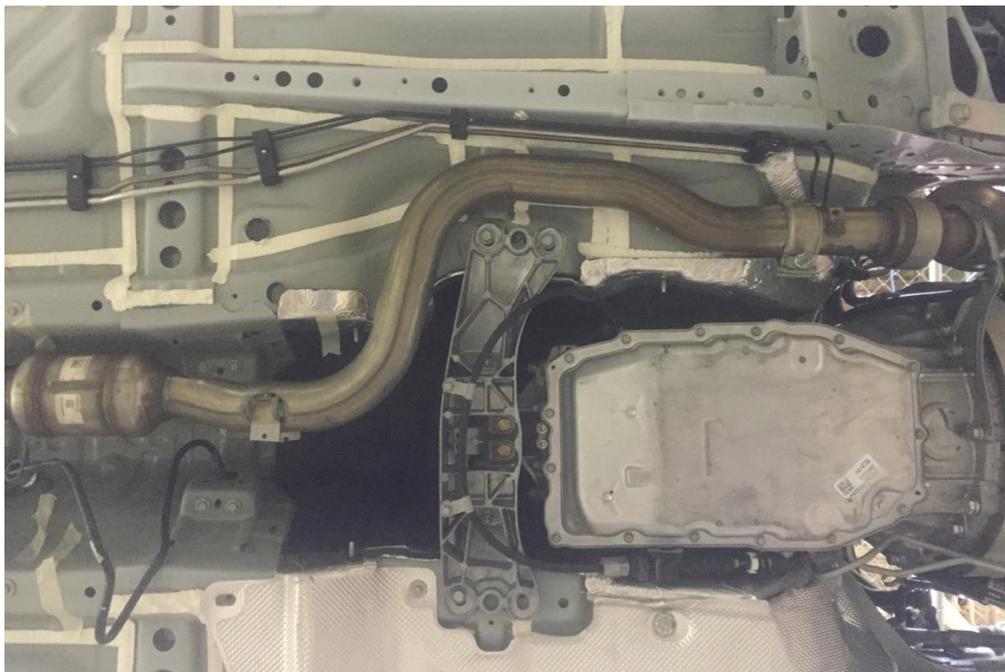
Figure 41: Retained length of stock exhaust routing

Air Conditioning Refrigerant Lines

As previously mentioned, DENSO, a supplier of automotive parts to numerous automakers, graciously donated both the electric AC compressors and DC/DC converters used by many EcoCAR 3 teams, including the Georgia Tech team. This compressor will be placed in the front of the vehicle for short connections to the stock refrigerant lines and will replace the stock belt-driven compressor. This placement requires high voltage wiring to be run from the junction box in the trunk underneath the length of the vehicle and around the back wall of the engine bay, but because this compressor requires much less current than the motor and inverter, its cabling is of a much smaller gauge, thus weighing less and having a much lower minimum bend radius than the thicker cables used for those two components. Extensions of the refrigerant lines will take the R134-a refrigerant through the stages of the vapor compression cycle in otherwise stock Camaro components, including condensation in the refrigerant radiator in the front of the engine bay, expansion in the stock AC expansion valve, and evaporation in the cabin heat exchanger before returning to the compression stage in the electric compressor.

Vacuum Lines

The vacuum pressure of the engine, produced by the intake stroke of its cylinders, has been used in vehicles over the years for a number of different tasks, most commonly to increase the application of braking force through the use of a brake booster. The brake booster, located on the brake master cylinder, applies the vacuum force of the engine to the brake piston when the brake pedal is pressed; its helpful pull is most missed when braking the vehicle without the engine running. In the Camaro, this routing is simple because the brake booster is the only device requiring vacuum power and its vacuum port almost directly faces the engine's vacuum port, requiring only a short run of tubing.

Fuel and Evaporative Emissions Lines

These two types of lines are included together because both their functionality in any vehicle and their routing in the stock Camaro are very similar, with the two lines running alongside each other for most of the vehicle. The fuel line, of course, carries fuel to the engine for combustion. The evaporative emissions system, or evap system for short, traps the fuel that has evaporated in the tank in a charcoal evap canister. It also simultaneously provides air to the tank to replace the volume of fuel that has been consumed by the engine, preventing a backwards vacuum pressure on the fuel lines. When the engine has warmed up during driving, the evap canister has a valve that can be opened, allowing the fuel stored in it to be sucked into the combustion chambers of the engine and consumed. This system allows evaporated fuel, a major source of pollution before the widespread implementation of evaporative emissions control in modern vehicles, to be redirected and combusted without venting to atmosphere.

As mentioned previously, the stock evap and fuel lines are used for the majority of their length, and shorter lengths of soft tubing are then used to connect them to the inlet ports on the engine, which aren't in the same location as those on the stock Camaro engine.

Weight Distribution and Mass Tracking

Weight distribution is a key and complicated aspect of vehicle handling, and the study of its effects on cornering, braking, and acceleration will not be discussed here. However, it is clear that a vehicle that is poorly weight-balanced will drive poorly, something that will be judged in this competition, and even more importantly it is required for competing in EcoCAR 3 dynamic events (those in which the car moves under its own power, such as acceleration testing and distance driving) that the vehicle not exceed the weight requirements set by the competition organizers and GM engineers. A summary of these requirements is presented below [4]. The limits with occupants assume two passengers at 91 kg each, and both sets of limits assume all fuel and fluids filled.

Table 19: Weight and Distribution Limits Without Occupants

Front Axle Weight Distribution (minimum)	43%
Rear Axle Weight Distribution (maximum)	57%
Total Weight	1938 kg

Table 20: Weight Limits With Occupants

Front Axle	1000 kg
Rear Axle	1170 kg
Total Weight	2120 kg

Predicting vehicle mass and weight distribution during the design and installation phase of the vehicle, before all components were installed, allowed for informed design decisions to reduce weight and redistribute weight as necessary. This was accomplished using a spreadsheet of all planned component additions and removals, featured in Appendix 1. Front/Rear weight distribution was tracked in this spreadsheet by including the longitudinal position of each of these components along with the mass of the stock Camaro with its center of gravity (CG) location as measured by the team. The weight on each axle was calculated by summing the forces and moments about the front axle to zero and solving for the reaction force on the rear axle and then the front axle.

Below are the predictions for these values as the model currently stands.

Table 21: Force on wheels (curb conditions)

Front (kg)	836.82
Rear (kg)	1089.63
Front %	43.44%
Rear %	56.56%
Total (kg)	1926.45

Table 22: Force on wheels (with passengers)

Front (kg)	918.01
Rear (kg)	1190.44
Front %	43.54%
Rear %	56.46%
Total (kg)	2108.45

As shown, the team vehicle is currently predicted to have too much weight on the rear axle when loaded with two passengers. Potential weight reduction strategies include composite panels, removal of trim, lighter aftermarket parts, and re-engineering of team parts, but for the Year 2 Competition weight reduction will not be prioritized.

CHAPTER 2: HYBRID POWERTRAIN MODELING

If automakers are to design more efficient vehicles, they must be able to predict vehicle performance and energy consumption prior to full-scale vehicle production. The sooner in the development cycle that accurately predicted values can be utilized to guide design decisions, the cheaper and more effective these decisions will be. Thus techniques have been developed to accurately simulate all aspects of vehicle operation, and as these methods have become more refined and as available computational power has increased, vehicle modeling has taken a greater importance in the development cycle. Because of this, vehicle modeling is one of the major components of the EcoCAR 3 competition, where it is referred to as Systems Modeling and Simulation (SMS), and great efforts were taken to follow a similar process to that of General Motors and other automakers in intelligently making and using predictions about the performance and efficiency of our vehicle to guide major design decisions.

For our simulation, the Simulink product suite by MathWorks was used extensively, and its operation and setup will be described in this section in detail. First, however, the principles of vehicle simulation and system simulation in general will be discussed.

Vehicle Longitudinal Motion

Vehicle motion in a longitudinal direction, where "longitudinal" refers to the vehicle's rear-to-front axis, can be modeled by the following equation:

$$\partial ma = F_t - \frac{1}{2}\rho A_f C_d v^2 - mgf_{rr} - mg\sin\alpha$$

Where ∂ is the rotational inertia mass factor, m is vehicle mass, a is vehicle logitudinal acceleration, F_t is total tractive effort at the tires' contact patches, ρ is the density of air, A_f is

the frontal area of the vehicle, C_d is the drag coefficient of the vehicle, v is vehicle velocity, g is gravitational acceleration, f_{rr} is rolling resistance (which can be calculated most accurately as a function of velocity, as shown below), and α is vehicle angle with respect to gravity. More simply, the above equation states that the net force on the vehicle is the sum of tractive effort provided by the powertrain, wind resistance, rolling resistance, and gravitational force.

The rotational inertia mass factor ∂ is required to include the effect of the rotational inertias in a vehicle (all shafts, gears, pulleys, clutches, etc. that are connected to the wheels) on the vehicle's apparent linear inertia. Because these components must gain rotational momentum as the vehicle gains linear momentum, the principle of conservation of energy states that the sum of linear forces on the vehicle must not simply be its mass multiplied by its acceleration but must include this effect as well.

The rolling resistance coefficient f_{rr} is often defined as a velocity-independent constant dependent on tire composition and road type. The following table gives some representative values for approximate rolling resistance calculations for various tire and road situations [10].

Table 23: Rolling Resistance Coefficients

0.001-0.002	Railroad steel wheels on steel rails
0.004	Bicycle tire on asphalt road
0.01-0.015	Car tires on concrete
0.2-0.4	Car tires on loose sand

While this approximates the effect of rolling resistance, a more detailed model known as the Magic Tire formula can be used for more accurate results. This formula, developed empirically by Hans B. Pacejka, a professor at Delft University of Technology [11], is used by MathWorks in various forms in their Simulink "Tire (Magic Formula)" block [12]. GM released tire data in the

form of coefficients for the pressure- and velocity-dependent form of the Magic Equation, which is shown below.

$$F_{rr} = \left(\frac{P}{P_0}\right)^\alpha \left(\frac{N}{N_0}\right)^\beta N_0\left(A + B|v_{hub}| + Cv_{hub}^2\right)$$

Where P is tire pressure, v_{hub} is vehicle velocity, P_0 is 1 Pascal, N_0 is 1 Newton, and α, β, A, B, and C are the experimental coefficients provided by GM (omitted in compliance with non-disclosure agreements). The value computed by this equation, referred to from now on by $F_{rr}(v)$, replaces the $mgf_{rr}(v)$ term, as it carries units of force.

Modeling Driver Input

The longitudinal vehicle dynamics equation can be restated in standard calculus notation as

$$\partial m\ddot{x} = F_t - {}^1\!/_2\, \rho A_f C_d \dot{x}^2 - F_{rr}(\dot{x}) - mg \sin \ddot{x}$$

Where x is the distance traveled by the vehicle in the longitudinal direction. While at first glance this appears to be a 2nd order partial differential equation with respect to x, the term F_t will vary irregularly over time as a driver requests varying levels of torque, making an analytical solution unrealistic.

Driver control of a vehicle is a difficult effect to model, with numerous factors such as psychology, road and weather conditions, and the behavior of other drivers affecting driving style just as much as easier-to-quantify inputs such as speed limit and traffic conditions. Thankfully, an easy-to-use method for simulating driver behavior is found in the use of drive cycles. To allow for standardization of vehicle testing in laboratory conditions, the United States Environmental Protection Agency (EPA) has developed a series of desired velocity profiles known as Dynamometer Drive Schedules, or drive cycles for short. A drive cycle pupports to give a representative speed trace of an average driver traversing a given generic situation such

as highway, city, or mixed driving. Actual vehicles' fuel economy and emissions characteristics are tested on dynamometers using a combination of these standard drive cycles, and the results of these tests are weighted to predict the average American's long-term driving habits. Thus these drive cycles are crafted to be representative of real-world driving patterns, and because they come in the form of a time series of vehicle speed demands, they are easy to implement in driver modeling.

In EcoCAR 3, the following drive cycles are used in both vehicle modeling and, in later years of the competition, on-road vehicle testing. The 505 cycle, so called because it is the first 505 seconds (out of 1352) of the EPA Urban Dynamometer Driving Schedule, simulates normal city driving conditions, and the EPA Highway Fuel Economy Test Driving Schedule, or HWFET, simulates normal highway driving under 60 miles per hour. The EPA's US06 cycle represents more aggressive driving, and for EcoCAR 3 is broken into two cycles known as US06 City, with steep stop-and-go driving, and US06 Highway, with a quick acceleration to sustained fast driving at a max speed of 80 miles per hour. The emissions and energy consumption results of these four EcoCAR 3 drive cycles are averaged with the following weightings per competition rules [4].

Table 24: EcoCAR 3 Drive Cycle Weighting

Drive Cycle	Weighting
505	0.29
HWFET	0.12
US06 City	0.14
US06 Highway	0.45

By providing a desired velocity profile, these drive cycles take care of all of the aforementioned effects that influence a driver's behavior and simplify them to speeds that must be maintained at

given times. This enables very simple modeling of driver input as a PID controller operating on the difference between vehicle speed and drive cycle-determined desired speed over time.

Since now this driver input can be converted into tractive force by some modeling of the driver's foot on the pedals and how the powertrain responds to such, the vehicle longitudinal dynamics equation can be revisited for the purposes of full vehicle modeling.

Discrete Time Simulation

In a discrete time simulation, equations that represent the system under investigation, often differential equations or systems of such, are solved at discrete time intervals using computational methods. This is opposed to continuous time simulations, in which analytical solutions may be used, as more complex systems with nonlinearities and discontinuities, such as driver input, cannot be represented in that more mathematically "pure" form. In a discrete time simulation, the system is assumed to behave according to some form of "steady state" behavior over the discrete time interval, with the order of behavior represented defined by the computational method employed. For example, in a first-order model of a body with mass m in motion with a single force F acting upon it, a position x (0^{th} order derivative of position) and velocity v (1^{st} order derivative of position) will be computed based on the position and velocity of the prior time step (where position is defined as the 0^{th} order description of the system). The governing equations at each time step will be:

$$x_{i+1} = x_i + v_i t$$

$$v_{i+1} = v_i + \frac{F}{m} t$$

Where t is time.

For this simple situation the first-order model is sufficient, as no higher-order derivatives are necessary to describe the body's behavior and they would only incur greater computation cost if

83

included. However, for more complex systems it is necessary to employ higher order derivatives to accurately model system changes over time.

For complex vehicle simulation, variables like fuel economy and vehicle speed are dependent on a number of interrelated effects including engine and motor efficiency, engine manifold air pressure, battery state of charge, and themselves. However, the order of derivative required for accurate solution of the representative equations of this system is not immediately apparent.

Simulink features a number of solvers of varying orders, summarized in Table 25 below. The higher the order of accuracy, the greater the computational time, so it was desired to choose the lowest-order solver possible that would generate accurate results.

Table 25: Simulink Fixed-Step Solvers

Solver	Integration Technique	Order of Accuracy
ode1	Euler's Method	First
ode2	Heun's Method	Second
ode3	Bogacki-Shampine Formula	Third
ode4	Fourth-Order Runge-Kutta (RK4) Formula	Fourth
ode5	Dormand-Prince (RK5) Formula	Fifth
ode8	Dormand-Prince RK8(7) Formula	Eighth

To determine what exactly "accurate" meant, the vehicle simulation was first run with a high-order variable step solver that generated the benchmark results. Variable step solvers use a background algorithm to detect the amount of change in the system from one time step to the

next. If the change is great, the next time step will be shortened to capture the effects of this change, and if it is small, the next time step will be lengthened to decrease computational time. While using a variable-step solver is a good way to speed up a simulation's run time while retaining or even increasing accuracy, parts of the vehicle simulation must be compiled to devices which require fixed step computation (this process will be described in detail in the following chapter). While frequently switching between fixed- and variable-step solvers is technically possible, certain mechanics throughout the model require changes when this happens and thus a fixed-step solver was used instead.

A (generally) binary search was carried out to find the lowest order (therefore computationally fastest) fixed-step solver that would still generate the same results as the benchmark variable-step results, with the results of this search shown in Table 26 below. The 4th-order Runge-Kutta Formula was the fastest solver that yielded the same results as the benchmark. The 3rd-order Bogacki-Shampine Formula was even faster and was the fastest tested, but it deviated from the benchmark and was thus considered insufficient for this system.

Table 26: Solver Binary Search Results

Solver	CD Electrical Energy Usage (Wh/km)	CS Fuel Energy Usage (Wh/km)
ode45 (benchmark)	157.2	396
ode8	157.2	396
ode5	157.2	396
ode3	159.9	397.6
ode4	157.2	396

With the overall model framework set up, the modeling of the individual components and

subsystems of the vehicle will now be discussed.

Hybrid Vehicle Model

Simscape Overview

Simulink is a powerful block diagram modeling tool that allows for simulation of any system that can be represented by any of a number of mathematical regimes, including algebra, discrete mathematics, calculus, differential equations, stochastics, and more. However, deriving the equation for a complex system is often an onerous task that requires significant rework if even a small addition or subtraction is made to the system. So MathWorks has also provided numerous extra "toolboxes," or sets of blocks that feature built-in equations for common applications. Connecting these blocks correctly will, under the hood, create the same system of equations to model the system as would be created if all of the mathematical elements required were manually connected instead.

For physical systems, one of these toolboxes is known as Simscape, and unlike many of the other toolboxes, Simscape has multiple special types of connections in addition to specialized blocks. When a shaft connects two rotating components in the real world, for example, the shaft transmits torque and shares speed between the two components. Simscape rotational blocks not only feature rotational domain math (a gear coupling block, for example will enforce a gear ratio if specified), but also connect to each other with rotationally conserving "shaft" connections that satisfy:

$$T_1 = T_2$$

$$\omega_1 = \omega_2$$

Where T is torque and ω is angular velocity, and points 1 and 2 are two points within a system connected by this shaft. These connections themselves have no mass or rotational inertia and are infinitely stiff and strong, guaranteeing that the above relationships hold true.

Similarly, for the blocks in the Simscape translational domain,

$$F_1 = F_2$$

$$v_1 = v_2$$

Where F is force and v is velocity. A connection in this domain represents a rigid, massless connection that transfers force and shares movement between two or more points.

Simscape is extremely useful for modeling multi-physics systems like hybrid electric vehicles where a full model could easily include translational, rotational, thermal, fluid, and electrical domains. Blocks from some of these domains are used extensively in the team vehicle's model, with future plans to incorporate elements from each of them.

To aid in model visualization, Simscape connections of different domains are shown in different colors. A guide to aid in understanding the Simulink images that follow is included here in Figure 42.

Default Simulink variable connection

Physical Signals (similar to normal Simulink signals but carrying a unit along with their value)

Translational domain

Rotational domain

Pneumatic domain

Figure 42: Simscape Connections Legend

Overall Layout

The Georgia Tech team vehicle model is laid out in a way that is meant to be both visually clean and representative of the actual vehicle architecture. At the plant level of the vehicle, as shown below in Figure 43, the engine connects to the transmission with a shaft connection. The output

of the transmission connects to the drivetrain, as does the motor through a Simple Gear block representing the transfer case. The battery appears free-floating but is connected to the other systems via Goto and From blocks that carry signals without visible (and sometimes cluttering) connecting lines.
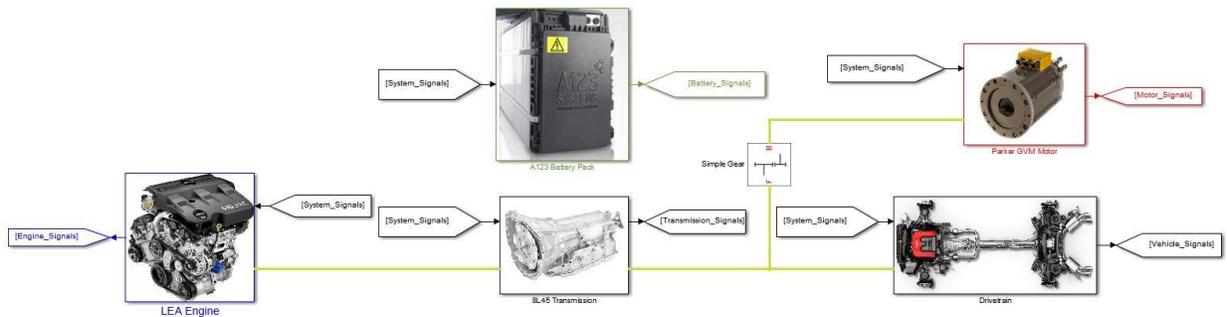
Figure 43: Vehicle Model, Plant Level

These subsystems' interiors will now be explored.

Vehicle Body and Drivetrain

The Drivetrain subsystem contains the rear differential, half-shafts and their inertias, front and rear wheels and brakes, and a vehicle body block that implements the vehicle longitudinal motion equation. This is shown in Figure 44 below.
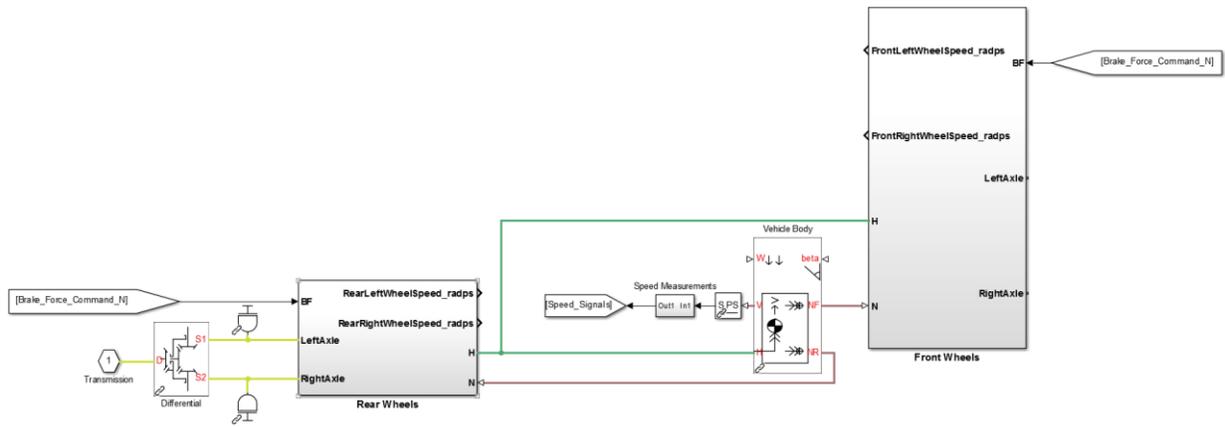
Figure 44: Vehicle Model, Drivetrain Subsystem

Diving one layer further into the wheel subsystems, Figure 45 shows that the Simulink tire

blocks receive the yellow rotational connections from the powertrain. This rotation is converted

into the car's translational motion conveyed on the green lines, which go to the Hub port on the

vehicle body block above, with the knowledge of the tire's static coefficient of friction, enabling

Simulink to calculate at each time step whether the tire is "sticking" or not, and if not, how much

forward force is being passed on to the vehicle.
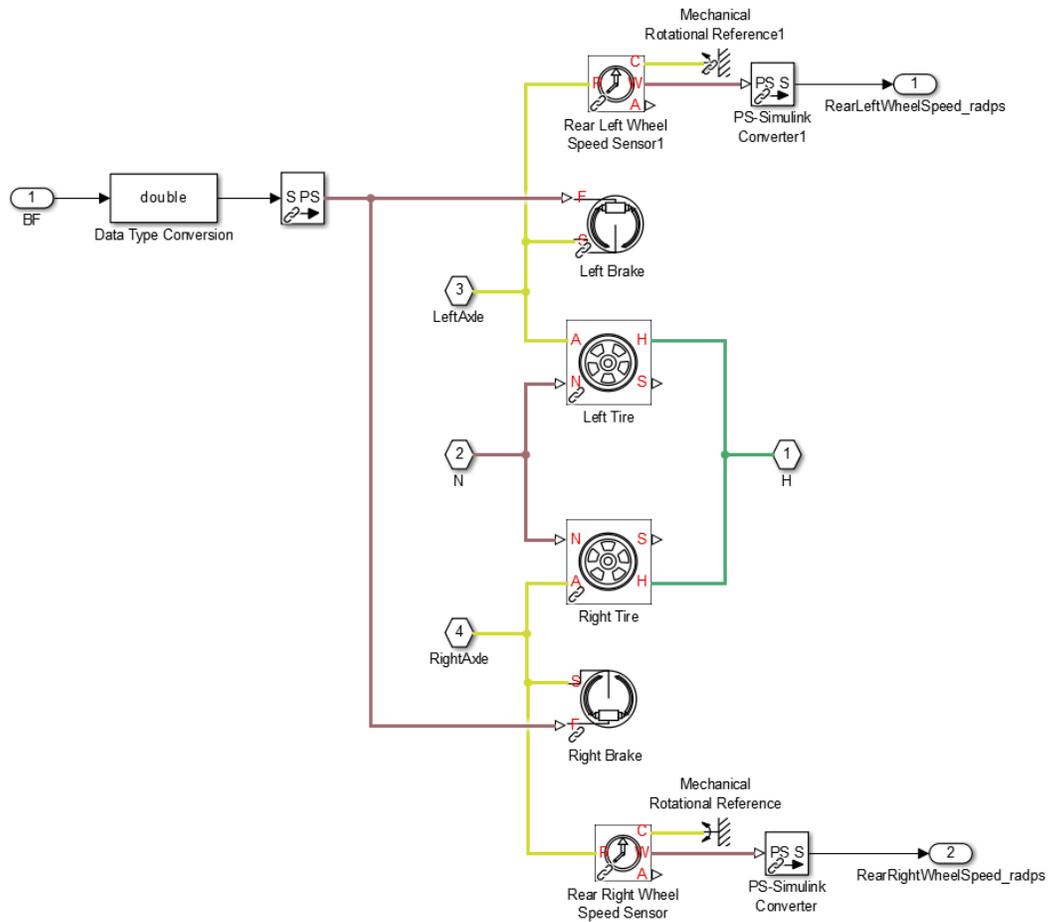
Figure 45: Vehicle Model, Wheels Subsystem

Engine

The engine model shown in Figure 46 below was developed for EcoCAR 3 teams and donated by MathWorks, with very little modification from the Georgia Tech team. It is shown here only to demonstrate the level of powertrain modeling accuracy featured in the team's vehicle model.
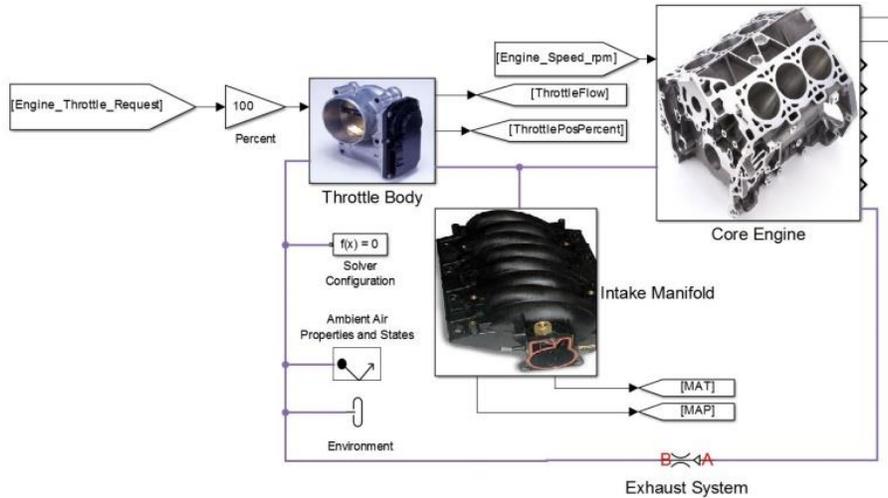
Figure 46: Vehicle Model, Engine Subsystem

As shown, Pneumatic blocks are used to model the airflow through the intake, throttle, engine, and exhaust system. The amount of fuel injected and torque produced is a function of engine speed and this airflow, which in this model and in a real vehicle is controlled by the opening and closing of the throttle valve. Including this level of throttle response builds confidence in this model's ability to accurately predict acceleration.

Transmission

The transmission model is also MathWorks-developed and donated, though it came as a model of the 6-speed 6L45 from the 2015 Camaro and had to be expanded to the 8-speed 8L45 for the team's vehicle.

As shown in Figure 47 below, the transmission is comprised of a series of eight gears attached to the same shaft at the bottom of the diagram. These are separated from the output shaft to the right by four double-sided synchronizers connected to shift levers, which model the clutch action within a manual transmission, a simplification from the operation of an automatic transmission but capturing the relevant physics to a satisfactory degree.
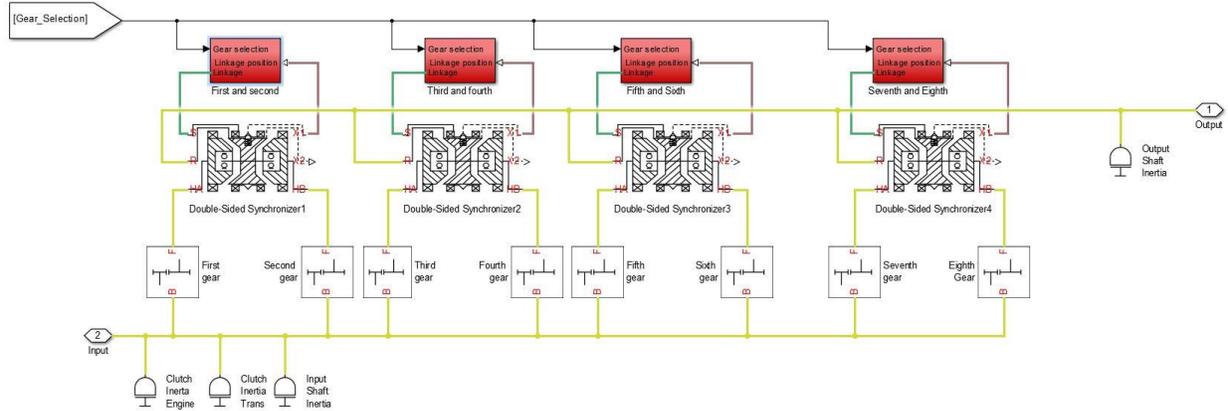
Figure 47: Vehicle Model, Transmission Subsystem

When the "Gear_Selection" signal in the upper left corner goes to a number 1 through 8, each of the shift lever subsystems, the red blocks towards the top, receive that value. Inside these blocks are levers and actuators, connected to the clutches in the middle of Figure 47, that are programmed to only respond to the signal corresponding to the correct clutch.

For example, when the gear selection signal reads "3," the actuator in the subsystem labeled "Third and fourth" pushes a lever that moves a cone clutch, synchronizing the torque converter output shaft and the 3rd gear shaft, before engaging a dog clutch to finish the gear shift. If the gear selection signal were to be 4 instead, the actuator would have acted in the opposite direction to engage the fourth gear. For either case, all the other actuators are commanded to remain neutral, neither going one way or the other to engage their gears' clutches.

This model allows for accurate simulation of gear shift events, which again is helpful for building confidence in acceleration simulations.

Motor

The motor thus far has been modeled at a much simpler level than the previous two subsystems, utilizing Simulink blocks without any Simscape capabilities beyond simple torque actuation on the rest of the vehicle.

Because the time constant of a motor is much smaller than the effects of throttle lag and gear shifting, the motor is assumed to have instantaneous torque production on command. Current drawn by the motor is calculated as follows,

$$I_{motor} = \frac{1}{\eta(T,\omega)}\left(\frac{T\omega}{V_t}\right)$$

Where T is motor torque, ω is motor angular velocity, η is motor efficiency, and V_t is battery terminal voltage, to be discussed in the following section. Motor efficiency is included as a function of torque and speed and can be referenced from manufacturer data or treated as a constant. While currently a value of 85% is used as a constant motor efficiency, an efficiency lookup table is soon to be included, and simultaneous development is underway to convert the entire motor model into a much more accurate implementation of 3-phase synchronous machine equations.

Battery

The battery model is similar to the motor model in not using Simscape features to represent its behavior. Instead it is represented as a voltage source with terminal voltage V_t that is calculated as a function of its open-circuit voltage V_{OC}, current draw, and resistance.

$$V_t = V_{OC}(SOC, T_{battery}) - IR(SOC, T_{battery})$$

As stated in these equations, open-circuit voltage and resistance are both functions of SOC and battery temperature and are calculated in the team vehicle model using lookup tables with

manufacturer data. Thus far, battery temperature has been assumed to be a constant 25 °C, a value that, practically, would only be realistic on a cool day with very light driving.

Future Model Improvements

As has been suggested, numerous opportunities to improve this model exist. These include:

- Converting the motor and battery into Simscape electrical elements that more accurately represent their behavior, which of course then requires an inverter model as well
- Developing a model variant of the engine subsystem to leave out pneumatic effects. These effects are not nearly as useful for the longer emissions and energy consumption simulations as for acceleration simulations and should be made optional to speed up simulation time.
- Integrating cooling systems into the motor, inverter, and battery subsystems so that their performance in a range of temperatures may be accounted for
- Numerous other additions that will be discussed in the chapter on SIL and HIL development

Thus while much has been done to create as accurate and useful a vehicle model as possible, work remains to fill up the final two years of EcoCAR 3 for the Georgia Tech team's modeling and simulation team.

The results of the model at its current stage of development, averaged over the four EcoCAR 3 drive cycles, are included below.

Table 27: Vehicle Modeling Results

Specification	Target	Requirement	Team Vehicle
Acceleration, IVM–60 mph	5.9 sec	7.9 sec	7.3 sec
Acceleration, 50–70 mph (Passing)	7.3 sec	9.9 sec	2.7 sec
Acceleration Events Torque Split (Front/Rear)	0% F, 100% R	49% F, 51% R	0% F, 100% R
Total Vehicle Range	**	150 mi (226 km)	167 mi
CD Mode Range	**	N/A	25.3 (Blended)
CD Mode Total Energy Consumption	**	N/A	329 Wh/km (Blended)
CS Mode Fuel Consumption	**	N/A	676 Wh/km
UF-Weighted Fuel Energy Consumption	**	N/A	398 Wh/km
UF-Weighted AC Electric Energy Consumption	**	N/A	120 Wh/km
UF-Weighted Total Energy Consumption	700 Wh/km (30 mpgge)	840 Wh/km (25 mpgge)	517.5 Wh/km (40.5 mpgge)
UF-Weighted WTW Petroleum Energy Use	420 Wh PE/km	750 Wh PE/km	113 Wh PE/km
UF-Weighted WTW Greenhouse Gas Emissions	225 g GHG/km	250 g GHG/km	156 g GHG/km
UF-Weighted WTW Criteria Emissions	TBD* (Base Vehicle)	TBD* (Base Vehicle)	.039 g/km

Thermal System Modeling

While the vehicle's engine and transmission will be cooled by stock cooling systems at least through Year 2 of the competition, the electric motor, inverter, and battery pack each need their own cooling systems due to the high currents that flow through them. Modeling these cooling systems was necessary to select pumps, piping, and radiators for the liquid coolant to be used.

While Simulink is a fantastic tool for vehicle modeling, its thermal systems capabilities are less full-fledged than its mechanical systems capabilities. Because of this, the cooling systems were modeled using fundamental heat transfer equations in a code-based interface known as Engineering Equation Solver instead. This allowed for intelligent component sizing based on worst-case scenario analysis without simultaneously simulating the various other time-consuming dynamics of the vehicle, though as mentioned, future work for the team will involve integrating the two approaches so that the thermal systems may be accurately modeled in the context of realistic vehicle behavior. The use of Engineering Equation Solver will be described in greater detail shortly, but first the engineering analysis of these systems will be discussed.

Battery Cooling Overview

Resistive heating in the modules of the vehicle's battery pack causes temperatures to rise above safe levels if the system is not cooled. The manufacturer, A123, specifies that 60 °C is the maximum allowed module temperature, above which the battery control module (BCM) will disable any current flow from or to the modules. To prevent this worst-case scenario, the BCM will first decrease allowable current output linearly when the temperature ranges from 50 °C to 60 °C. Thus for a robust hybrid control strategy it is desirable to keep the maximum temperature within any module lower than 50 °C so that full current is always available for driving and regenerative charging.

The modules can be liquid- or air-cooled, each having associated pros and cons. Air flowing across a surface will generally have a convective heat transfer coefficient orders of magnitude lower than fast-moving liquids on solids, but a source of already-cooled air is only inches away from the trunk in the passenger cabin. Blowing conditioned air from the cabin into the trunk to keep the modules cool enough for safe operation makes for a simple design that wouldn't require radiator selection or packaging. However, air from the battery box could contain dangerous gases in the event of battery failure, so per competition rules the passage from the cabin into the trunk would need to be shut off to airflow whenever the blowing fans weren't on. This requires a control algorithm and additional hardware and wiring.

A liquid cooling strategy, however, requires a radiator of adequate heat transfer capability to be placed in a location with adequate airflow and piping to be routed to that location. A pump would then need to be selected and placed in the cooling loop as well, adding a power draw on the battery. The entire system, of course, would need to be leak-tested to ensure that liquid coolant doesn't leak onto the numerous sensitive components in the vehicle.

While the packaging and engineering effort required for the liquid cooling solution makes it more complex on the whole than air cooling, the benefits of having more powerful heat transfer abilities to guarantee able battery operation in a variety of climates persuaded the team to select liquid cooling.

A cooling plate was designed to sit under the modules through which copper tubes could be routed. The plate, shown in Figure 48 below, is made of aluminum to take advantage of that metal's high thermal conductivity and low density.
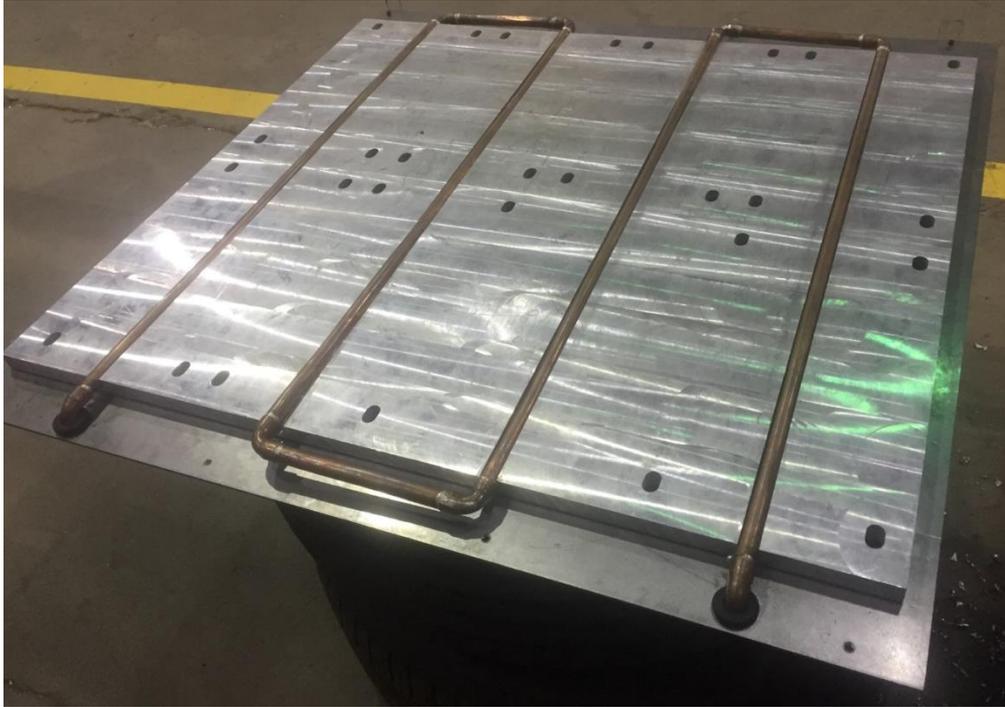
Figure 48: Battery Cooling Plate

The design of this system and the selection of the numerous components required to cool the battery modules will be presented in parts, starting with the analysis of the battery modules' internal heat generation and then moving to the cooling plate, radiator, and piping, before solving these dependent sub-problems together. Because much of the analysis is similar for the inverter and motor, besides the simpler analysis of these components' internal heat generation, the format used for solving the battery system will then be applied to those components.

Battery Thermal Solution- Initial Setup

To understand the internal heat generation of the battery modules, the problem statement is as follows. A battery module in the vehicle's 7-module pack is to be cooled by a fluid which will enter the system at a temperature T_1. The battery module sits on top of the cooling plate through which the coolant flows. Heat is transferred from the sides of the module to the

ambient air ($\dot{Q}_{sides}$) and below the battery to the coolant ($\dot{Q}_{bottom}$), and generated internally due to electrical current ($\dot{Q}_{gen}$). The coolant then flows to a radiator in the front of the vehicle and is cooled by the ambient air ($\dot{Q}_{rad}$). Leaving the radiator, it then returns to the battery, completing the loop. A schematic representing this layout and showing heat transfers and coolant temperatures is shown in Figure 49 below.
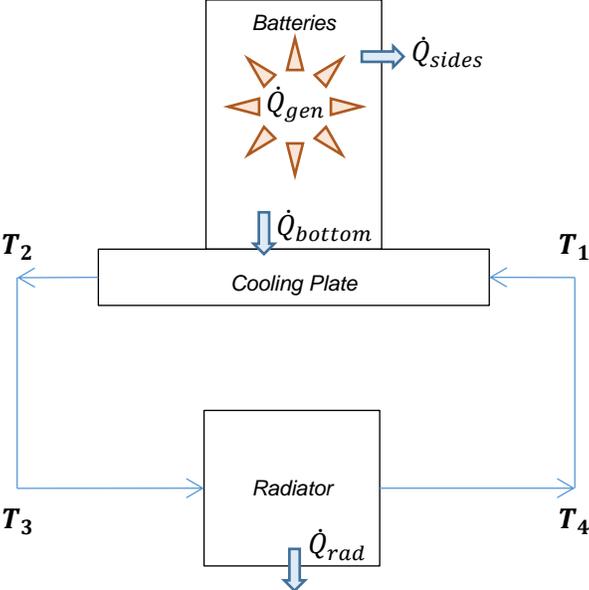
Figure 49: Battery Cooling Schematic

A selection of known and unknown parameters for the battery module and cooling plate portion of this problem is tabulated below. The full listing of the 212 variables in this solution can be seen in the Engineering Equation Solver codes in Appendix 2.

Table 28: Control Variables

Parameter	Units	Notes
T_{air}	°C	Ambient air temperature
I_{pack}	A	Total pack current
T_1	°C	Coolant inlet temperature
v_{air}	mph	Speed of vehicle moving through air

Table 29: Selected known parameters

Parameter	Value	Units	Notes
$A_{battery,top}$	.05	m²	Surface area of battery top or bottom
$A_{battery,sides}$	.147	m²	Surface area of battery sides
$height$	.243	m	Battery module height
$length$	.303	m	Battery module length
$width$	.165	m	Battery module width
$V_{battery}$	.0121	m³	Battery module volume
h_{top}	5	$W/{m^2 K}$	Convective coefficient for battery-air interface, estimated
ID	.01092	m	Pipe inner diameter for 3/8" Type L copper pipe

Table 30: Selected Unknown Parameters

Parameter	Units	Notes
$T_{battery}$	°C	Temperature battery module
$\dot{Q}_{sides}$	W	Convective heat transfer to ambient air
$\dot{Q}_{bottom}$	W	Heat transfer to coolant pipe and to coolant
T_2	°C	Coolant outlet temperature
h_{fluid}	$W/{m^2 K}$	Convective heat transfer coefficient for coolant-pipe interface
U_{rad}	$W/{m^2 K}$	Overall heat transfer coefficient for radiator

A few assumptions for this solution are now stated:

- The battery is at a single temperature throughout its volume, an assumption informed by another team's experience with this pack in EcoCAR 2

- The surface temperature of the pipe in the cooling plate is at a constant temperature

- No heat transfer occurs within the cooling fluid along the axis of flow, an assumption generally valid for non-metallic fluids

- The thermal capacity of the battery is large enough that only steady state effects need to be analyzed

- No heat transfer to the air occurs along the copper piping between the radiator, pump, and cooling plate, a conservative assumption since these pipes will actually provide additional cooling

- No heat is added to the fluid by the pump

To justify the final assumption, if the thermal mass of the battery pack (123 kJ/K) is subjected to the max heating load of 1512 W, it will take 13.5 minutes for the pack temperature to rise 10 °C.

To solve this problem, a number of relationships must be stated. First, to not violate conservation of energy, the heating load within the battery must be equal to the heat dissipated to the air through convection above the modules and below to the coolant, and the heat transferred to the coolant must equal the heat transferred to the air by the radiator. In other words, with heat transferred into the battery defined as positive,

$$\dot{Q}_{gen} = I_{pack}{}^2 R$$

$$\dot{Q}_{gen} + \dot{Q}_{sides} + \dot{Q}_{bottom} = 0$$

$$\dot{Q}_{rad} = \dot{Q}_{bottom}$$

Next, this heat flow is not driven by itself- it must be driven by a temperature differential across a medium with a thermal resistance. Thus the heat transfer can be restated as:

$$\dot{Q}_{bottom} = h_{pipe} A_{pipe} T_{LM,pipe}$$

$$\dot{Q}_{top} = h_{battery} A_{battery,sides} (T_{air,batterybox} - T_{battery})$$

Where T_{LM} is the appropriate form of the log mean temperature difference for the given scenarios.

For these equations' solution, the second assumption is restated. Because of the low thermal resistivity of aluminum and because the aluminum cooling plate will surround the copper pipes on most sides, it is assumed that the heat transfer within the cooling pipes can be approximated as having a constant surface temperature T_s. This assumption is backed up by the low temperature change observed in the coolant from inlet to outlet shown in the results at the end of this section.

The general definition of $T_{LM,pipe}$ is

$$T_{LM,pipe} = \frac{\Delta T_1 - \Delta T_2}{\ln\left(\frac{\Delta T_1}{\Delta T_2}\right)}$$

The appropriate definitions of its terms for this the constant surface temperature assumption is

$$\Delta T_1 = T_s - T_1$$

$$\Delta T_2 = T_s - T_2$$

The radiator, however, gives a different heat transfer scenario entirely. Because one fluid medium is passing over another and their flows are perpendicular to each other, a radiator can be thought of as a cross-flow heat exchanger [13]. Cross-flow heat exchangers behave in a

similar way to counterflow heat exchangers but with a correction factor *F* applied to the total

heat transfer, as below.

$$\dot{Q} = FUAT_{LM}$$

First, for a cross-flow heat exchanger, ΔT_1 and ΔT_2 are redefined as (using this solution's

variable names):

$$\Delta T_1 = T_3 - T_{air,out}$$

$$\Delta T_2 = T_4 - T_{air}$$

Next, the correction factor F is an experimentally derived function of two parameters commonly

known as P and R, defined for this situation as follows.

$$P = \frac{T_4 - T_3}{T_{air,in} - T_3}$$

$$R = \frac{T_{air,in} - T_{air,out}}{T_4 - T_3}$$

Finding F from these two values by hand is aided by the use of this plot [13]
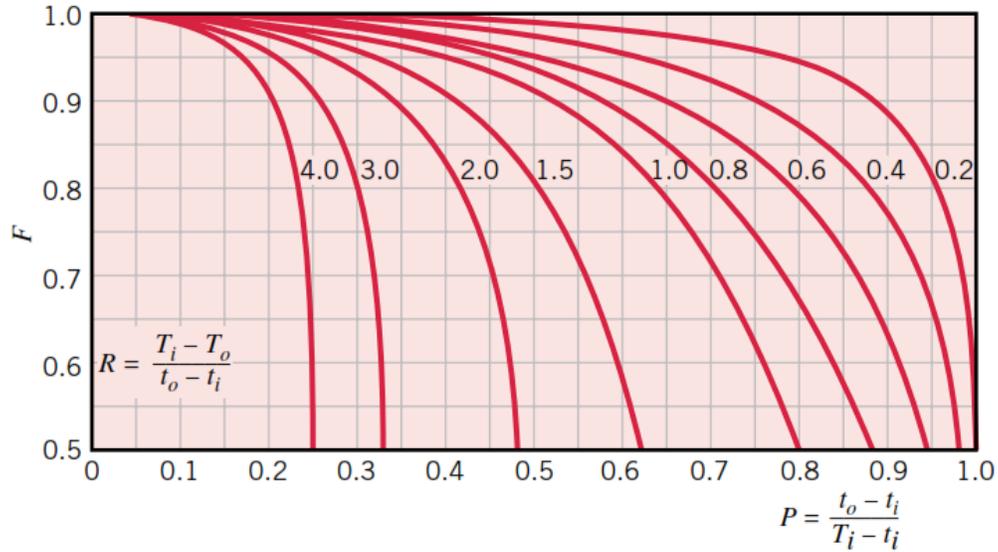
Figure 50: Correction factor for a single-pass, cross-flow heat exchanger with both fluids unmixed

For this solution, however, the corresponding function built into Engineering Equation Solver was used to find F with more precision.

A few more basic relationships can be defined with the assumption that no heat transfer occurs along the length of the piping bringing coolant to and from the radiator. Thus,

$$T_1 = T_4$$

$$T_2 = T_3$$

As well, heat transfer to and from these fluids must cause a corresponding change in temperature, and thus the following equation must be satisfied for coolant in the cooling plate, coolant in the radiator, and air over the radiator.

$$\dot{Q} = \dot{Q}_{bottom} = \dot{Q}_{rad} = \dot{m}c_p(T_{in} - T_{out})$$

This scenario is almost able to be solved, with two pieces still missing: the heat transfer coefficient within the cooling plate, and the heat transfer area and overall heat transfer coefficients of the radiators.

Heat transfer coefficient within cooling plate

The heat from the battery must pass through a series of resistances to transfer to the coolant. The analysis of the fluid-to-pipe heat transfer coefficient will be considered first, followed by the more straightforward addition of the solid medium resistances around it.

For internal flow such as this, the heat transfer coefficient *h* will be a function of Nusselt number Nu_D that will itself be found through an experimentally-found correlation with Prandtl number Pr. A correlation that works well for a wide range of flows, including the laminar-turbulent transition region, is the Gnielinski correlation below [13], which has been found to yield smaller errors than the more commonly used Dittus-Boelter equation.

$$Nu_D = \frac{(f/8)(Re_D - 1000)Pr}{1 + 12.7(f/8)^{1/2}(Pr^{2/3} - 1)}$$

Where f is the tube's friction factor and Re_D is the Reynolds number.

Nu_D is defined as the ratio of internal convection to internal conduction in a flow,

$$Nu_D = \frac{hD}{k}$$

Where D is the inner diameter of the pipe and k is the thermal conductivity of the fluid, both known values. Solving for Nu_D with the first equation thus enables h to be found.

While this value will be determined by an iterative solution, a few thermal resistances must also be added to correctly determine total heat transfer to the fluid. As shown in Figure 51 below, between the battery module and the coolant are a number of layers that add resistance to the heat transfer's path.
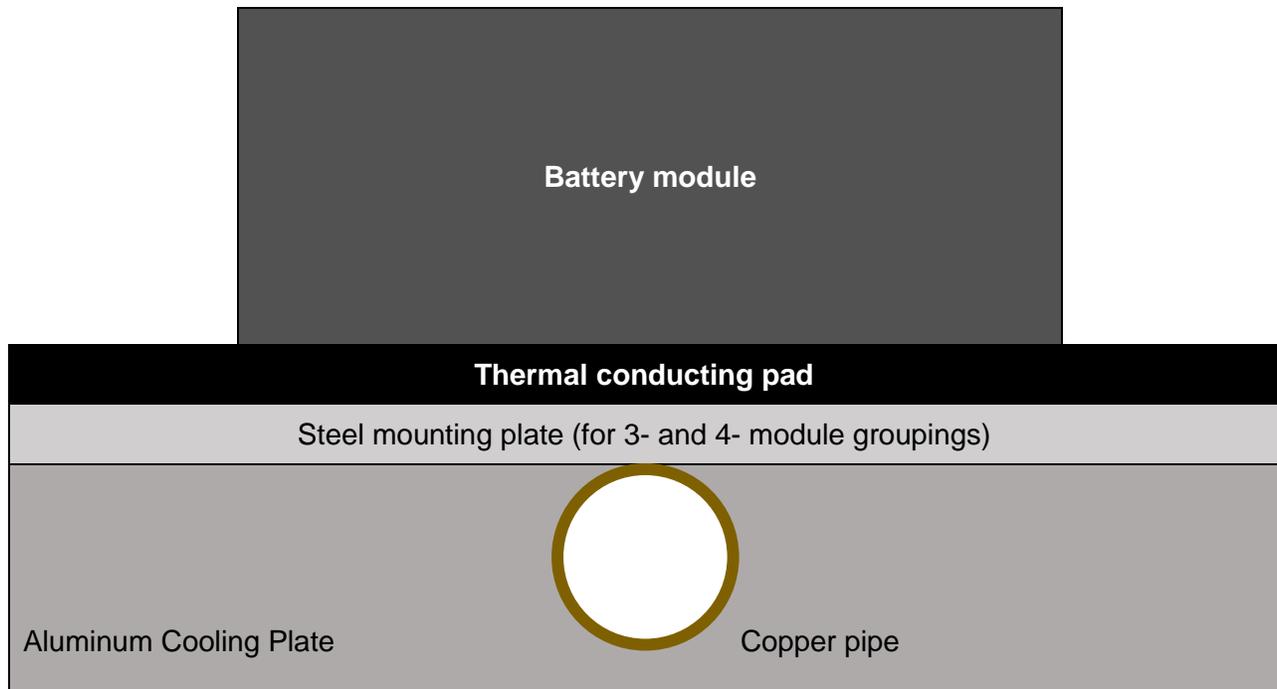
Figure 51: Battery Thermal Resistances

Calculating the correct resistance between the battery module and the fluid is complicated by two factors. The first is obvious in the image above- this is a complex 3D construction without an analytical solution. However, the high thermal conductivities of aluminum (~210 W/m-K [14]) and copper (385 W/m-K [15]) allow their resistances to be neglected, simplifying this problem to the series resistances of the module, thermal conducting pad, and steel plate.

Secondly, the value given by A123 for module thermal resistance is unrealistic in the form given, as it would require an impossible temperature differential across the battery and its surroundings, meaning this crucial value must be omitted from this analysis pending further dialogue with representatives from A123.

Radiator Heat Transfer

Numerous equations and relationships exist to correlate the various features of heat exchangers to their effective heat transfer area A and overall heat transfer coefficient U. However, to use

these equations to calculate the performance of an off-the-shelf radiator is difficult without knowing the precise counts and dimensions of channels, fins, and piping, and even then will be prone to error because of potential manufacturing tolerances issues and errors coming from experimental correlations. Experimental data on the specific product, however, can be used to take data from one heat transfer situation and predict performance for another.

Because of the size constraints inherent to packaging additional radiators in a vehicle, aftermarket transmission coolers were seen as a viable option for their size and availability. Trucool offers a line of radiators for this market and provided test results for each, summarized below.

Table 31: Trucool Radiator Test Conditions

Fluid	Automatic Transmission Fluid (ATF)
Air speed	15 mph
Flow rate	3 gallons per minute

Table 32: Trucool Radiator Test Data

Radiator	Cooling Rate [BTU/hr]
B7B	7500
L7B	15000
M7B	20500
H7B	29200

This data was used by making a few assumptions.

1. Air-to-metal thermal resistance dominates the sum of other resistances (metal interior and metal-to-ATF), meaning that overall heat transfer coefficient U can be assumed constant between the manufacturer's test data (with ATF) and the team's application (with ethylene glycol).

2. Air temperature for these manufacturer tests was carried out at 25 °C. This is a
   conservative assumption, as achieving these cooling rates at higher temperatures
   implies better cooling capability.

3. ATF temperature coming into the transmission was 170 °F, a normal ATF fluid
   temperature [16].

With these assumptions, the following system of equations was solved to find the product of
these radiators' U and A.

$$\dot{Q} = \dot{m}c_{p,ATF}\left(T_{in,ATF} - T_{out,ATF}\right)$$

$$\dot{Q} = \dot{m}c_{p,air}\left(T_{in,air} - T_{out,air}\right)$$

$$\dot{Q} = FUAT_{LM,ATF}$$

Where F and T_{LM} are found in an identical manner to that described in the preceding section.

Solving these equations yielded a single value for the product UA, which is sufficient for this
analysis because of the assumption that U would not change for the team's cooling (and of
course neither will radiator heat transfer area A). This parameter will be referred to as $(UA)_{rad}$.

All values necessary for full system solution have now been found.

Battery Thermal Solution

The following equations have now been set by the previous sections.

$$\dot{Q}_{gen} + \dot{Q}_{sides} + \dot{Q}_{bottom} = 0$$

$$\dot{Q}_{rad} = \dot{Q}_{bottom}$$

$$\dot{Q}_{gen} = I_{pack}{}^2 R$$

$$\dot{Q}_{bottom} = h_{pipe}A_{pipe}T_{LM,pipe}$$

$$\dot{Q}_{top} = h_{battery} A_{battery,sides} \left(T_{air,batterybox} - T_{battery} \right)$$

$$\dot{Q}_{rad} = (F)(UA)_{rad} T_{LM,rad}$$

$$\dot{Q}_{bottom} = \dot{m} c_p (T_2 - T_1)$$

$$\dot{Q}_{rad} = \dot{m} c_p (T_4 - T_3)$$

$$T_1 = T_4$$

$$T_2 = T_3$$

This is 10 equations and 10 unknowns ($\dot{Q}_{gen}$, $\dot{Q}_{sides}$, $\dot{Q}_{bottom}$, $\dot{Q}_{rad}$, T_1, T_2, T_3, T_4, F, and $T_{air,out}$, which contained in $T_{LM,rad}$). While a simple solution of this system is obviously possible, it is more accurate to account for how the physical property constants in these equation change at different temperatures. Although manual iteration is achievable and it could be possible to reach satisfactory convergence of these variables' values by consulting the physical property tables at each iteration, the solution is greatly aided by the previously-mentioned Engineering Equation Solver (EES) software. The main functionality of EES is to solve systems of equations iteratively, and it contains built-in functions for physical properties such as density, specific heat, and viscosity along with helpful heat transfer-related functions like finding the cross-flow heat exchanger factor F from above.

Referencing the control variables listed at the beginning of this section, the ambient temperature was set at 40 °C and airspeed at 30 mph. The problem was then solved with battery current iterated up to 120 amps, the maximum continuous current the battery could produce, to find maximum interior temperature at various loads. This was repeated for a range of ambient temperatures to find the maximum continuous current allowed at these temperatures, a result that allows for checking the performance of the team's cooling system and also will inform the hybrid powertrain control strategy. The worst-case scenario results are presented here in

graphical form using the interface developed within EES for easy results visualization (Figure 52).

$I_{battery}$ = 120 [A]
$T_{battery}$ = 47.98 [C]
$\dot{Q}_{gen}$ = 1512 [W]

$T_{air,box}$ = 42.83 [C]
$\dot{Q}_{trunk}$ = -6.543 [W]

T_2 = 44.99 [C]

$\dot{Q}_{bottom}$ = -1505 [W]
T_{surf} = 46.95 [C]

T_1 = 43.34 [C]

T_6 = 43.34 [C]

LPM = 15.15 [L/min]

Pump

Airspeed = 30 [mph]
$\dot{Q}_{rad}$ = -1505 [W]
T_{airout} = 40.87 [C]

T_5 = 43.34 [C]

T_4 = 43.34 [C]

T_3 = 44.99 [C]

T_{air} = 40 [C]
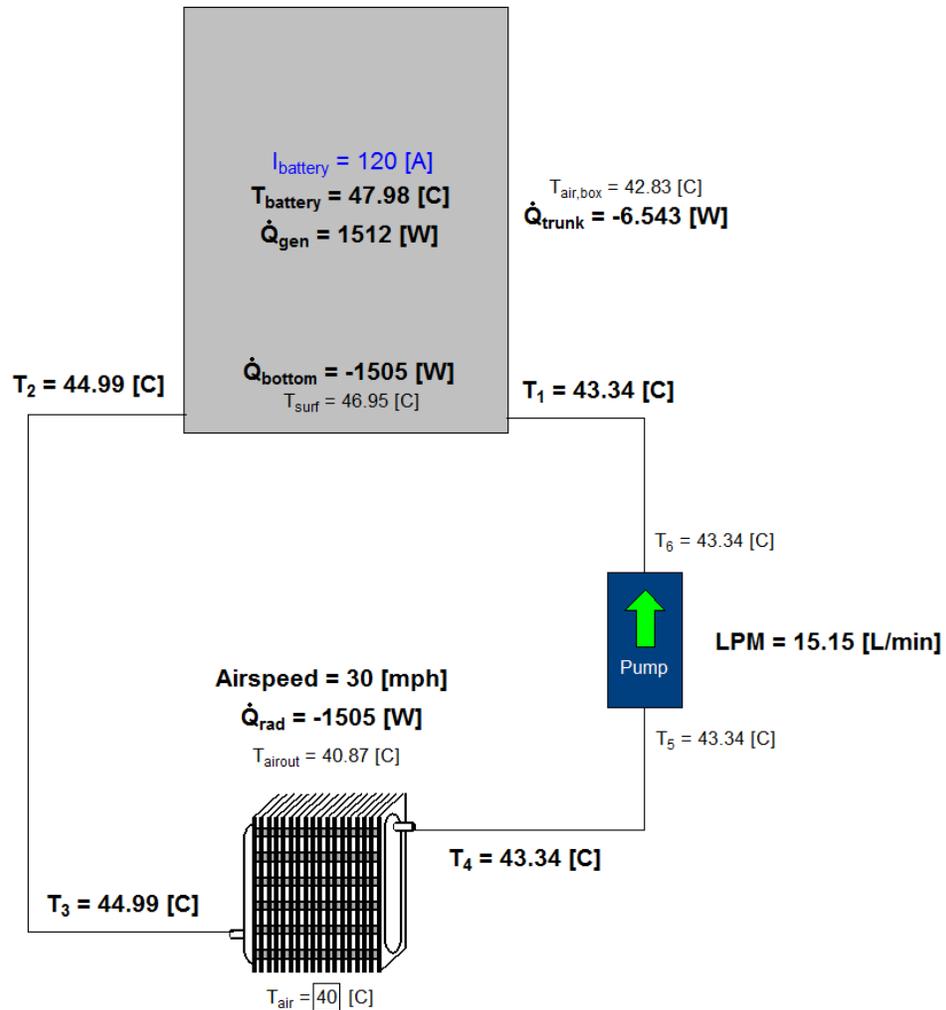
Figure 52: Worst-Case Scenario Battery Cooling

These results feature the coolant running through two Trucool M7B radiators in series. These radiators will be placed in front of the engine radiator in the front of the vehicle and should provide sufficient cooling even at the elevated temperature simulated. This radiator model was compared against other Trucool models and found to be the optimal selection for cooling the battery while also fitting in the limited space of the front of the vehicle. By comparison, one M7B

radiator or multiple smaller radiators would yield higher battery temperatures due to the overall decrease in radiator heat transfer area.

The battery temperature predicted with a maximum steady-state current at various temperatures is shown in Figure 53 below.
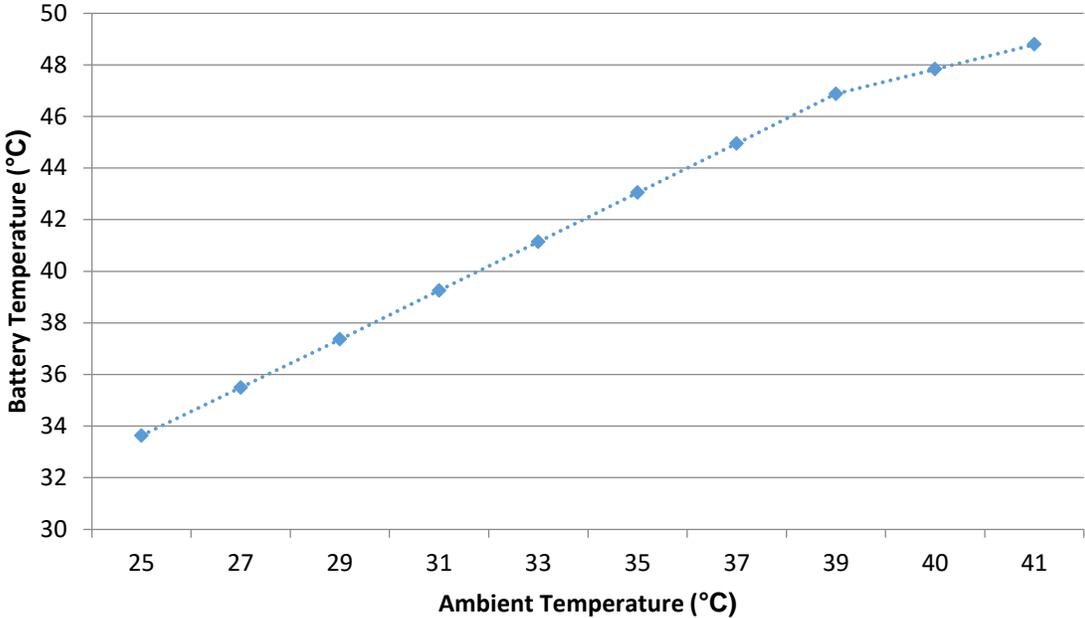
Figure 53: Battery Temperature vs Ambient Temperature

Motor and Inverter Thermal Analysis

As said, the analysis for the motor and inverter shared common elements with that of the battery but featured simplified studies of the heat transfer that takes place within these components themselves. Because their cooling channels are built-in and were not designed by the Georgia Tech team, this analysis trusts manufacturer datasheets instead of making any poorly-informed assumptions about them that could lead to erroneous results.

To relate the heat produced by the motor and inverter to the temperatures of the coolant, the following relationship was used.

$$\dot{Q}_{gen} = \dot{m}c_p(T_{in} - T_{out})$$

This ignores the convective heat transfer within these components and simply states that the temperature rise in the coolant will be proportional to the manufacturer-stated heat production.

With this, a simplified system of equations is developed for the motor and inverter:

$$\dot{Q}_{rad} = -\dot{Q}_{gen}$$

$$\dot{Q}_{rad} = (F)(UA)_{rad}T_{LM,rad}$$

$$\dot{Q}_{gen} = \dot{m}c_p(T_{in} - T_{out})$$

$$\dot{Q}_{rad} = \dot{m}c_p(T_{out} - T_{in})$$

$$\dot{Q}_{rad} = \dot{m}c_p(T_{in,air} - T_{out,air})$$

In these five equations are found five unknowns: $\dot{Q}_{rad}$, F, T_{in}, T_{out}, and $T_{air,out}$. These were solved using EES with ambient conditions of 40 °C and 25 mph vehicle speed, with results presented below.

Table 33: Motor and Inverter Cooling Results

	Motor	Inverter
Maximum allowable inlet temperature	65 °C	65 °C
Calculated inlet temperature	64.88 °C	51.24 °C
Heat generation	3596 W	1642 W
Flow rate	6.3 L/min	6.0 L/min
Radiator selected	Trucool L7B	Trucool L7B

Pump Selection and Flow Rates

In the preceding sections, coolant flow rate has been mentioned as if it were a natural constant. In reality, a pump can only produce a finite flow rate at a given system pressure drop. Because this itself is dependent on flow rate, the relevant fluid dynamics equations were integrated into each of the above analyses for iterative solution.

First, Bernoulli's equation states

$$P_1 + \tfrac{1}{2}\rho v_1{}^2 + \rho g h_1 = P_2 + \tfrac{1}{2}\rho v_2{}^2 + \rho g h_2 - h_L + h_p$$

Where P is pressure, ρ is fluid density, v is fluid velocity, g is gravity, h is height, h_L is head losses, and h_p is pump head. For a cooling loop, when points 1 and 2 are taken to be on either side of a pump mounted horizontally and with the same inlet and outlet pipe sizes, this equation simplifies to:

$$h_L = h_p$$

In other words, the pump must simply overcome the losses caused by fluid friction and compressibility effects. These losses can be split into major and minor losses, where major losses stem from the friction of the fluid against its piping and are calculated as:

114

$$h_{L,major} = f\left(\frac{l}{D}\right)\frac{v^2}{2g}$$

Here f is the friction factor, l is the length of the piping, and D is the diameter of the piping. This value is multiplied by ρg to give it units of pressure commensurate with the form of Bernoulli's equation written above.

Minor losses are calculated for individual features that cause flow irregularities, such as elbows, nozzles, and valves. The cooling systems of the motor and inverter are fairly straight outside of their internal channels. The battery cooling plate, however, has numerous elbows whose effect on the system pressure drop was calculated as

$$h_{L,major} = k_{elbow}\frac{v^2}{2g}$$

This value, as well, must be multiplied by ρg for units of pressure.

Finally, the internal channels of the radiators, motor, and inverter were accounted for using manufacturer data.

These losses were summed and EES was able to find where the system would operate based on manufacturer pump curves.

CHAPTER 3: MIL/SIL/HIL DEVELOPMENT PROCESS

The advent of electronic control units (ECUs), computers of varying size and capability used in embedded systems to control one or more functionalities of the system, has led to entirely different system development cycles as compared to previous eras. This section will discuss the effect of the semiconductor industry on system control, specifically for automobiles, and how the Georgia Tech EcoCAR 3 team follows modern industry practices for system controls code testing.

It should be noted that the word "controls" in this context refers to the same idea but a different application than what is often discussed in mechanical engineering circles. "Controls" in both contexts refers to commanding a system with a desired response in mind, but in this thesis the term refers to doing so with programmed computer logic, not to the engineering discipline of analyzing a system's physical properties and behavior to understand its response to physical stimuli of varying magnitudes or frequencies.

Before the semiconductor revolution enabled designers to program complex logic with any number of inputs and outputs onto compact memory and processing devices, large and dynamic systems had to be controlled by their human operators with physical connections to the component that was to be controlled. For example, earlier generations of airplanes were controlled, taking the example of changing wing flap angle, by physical cabling actuated by manual levers and acting upon various pulleys and levers that moved the flaps. Numerous physical backups were in place to ensure that core functionality and safety weren't lost in the event of a cable snapping, adding significant weight in an application where weight was highly undesirable. To test the efficacy of controlling an airplane often required experimenting with the

actual setup to be used, which could be very expensive and dangerous, especially in the event of component failure.

Modern such systems are controlled in an entirely different manner. Instead of a manual lever actuating a heavy cable that tugs on another lever to move a component, the connection is a comparably lightweight link between two computers. A human input to one computer, for example the press of a button or the depression of a pedal, can be translated by the computer into a command, logic can be applied to it, and a digital signal can be sent to the other computer. The second computer then applies a voltage to cause whatever component behavior is desired. For example, in modern vehicles, pushing the start button sends a signal to one ECU, often called the Ignition Control Module or sometimes the Body Control Module, which applies logic to it, namely making sure that the vehicle is in Park or Neutral. If so, a signal is sent to the Engine Control Module telling it to start the engine. The Engine Control Module interprets this command and applies a voltage to the engine's starter motor, starting the vehicle. Alternatively, a command can go straight into an ECU and then to a component, as in the case of a car's accelerator pedal. Signals from this driver input usually go straight into the Engine Control Module, which then opens up the throttle valve to allow more air to flow into the engine.

These advances, of course, do not obviate the need to rigorously test system controls, but they do make such testing much more cost-effective and safe. When combining these controls advances with the ability for modern computers to simulate complex systems like vehicle powertrains, a new paradigm for controls development emerges known as xIL, or "x in the Loop." In this EcoCAR 3 competition, the stages of this process that are highlighted are Model in the Loop (MIL), Software in the Loop (SIL), and Hardware in the Loop (HIL).

Model in the Loop (MIL)

MIL testing comprises the first stage of powertrain controls development and consists of the work described in the preceding section on vehicle modeling. A computer program, such as Simulink, is used to model the vehicle's physics and capabilities, and control logic is included to allow the system to be simulated under the commands it will receive in real life. In this stage of development, no thought is required on the organization or naming of signals beyond normal practices, the ability of an actual ECU to run the control logic programmed, or what-if scenarios testing the control logic's response to faults and adverse conditions. The purpose of controls code at this point is to give the system commands that will be acted on to ensure that the system model is accurate enough for further progress.

The Georgia Tech EcoCAR 3 team used this step in the process to model system performance and efficiency, as described earlier. The vehicle model was developed in such a way as to have interchangeable parts- different engine, transmission, electric motor, and battery models could be easily inserted to see the effect on these two metrics. These components were modeled as having identical inputs and outputs- every engine receives a torque request and outputs a fuel flow measurement, for example- even though the formats, signal names, and signals themselves will be different for the actual components, a discrepancy to be rectified in the next step of development.

Software in the Loop (SIL)

What is specifically ignored in the MIL stage of development is now required for the SIL stage. In the real world, the system being developed will have ECUs with actual inputs and outputs, and the supervisory controller will need to look for signals that are exactly as the various ECUs on the vehicle will send them. To convert the team's model to SIL, all control decisions, whether those modeled and performed in the various component ECUs or those in the supervisory

controller logic, were segregated using model subdivision. Next, each of these ECUs was augmented to send and receive the full list of signals that will be present in the actual vehicle with the chosen components. As an example of what this entails, at the MIL stage of development, the team wasn't concerned with the difference in, for example, gasoline and diesel engine operation. Torque requests were sent, and manufacturer data was used to calculate emissions and energy consumption for a given torque and speed for both types of engine, and that was the extent of engine communication and control modeled. However, the signals sent from the actual engines would be very different based on differences in their physical operation. In another example, major electrical parameters like voltage, current, and state of charge were passed from the battery subsystem to the supervisory control logic. These were passed from the Bosch and A123 battery models, and model structure was optimized for easy comparison between the two brands for the purpose of architecture selection. However, if the actual supervisory controller computer on the car looks for some signal named "Battery_Voltage" as it was named during MIL development, it will be disappointed- the A123 Battery Control Module used in our pack outputs that parameter under a completely different name (retained in compliance with non-disclosure agreements).

Thus the team model was filled out with the selected components' controllers and their various signal names and in the process was transformed from a relatively simple model of the vehicle's physical behavior into a much more complete one incorporating electronic and computational behavior as well. Now not only was the supervisory logic present, but some basic logic on component ECUs had to be reverse engineered as well (with this portion of the work only in its initial stages). It's worth mentioning that only those signals with impact on supervisory controller logic were included in this effort- a modern vehicle has thousands of signals being passed around multiple CAN busses and it does the team little good to model them all, such as, for example, exactly how the engine ECU (or Engine Control Module "ECM") modulates spark

timing with respect to engine speed. But many things were needed for eventual supervisory controller deployment that were not present at the MIL stage, one example being all the signals that the ECM needs to start the engine, such as signals from the transmission confirming its own safe-to-operate status. Because the engine will be electronically disconnected from the rest of the vehicle in our application in order to allow for full supervisory control of its operation, those transmission signals need to be taken into the supervisory control logic and specifically passed through to the engine, and this behavior must be present in the model.

With this level of fidelity, where not only the physical behavior of the vehicle but its electronic communications behavior is simulated, it is possible to do advanced algorithm testing through fault insertion. A "fault" occurs when something undesired happens- a communications wire is shorted to ground, for example, or a component goes haywire. Because these things can and do happen in real world systems, it is necessary to develop supervisory code to recognize and mitigate these occurrences. While it is impossible to predict everything that could go wrong, the things that are likely to go wrong are well known and can be caused, or "inserted" (thus "fault insertion"), in the model.

For example, the electric motor used in the team vehicle, the Parker GVM 210-150P, is controlled by a third-party inverter, the Sevcon GEN4S8. While Sevcon is a reputable brand whose products can be trusted, it is not out of the realm of possibility for a component in the inverter to somehow fail or for the control code in the inverter itself to not be adequately robust to the effects of various situations, such as high temperatures. In any case, if the motor were to start producing torque in excess of what is commanded, it is not enough to simply hope the driver notices and presses the brake pedal. The torque produced by the motor could potentially overcome the torque applied by the brakes, and even if it couldn't, that is still a situation which should be addressed immediately. In fact, by modern automotive standards, if a fault mitigation strategy for this fault has not been implemented, the vehicle should not be considered safe to

operate. So a mitigation strategy was developed for the team vehicle as follows, and this kind of work is indicative of what is added to the vehicle model and controls code for other such potential faults during the SIL process.

One way to recognize the situation is to continuously check that the current drawn by the motor is within a small margin of error of what it should be drawing at a given torque and speed, which can be known using basic electric machine parameters and equations or by referencing an efficiency lookup table. The correct supervisory controller response in this situation, then, is to disable high voltage power by opening the battery pack's contactors. With that logic and response included in the supervisory controller code, the fault, the motor producing excess torque, can be inserted in a number of ways depending on the way the motor model is constructed. For the team vehicle model, this specific fault was inserted by adding 10 amps to the current reported by the inverter. The supervisory control code was tested with and without the extra 10 amps reported to ensure that its detection mechanism and response were both successful.

In the greater world of xIL development outside this team's efforts, the preceding model enhancements and additions comprise a small part of the SIL stage, with a greater emphasis on turning one's "model" of supervisory logic (so called because it is thus far contained in a proprietary format, Simulink, that wouldn't run on almost any embedded controller and isn't speed- or memory-limited beyond the limitations of the host computer) into actual software through compiling. This step is important, as real-world code issues such as variable scope and naming, improper function calls, memory allocation and addressing errors, and more may not be caught until the software is compiled. In addition, the computer that will actually run the supervisory code, or the "target hardware," will almost certainly not possess the memory and speed capabilities of the host machine. For example, the model was initially developed on a professional desktop workstation with 12 CPU cores running at 3.5 GHz and could utilize up to

16 GB of memory. The ETAS ES910 supervisory controller in the vehicle, however, has a single-core 800 MHz processor [17] with 512 MB of RAM [18]. It is thus not guaranteed that logic developed on the host simulation computer will be able to run in real-time on the target hardware, and in fact this issue is a common problem if it isn't addressed at every stage of the development cycle.

That said, the Georgia Tech EcoCAR 3 team chose not to spend time on this aspect of SIL development, focusing instead on implementing vehicle signals and fault insertion. This was done largely because our target hardware could be used so easily in the next stage of model development, HIL.

Hardware in the Loop (HIL)

So far what has been left out of the process to develop efficient, robust code to control a vehicle powertrain has been the controlling computer itself, the target hardware whose limitations, quirks, and needs may be the most important factors in ensuring full functionality. The supervisory controller will be sending and receiving data over actual wires, and this transition from virtual to actual signals is extremely important to test, especially since some of those signals won't be data at all, but rather actual physical voltages. For example, in the team's vehicle the accelerator pedal's two potentiometers, producing a voltage and nothing more, will be plugged directly into the supervisory controller, which will be responsible for deciphering the driver's acceleration wishes solely from those voltages. Carrying out this level of testing in the actual vehicle will retain the unwanted risk and cost inherent to physical experimentation. But if the supervisory controller can be plugged into a wiring harness that will accurately replicate the environment of the actual vehicle, with both CAN signals and voltages coming in and out to interact with the controller's logic, then the control setup can be tested to a degree that is very close to actual in-vehicle testing. This setup is found in Hardware in the Loop simulators.

Hardware in the Loop Simulators are a class of specialized computers produced for industries, such as aviation and automotive, for whom live testing of controllers would be prohibitively expensive and/or dangerous. These computers run system models, just like the previously described host PCs, but also feature system-specific I/O capabilities such as electrical power supplies and CAN ports. A dSPACE Simulator Mid-Size HIL simulator, similar to the larger versions ("dSPACE Simulator Full-Size") that are used by GM and other automakers, was donated for use in this competition. The dSPACE simulator can be set up such that when its team-created wiring harness is plugged into the ETAS ES910 supervisory controller, the ETAS believes that it is in the actual vehicle because all CAN messages and actual voltages that it will see in the vehicle will be generated by the simulator. While external I/O can be added to the testing suite, including pedals, buttons, and switches, the dSPACE simulator is capable of producing all signals internally and sending them through wires to the supervisory controller. Software included in the sponsorship, dSPACE ControlDesk, allows for layouting and programming a user interface with virtual inputs and indicators so that a programmer can manually generate inputs and watch how their control code responds on the target hardware in real-time.

The usefulness of the xIL development model is found in its progression from system modeling through hardware in the loop controls testing without requiring the actual system at any step of the way (given the accurate and complete understanding of the system's functionality required to model it, of course). At each level, the time spent on issues that are secondary to logic development increases. At the MIL level, for example, it is easy to write logic that can be quickly tested without regard for the various concerns of embedded hardware implementation. If one's logic is accidentally flipped to cause the car to go backwards undesirably, this can be caught without any significant software compilation times or physical (and potentially destructive) testing. Moving to the SIL level allows programmers to increase functionality and catch another

group of possible errors without spending the time on flashing their compiled code to the target hardware or, again, potentially destructive testing. The control logic will be refined to include accurate variable names and complete I/O, and if developers chose to compile their logic to resource-limited software they can ensure that it will be able to run in real-time. And finally, in HIL, almost all elements of system control can be verified with a virtual source of real I/O, including its responses to physical fault insertion. When the supervisory control computer is installed in a vehicle after passing HIL, it can be known with a degree of certainty that safe and non-destructive testing of the vehicle is possible.

xIL Simulation on the Georgia Tech EcoCAR 3 Team

The Georgia Tech EcoCAR 3 team utilizes the xIL development process to quickly and efficiently develop safe controls code. In this section, the team's practices for doing so will be discussed.

In Simulink, a library object is a piece of a simulation that can be instanced in multiple files. Libraries are a good tool for smart development because if a file consists of library links to other files, multiple people may work on this file without their changes conflicting by just working on the relevant library-linked files instead. More concretely, if the vehicle body model is to be changed, for example by changing the total vehicle weight because of an added component, the vehicle control strategy can be changed at the same time on a different computer with no coordination between the two people working on the model. When the changes to the vehicle body and controls portions are pushed to the cloud, the full model file will also be updated with both sets of changes.

Library links are also useful for xIL development by allowing multiple versions of the model file to be used without requiring changes made in one to be manually propagated to another. On the Georgia Tech EcoCAR 3 team, this functionality is used to address the issues inherent to

switching between SIL and HIL simulation. The vehicle plant model and controls code are each

contained within their own Simulink blocks, or subsystems, and links to these subsystems

appear in numerous files for different purposes. When the vehicle is to be simulated in SIL, the

inputs and outputs of these two blocks must connect to each other in the same Simulink file, a

file simply named "Model", whose general layout is shown in Figure 54 below.
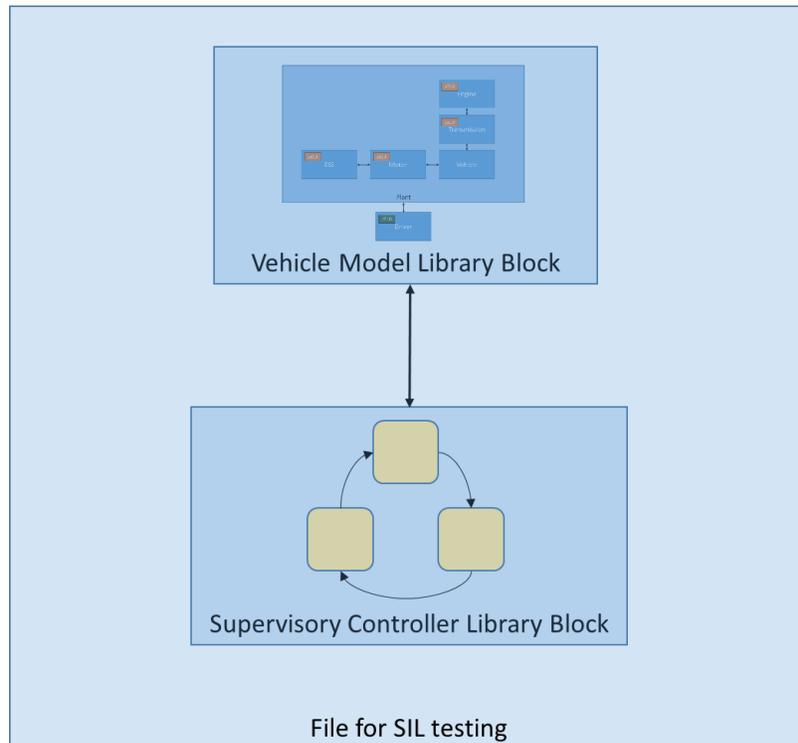
Figure 54: SIL Model File Layout

When it's time for HIL simulation, however, these two pieces of simulation must be

independently compiled onto separate devices, the physical model onto the dSPACE HIL

Simulator and the controls code onto the ETAS ES910 supervisory controller. In Simulink,

compiling to a device requires file-wide device-specific settings, meaning that compiling one file

to two different devices is difficult and onerous. The simplest solution is thus to have separate

files for ETAS and dSPACE, each containing library links to the corresponding features to be

compiled. In each of these files, simply named "ETAS" and "dSPACE," the inputs and outputs are connected to corresponding CAN message blocks so that the devices will be ready to communicate with each other on a shared CAN bus. This layout is shown in Figure 55 below.
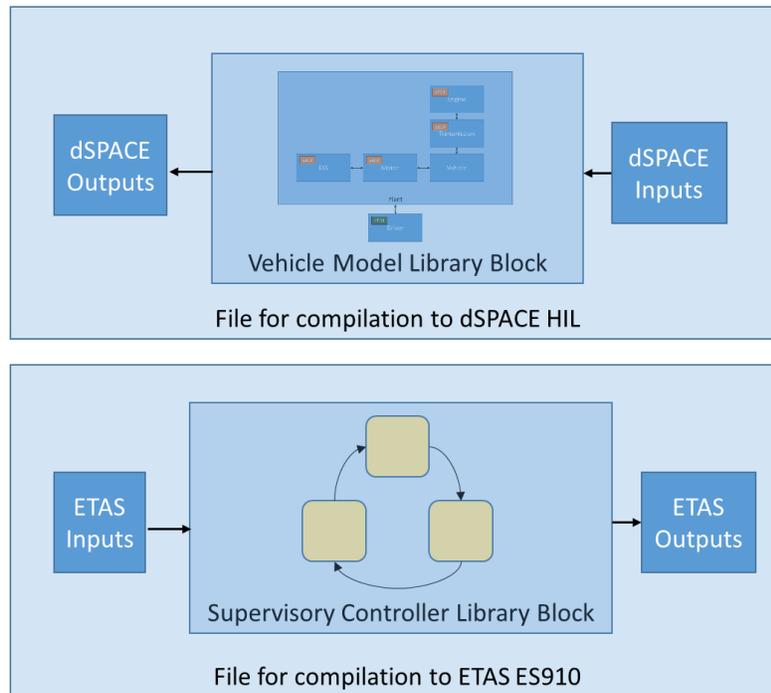
Figure 55: HIL Model File Layout

In terms of utilization of HIL-specific features, work is under way to complete the HIL harness through which the supervisory controller can communicate with what it will think is the vehicle. Simultaneous work in progress includes generating appropriate physical voltages for analog signals in the vehicle, such as the accelerator pedal and the transmission internal mode switch signals. Because the goal for the quickly upcoming Year 2 competition is to have a vehicle that simply passes the appropriate signals from the stock Camaro CAN bus to the LEA engine to run as a conventional engine-powered vehicle, with little or no electrical powertrain utilization, HIL simulation does not yet take priority over physical vehicle integration for the limited time resources available. However, the foundation has been laid for successful xIL simulation for the

final two years of the competition, ensuring safe and tested controls code and graduating

controls engineers with real-world development experience.

CHAPTER 4: ECOCAR 3 VEHICLE PROGRESS

As of this writing, the Georgia Tech EcoCAR 3 team is hard at work integrating the various components required to make a full-functioning hybrid Camaro. Because the distribution of points awarded at the upcoming Year 2 Competition are more heavily weighted towards integration completeness and quality than vehicle operation, the team's priority is to install as much as possible before spending too much time achieving even simple functionality, though the goal still remains to have simple driving enabled by the end of competition. In this section, the team's current progress will be detailed.

Engine Bay

The LEA engine and 8L45 transmission are currently installed in the vehicle, as shown.

Figure 56: LEA Engine in Camaro

Figure 57: Stock LGX Engine removed from Camaro

It was desired that the installation process require the removal of the fewest number of stock components possible, so once the stock LGX engine and transmission were removed, shown in Figure 57 above, the LEA and 8L45, bolted together and fully coupled (Figure 58), were inserted through the top of the engine bay together (Figure 59) rather than removing the front cradle and inserting from below.

130
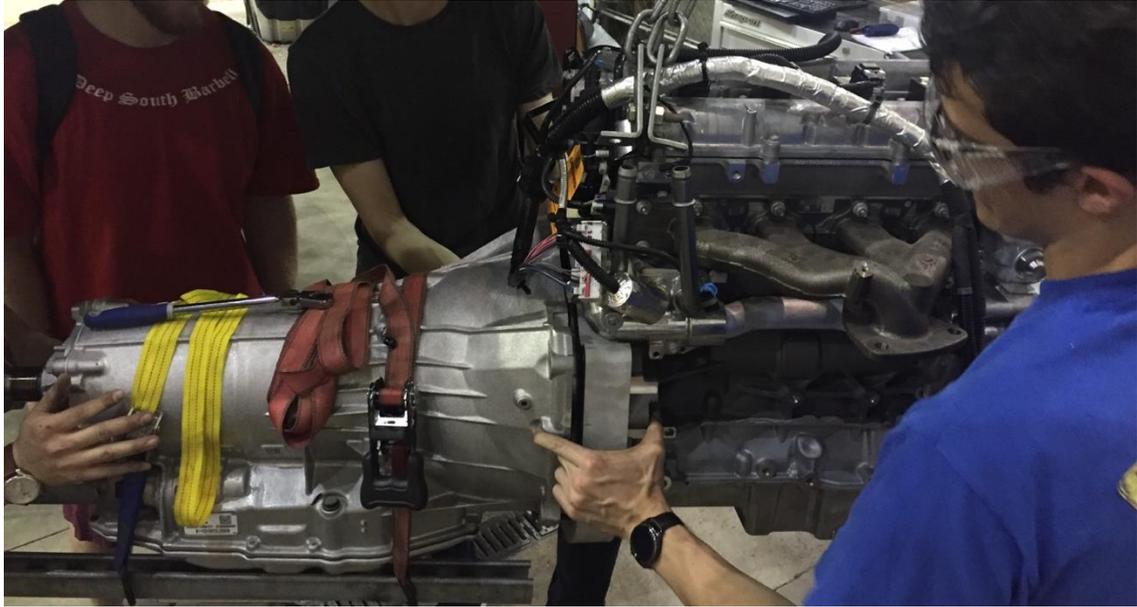
Figure 58: Mating the LEA and 8L45
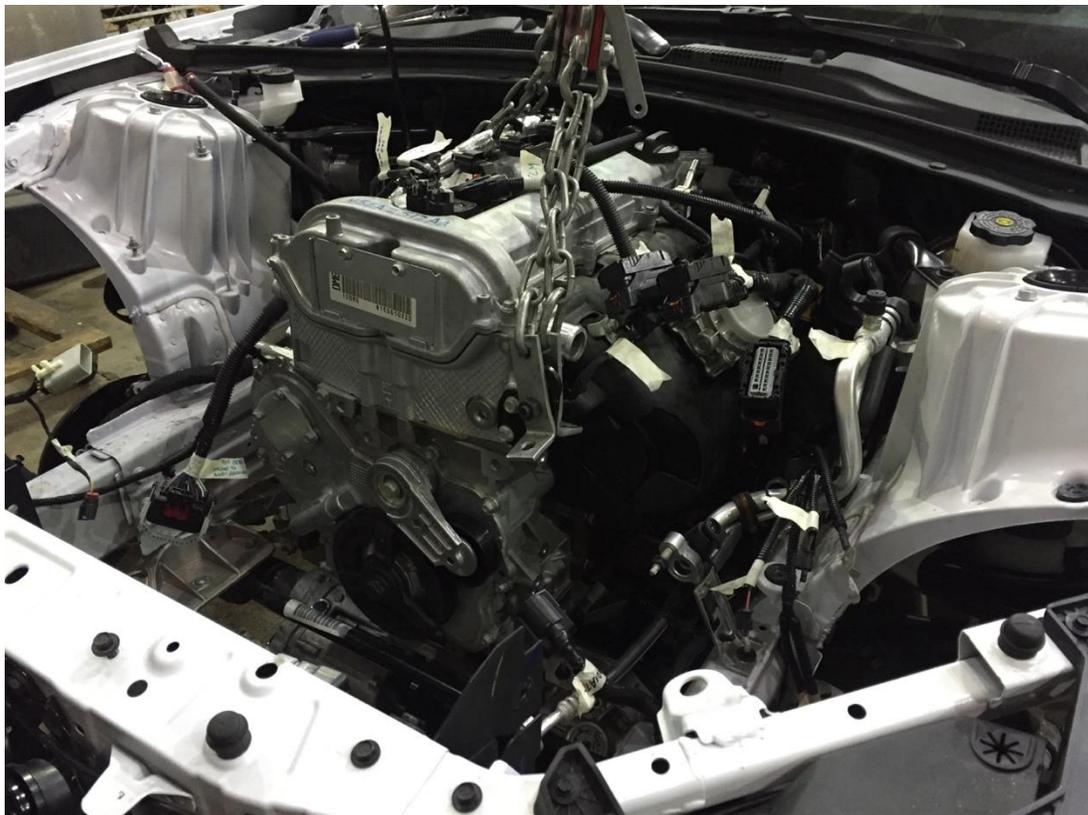
Figure 59: Inserting the LEA and 8L45 through the engine bay

Routing the various lines an engine requires to function has been a detailed task requiring matching of ports on the engine to features on the Camaro. First, the engine coolant lines were routed, with the engine outlet line shown curving in front of the engine in Figure 60 below.
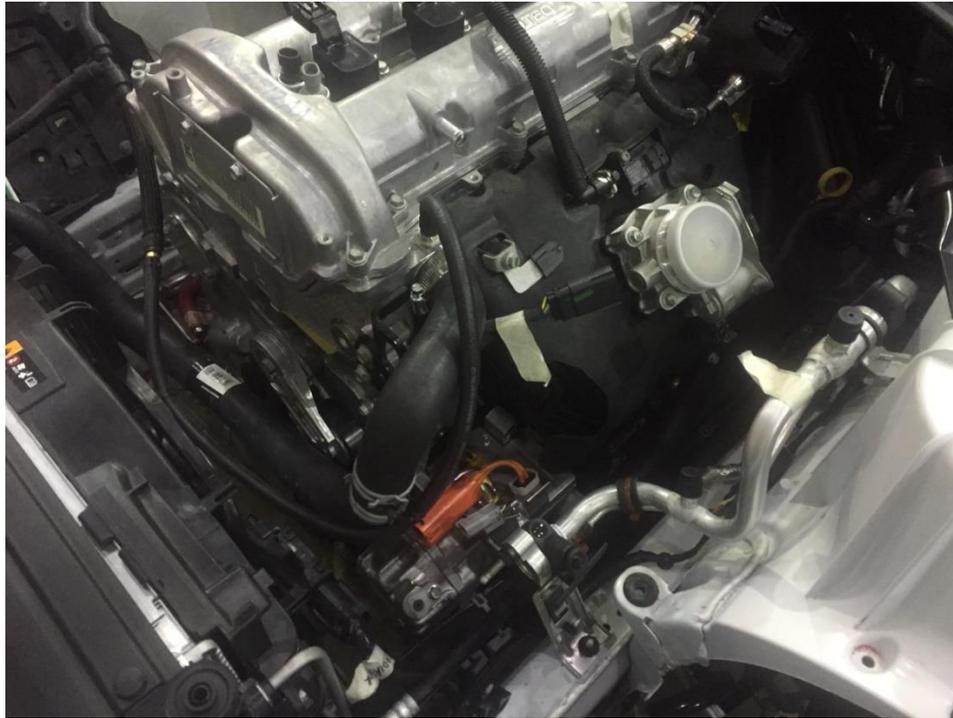
Figure 60: Engine coolant routing

Both the engine inlet and outlet lines were difficult to route because of the 90 degree rotation of the engine from its usual transverse applications to the team's longitudinal orientation. The engine inlet hose, in particular, required two lengths of custom-cut tubing and a third straight length to route from one corner of the engine block to the opposite corner.

The electrical harness required re-soldering of the wires coming from the Camaro to interface with the LEA's harness and connectors.

The new fuel lines were made of stainless steel to prevent ethanol-induced corrosion and routed from the Camaro's stock hard lines to the inlet on the engine. The evaporative emissions lines were routed in a similar way alongside them.

Figure 61: Exhaust manifold and routing as installed

The exhaust manifold that comes on the LEA was replaced with one from the Saturn Sky, one of the few vehicles to feature a longitudinal Ecotec engine and shown in Figure 61 above. This points the exhaust in a rearward direction, helpful from a routing perspective. The exhaust continues to meet with the stock Camaro routing until the first catalytic converter, shown in Figure 62.
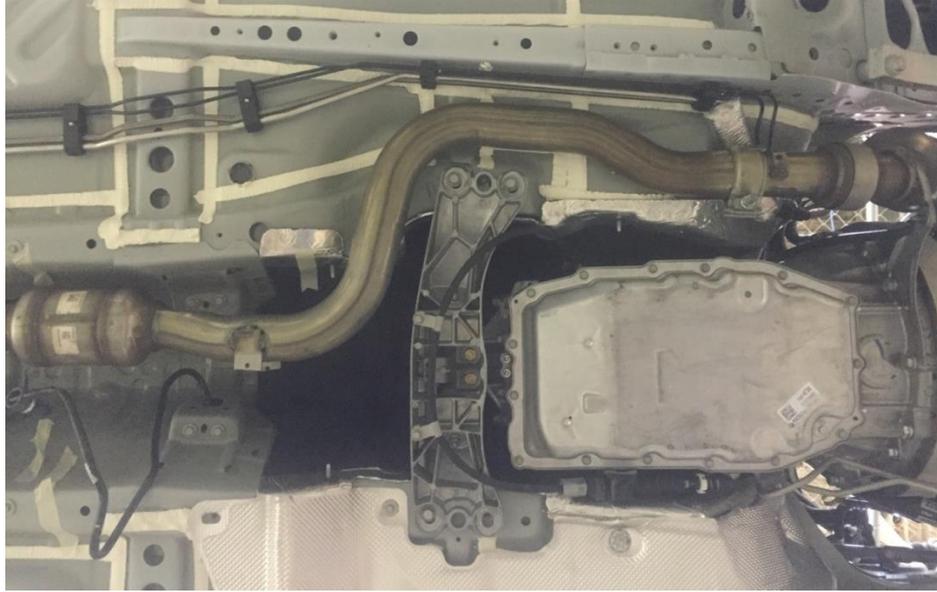
Figure 62: Front exhaust routing

The air intake is the same as that used on the 2015 Buick Regal, a vehicle that features the LEA and has overhead engine bay space that is similar to what's in the Camaro (almost none). This routing allows for intake from the stock air inlet.

Figure 63: Air Intake

The single vacuum line runs from the engine vacuum port to the brake booster by the master cylinder. This line is shown in Figure 64 below, curving through the middle of the image to the right of the yellow dipstick.

Figure 64: Vacuum Lines

Electric Drive

The motor and transfer case will be integrated soon using mounts that bolt into plates welded onto the vehicle. The mounts themselves are done and shown in Figure 65 below and are awaiting completion of the new driveshaft before final installation into the vehicle.

Figure 65: Motor (bottom) plugged into Transfer Case (top) and attached with middle MTC mount

The high voltage wires from the motor loop backwards and will go over the rear cradle before coming through the lower trunk floor to plug into the inverter.

Driveshaft

The driveshaft is currently being fabricated using the stock Camaro driveshaft shortened in sections, with pieces from the transmission used in the 2011 Cadillac CTS used to plug into the transfer case. The driveshaft will feature two sets of universal joints on the front segment and a single flex coupling on the much shorter and straighter rear segment.

Trunk

The trunk is currently in a mostly unfinished state of progress, with the battery pack half-

assembled outside of the car and very little mounted in the trunk so far. Progress on this area

will move quickly once all battery pack structural elements are completed.

CHAPTER 5: FUTURE WORK AND CONCLUSIONS

Figure 66: Members of the Georgia Tech EcoCAR 3 team after removing the stock engine

At the end of Year 2 of the EcoCAR 3 competition, the Georgia Tech team is hard at work

integrating a new hybrid powertrain and all its supporting components into the team Camaro.

The team has come a long way since the beginning of Year 1- for review, here is a summary of

the accomplishments detailed in this paper:

- P3 Parallel Hybrid Electric Vehicle architecture selected

- Components including engine, motor, and battery pack selected for maximum
  performance and efficiency

- Detailed vehicle model developed and improved for use in MIL, SIL, and HIL

- Battery pack designed and built for A123 module kit

- Structural and thermal analysis performed on numerous vehicle components

- Powertrain mounts and torque-carrying couplings designed and built

- Vehicle partially integrated and becoming more complete every day

To achieve the competition goal of having a consumer-ready vehicle by the end of Year 4, much work remains to be done, a sampling of which will be given here.

- Finish vehicle integration and achieve baseline functionality from every component

- Develop a controls strategy that optimally uses the engine and motor in tandem

- Reduce vehicle weight through the redesign of both team-fabricated and stock components

- Optimize cooling algorithms to reduce unnecessary pumping losses

- Further the development of the vehicle model and controls code for extensive and automated HIL testing

The team finishes Year 2 of the competition in the strongest position it has ever been, and through this competition the Georgia Tech EcoCAR 3 team is likely to be the source of numerous gifted and well-trained automotive engineers who will advance the state of the art towards more environmentally friendly and efficient modes of transportation.

APPENDIX 1: WEIGHT TRACKING DOCUMENT

Component	Mass (kg)	Longitudinal distance from front wheels (mm)	Notes
LGX	-157	164	
LEA	120	164	
Stock	1552	1346	
Fuel Tank	27.5	2155	
Thermal	33.4	2665	
Transfer Case	30	2032	
Motor	46	2283	
HV Wiring	20	2665	Estimated wiring weight
ESS	197	3028	
Pumps	7.91	93	Estimated pump location
Radiators	5.5	162	Estimated radiators location
DC/DC Converter	10	3621	
AC Compressor	9	125	Estimated weight, not including mounts
ESS Charger	6.3	3365	
Junction Box	0.95	3286	Weight is without internals
Rear Motor Mount	3.54	2441	
T.C. Primary Support Assembly	9.46	1929	
Secondary Support Assembly	1.05	1902	
Engine Mounts	17.39	164	
Stock engine mounts removal	-5.05	164	
Stock exhaust removal	-18.18	1346	Assumed to be distributed across CG
New exhaust	Unknown	2487	Midpoint distance
New Muffler	Unknown	3397	
Driveshaft removal	Unknown	1807	Midpoint of driveshaft
Front Driveshaft	Unknown	1482	
Rear Driveshaft	Unknown	2362	
ETAS controller	0.5	3303	Estimated mass and location TBD
12 V battery moved up front	-18	3351	Estimated mass
12 V battery moved from rear	18	93	Estimated mass
Flexplate	1.02	352	Rough location
Rim adapter	8.15	352	Rough location

APPENDIX 2: BATTERY COOLING EES FILE

{Solver for cooling of Georgia Tech EcoCAR 3's electronic components.
Points:
1: Entering component or component cooling plate
2: Exiting component or component cooling plate
3: Entering radiator
4: Exiting radiator
5: Entering pump
6: Exiting pump}

{--}
"Ambient and Heat Load Settings"
{--}

"Ambient Settings"
Airspeed=30 [mph]
"T_air=40 [C]"

"Battery Heat Load"
I_battery=120 [A]

"Which Radiator?"
A_radiator=2*A_M7B
UA_rad=2*UA_rad_M7B

"How many passes of cooling pipes?"
Num_Pipes=4

{--}
"Dimensions and Properties"
{--}

{-------Piping dimensions-------}

"Cooling plate"
Num_Elbows=Num_Pipes*2
Length_1=0.636 [m] "Cooling plate straight length"
Length_2=0.024 [m] "Elbow length"
Length_3=0.143 [m] "Distance between elbows"
Length_CoolingPlate=(Length_1*Num_Pipes) + (Length_2*Num_Elbows) + Length_3*((Num_Elbows/2)-1)

"Outside cooling plate"
Length_Outside=8 [m] "Rough estimate of back to front to back of car"

Length_Total=Length_CoolingPlate+Length_Outside
OD_pipe=0.0127 [m]
ID_pipe=0.010922 [m]
Length_Pipe=Length_1*Num_Pipes
A_Pipe=Length_Pipe*pi*ID_pipe

"Flow variables"
Velocity_Fluid=LPM*convert(L,m^3)*convert(sec,min)/(pi*ID_pipe^2/4)
m_dot=LPM*rho_1*convert(Liter,m^3)*convert(sec,min)

{-------Batteries-------}

"Battery Dimensions"
Num_modules=7
height_battery=.243 [m]
length_battery=.303 [m]
width_battery=.165 [m]
volume_battery=height_battery*length_battery*width_battery
A_battery=width_battery*length_battery
A_battery_sides=2 [1] *height_battery*length_battery
A_battery_frontback=2[1]*height_battery*width_battery

"Battery Thermal Properties"
h_sides=5 [W/(C-m^2)]

Thickness_ThermalPad=0.0005 [m]
Thickness_MountingPlate=0.003175 [m]
Thickness_Acrylic=0.22*.0254 [m]
Thickness_Trunk=.002 [m]

A_CoolingPlate=7*A_battery "0.4452 is previous input"
A_acrylic=1.5 [m^2]
A_battery_air=A_battery+A_battery_sides+A_battery_frontback

k_acrylic=0.2 [W/m-C]
k_Steel=50 [W/m-C]
k_Cu=385 [W/m-C]
k_ThermalPad=3 [W/m-C]

R_ThermalPad=Thickness_ThermalPad/(A_CoolingPlate*k_ThermalPad)
R_MountingPlate=Thickness_MountingPlate/(A_CoolingPlate*k_Steel)
R_total=R_ThermalPad+R_MountingPlate

"Unused:"
R_CopperPipe=ln(OD_pipe/ID_pipe)/(2*pi*k_Cu*Length_Pipe)
R_fluid=1/(h_plate*A_Pipe)
R_module_sides=2.0 [C/W] /(Num_Cells_Series*Num_Cells_Parallel)
R_module_bottom=4.3 [C/W] /(Num_Cells_Series*Num_Cells_Parallel)

h_acrylic=5 [W/C-m^2]
h_trunk=5 [W/C-m^2]
h_outside=50 [W/m^2-C]
R_trunk=(1/(h_sides*A_battery_air))+(1/(h_trunk*A_acrylic))+(Thickness_Acrylic/(k_acrylic*A_acrylic))+(1/(h_acrylic*
A_acrylic))+(1/(h_trunk*A_acrylic))+(Thickness_Trunk/(k_Steel*A_acrylic))+(1/(h_outside*A_acrylic))

"Heat Generation"
Num_Cells_Parallel=2
Num_Cells_Series=15
Resistance_cell=0.002 [ohm]
Resistance_pack=Num_Modules*Num_Cells_Series*Resistance_cell
Q_dot_gen=2*(I_battery/Num_Cells_Parallel)^2*Resistance_pack

{-------Physical Properties-------}
g=9.81 [m/s^2]
P_atm=101.325 [kPa]

T_1=T_4
T_3=T_2
T_5=T_4
T_6=T_4

cp_air_in=specheat(air_ha,T=T_air,P=P_atm)
cp_air_out=specheat(air_ha,T=T_airout,P=P_atm)

142

```
mu_1=viscosity(EG,C=50[%],T=T_1)
mu_2=viscosity(EG,C=50[%],T=T_2)
mu_3=viscosity(EG,C=50[%],T=T_3)
mu_4=viscosity(EG,C=50[%],T=T_4)
mu_5=viscosity(EG,C=50[%],T=T_5)
mu_6=viscosity(EG,C=50[%],T=T_6)

Pr_1=Prandtl(EG,C=50[%],T=T_1)
Pr_2=Prandtl(EG,C=50[%],T=T_2)
Pr_3=Prandtl(EG,C=50[%],T=T_3)
Pr_4=Prandtl(EG,C=50[%],T=T_4)
Pr_5=Prandtl(EG,C=50[%],T=T_5)
Pr_6=Prandtl(EG,C=50[%],T=T_6)

rho_1=density(EG,C=50[%],T=T_1)
rho_2=density(EG,C=50[%],T=T_2)
rho_3=density(EG,C=50[%],T=T_3)
rho_4=density(EG,C=50[%],T=T_4)
rho_5=density(EG,C=50[%],T=T_5)
rho_6=density(EG,C=50[%],T=T_6)

cp_EG_1=specheat(EG,C=50[%],T=T_1)
cp_EG_2=specheat(EG,C=50[%],T=T_2)
cp_EG_3=specheat(EG,C=50[%],T=T_3)
cp_EG_4=specheat(EG,C=50[%],T=T_4)
cp_EG_5=specheat(EG,C=50[%],T=T_5)
cp_EG_6=specheat(EG,C=50[%],T=T_6)

k_1=Conductivity(EG,C=50[%],T=T_1)
k_2=Conductivity(EG,C=50[%],T=T_2)
k_3=Conductivity(EG,C=50[%],T=T_3)
k_4=Conductivity(EG,C=50[%],T=T_4)
k_5=Conductivity(EG,C=50[%],T=T_5)
k_6=Conductivity(EG,C=50[%],T=T_6)

Re_1=rho_1*Velocity_Fluid*ID_pipe/mu_1
Re_2=rho_2*Velocity_Fluid*ID_pipe/mu_2
Re_3=rho_3*Velocity_Fluid*ID_pipe/mu_3
Re_4=rho_4*Velocity_Fluid*ID_pipe/mu_4
Re_5=rho_5*Velocity_Fluid*ID_pipe/mu_5
Re_6=rho_6*Velocity_Fluid*ID_pipe/mu_6


{--------------------------------------------------------------------------------}
"Fluid Dynamics"
{--------------------------------------------------------------------------------}
"Loss factors"
k_elbow=0.3
roughness=0.000015 [m] "high value for rubber hose"
1/(ff^.5)=-2*log10(((roughness/ID_pipe)/3.7)^1.11+(2.51/(Re_5*ff^.5))) "Equation 8.35a (Colebrook Equation)"

"Minor Losses"
Loss_Elbows=k_elbow*Velocity_Fluid^2/(2*g)
Extra_Elbows=4 [1] "Estimate of how many elbows will be required on the round trip to the radiator and back"
Losses_minor=(Num_Elbows+Extra_Elbows)*Loss_Elbows

"Major Losses"
Losses_major=ff*(Length_Total/ID_pipe)*(Velocity_Fluid^2/(2*g))

"Radiator Loss"
Pressure_Drop_Rad=1.6 [psi] "Manufacturer says this is maximum"
```

Losses_radiator=Pressure_Drop_Rad*convert(psi,Pa)/(rho_1*g)

"System Total"
Losses_Total=Losses_minor+Losses_major+Losses_radiator
PressureHead=Losses_Total*g*rho_1
PressureHead_psqi=PressureHead*convert(Pa,psi)

GPM=32.10 [(gal/min)/(psi^0)] -1.02 [(gal/min)/(psi^1)] *PressureHead_psqi +0.05 [(gal/min)/(psi^2)]
*PressureHead_psqi^2 -0.003 [(gal/min)/(psi^3)] *PressureHead_psqi^3
LPM=GPM*convert(gallon,liter)

{--}
"Heat Transfer"
{--}

Q_dot_trunk=(T_air-T_battery)/R_trunk
Q_dot_bottom+Q_dot_trunk+Q_dot_gen=0

"Resistances between modules and coolant"
Q_dot_plate=(T_battery-T_surf)/R_total
Q_dot_plate=-Q_dot_bottom

"Average of inlet and outlet properties"
Re_plate=(Re_1+Re_2)/2
Pr_plate=(Pr_1+Pr_2)/2
cp_EG_plate=(cp_EG_1+cp_EG_2)/2
k_plate=(k_1+k_2)/2

"Heat transfer within battery cooling plate"
Nusselt=((ff/8)*(Re_plate-1000)*Pr_plate)/(1+12.7*(ff/8)^(1/2)*(Pr_plate^(2/3)-1)) "Gnielinkski correlation"
h_plate=Nusselt*k_plate/ID_pipe
Q_dot_plate=m_dot*cp_EG_plate*(T_2-T_1)
T_2=T_surf-(exp(-(A_pipe*h_plate)/(cp_EG_plate*m_dot))*(T_surf-T_1));

"Use classic Q=hAT formula to check work. Technically the expression above, T_2=T_surf... is this one rearranged,
but this builds confidence"
CheckDT1=T_surf-T_2
CheckDT2=T_surf-T_1
CheckLMTD=(CheckDT1-CheckDT2)/ln(CheckDT1/CheckDT2)
CheckQ=h_plate*A_pipe*CheckLMTD

"Convective heat transfer to air"
P=(T_4-T_3)/(T_air-T_3)
R=(T_air-T_airout)/(T_4-T_3)
F=LMTD_CF('crossflow_both_unmixed',P,R)
Denominator_LMTD=(T_3-T_airout)/(T_4-T_air)
T_LMTD=((T_3-T_airout)-(T_4-T_air))/LN(Denominator_LMTD)
Q_dot_rad=-F*UA_rad*T_LMTD

"Average of inlet and outlet properties"
cp_air_rad=(cp_air_in+cp_air_out)/2
cp_EG_rad=(cp_EG_3+cp_EG_4)/2

"Mass flow equation for heat transfer to air"
m_dot_air=Airspeed*convert(mph,m/s)*A_radiator*density(air_ha,T=T_air,P=P_atm)
Q_dot_rad=m_dot*cp_EG_rad*(T_4-T_3)
Q_dot_rad=m_dot_air*cp_air_rad*(T_air-T_airout)

```
rho_ATF=860 [kg/m^3]
cp_ATF=1968 [J/kg C]
T_in_ATF=((170-32)/1.8) [C]
T_in_air=30 [C]
airspeed_manu=15 [mph]
LPM_ATF=11.36 [L/min]
m_dot_ATF=LPM_ATF*rho_ATF*convert(L,m^3)*convert(sec,min)


"B7B"
Width_B7B=11*convert(in,m)
Height_B7B=2.75*convert(in,m)
A_B7B=Width_B7B*Height_B7B
Q_B7B_BTU=7500 [BTU/hr]

Q_B7B=Q_B7B_BTU*convert(BTU/hr,W)
m_dot_air_B7B=airspeed_manu*convert(mph,m/s)*A_B7B*density(air_ha,T=T_in_air,P=P_atm)
T_out_ATF_B7B=T_in_ATF-Q_B7B/(m_dot_ATF*cp_ATF)
T_out_air_B7B=T_in_air+Q_B7B/(m_dot_air_B7B*specheat(air_ha,T=T_in_air,P=P_atm))
P_B7B=(T_out_air_B7B-T_in_air)/(T_in_ATF-T_in_air)
R_B7B=(T_in_ATF-T_out_ATF_B7B)/(T_out_air_B7B-T_in_air)
F_B7B=LMTD_CF('crossflow_both_unmixed',P_B7B,R_B7B)
T_LMTD_B7B_1=T_in_ATF-T_out_air_B7B
T_LMTD_B7B_2=T_out_ATF_B7B-T_in_air
T_LMTD_B7B=(T_LMTD_B7B_1-T_LMTD_B7B_2)/(LN(T_LMTD_B7B_1/T_LMTD_B7B_2))
UA_rad_B7B=Q_B7B/(F_B7B*T_LMTD_B7B)


"L7B"
Width_L7B=11*convert(in,m)
Height_L7B=5.75*convert(in,m)
A_L7B=Width_L7B*Height_L7B
Q_L7B_BTU=15000 [BTU/hr]

Q_L7B=Q_L7B_BTU*convert(BTU/hr,W)
m_dot_air_L7B=airspeed_manu*convert(mph,m/s)*A_L7B*density(air_ha,T=T_in_air,P=P_atm)
T_out_ATF_L7B=T_in_ATF-Q_L7B/(m_dot_ATF*cp_ATF)
T_out_air_L7B=T_in_air+Q_L7B/(m_dot_air_L7B*specheat(air_ha,T=T_in_air,P=P_atm))
P_L7B=(T_out_air_L7B-T_in_air)/(T_in_ATF-T_in_air)
R_L7B=(T_in_ATF-T_out_ATF_L7B)/(T_out_air_L7B-T_in_air)
F_L7B=LMTD_CF('crossflow_both_unmixed',P_L7B,R_L7B)
T_LMTD_L7B_1=T_in_ATF-T_out_air_L7B
T_LMTD_L7B_2=T_out_ATF_L7B-T_in_air
T_LMTD_L7B=(T_LMTD_L7B_1-T_LMTD_L7B_2)/(LN(T_LMTD_L7B_1/T_LMTD_L7B_2))
UA_rad_L7B=Q_L7B/(F_L7B*T_LMTD_L7B)

"M7B"
Width_M7B=11*convert(in,m)
Height_M7B=8*convert(in,m)
A_M7B=Width_M7B*Height_M7B
Q_M7B_BTU=20500 [BTU/hr]

Q_M7B=Q_M7B_BTU*convert(BTU/hr,W)
m_dot_air_M7B=airspeed_manu*convert(mph,m/s)*A_M7B*density(air_ha,T=T_in_air,P=P_atm)
T_out_ATF_M7B=T_in_ATF-Q_M7B/(m_dot_ATF*cp_ATF)
T_out_air_M7B=T_in_air+Q_M7B/(m_dot_air_M7B*specheat(air_ha,T=T_in_air,P=P_atm))
P_M7B=(T_out_air_M7B-T_in_air)/(T_in_ATF-T_in_air)
R_M7B=(T_in_ATF-T_out_ATF_M7B)/(T_out_air_M7B-T_in_air)
F_M7B=LMTD_CF('crossflow_both_unmixed',P_M7B,R_M7B)
```

T_LMTD_M7B_1=T_in_ATF-T_out_air_M7B
T_LMTD_M7B_2=T_out_ATF_M7B-T_in_air
T_LMTD_M7B=(T_LMTD_M7B_1-T_LMTD_M7B_2)/(LN(T_LMTD_M7B_1/T_LMTD_M7B_2))
UA_rad_M7B=Q_M7B/(F_M7B*T_LMTD_M7B)

"H7B"
Width_H7B=11*convert(in,m)
Height_H7B=11*convert(in,m)
A_H7B=Width_H7B*Height_H7B
Q_H7B_BTU=29200 [BTU/hr]

Q_H7B=Q_H7B_BTU*convert(BTU/hr,W)
m_dot_air_H7B=airspeed_manu*convert(mph,m/s)*A_H7B*density(air_ha,T=T_in_air,P=P_atm)
T_out_ATF_H7B=T_in_ATF-Q_H7B/(m_dot_ATF*cp_ATF)
T_out_air_H7B=T_in_air+Q_H7B/(m_dot_air_H7B*specheat(air_ha,T=T_in_air,P=P_atm))
P_H7B=(T_out_air_H7B-T_in_air)/(T_in_ATF-T_in_air)
R_H7B=(T_in_ATF-T_out_ATF_H7B)/(T_out_air_H7B-T_in_air)
F_H7B=LMTD_CF('crossflow_both_unmixed',P_H7B,R_H7B)
T_LMTD_H7B_1=T_in_ATF-T_out_air_H7B
T_LMTD_H7B_2=T_out_ATF_H7B-T_in_air
T_LMTD_H7B=(T_LMTD_H7B_1-T_LMTD_H7B_2)/(LN(T_LMTD_H7B_1/T_LMTD_H7B_2))
UA_rad_H7B=Q_H7B/(F_H7B*T_LMTD_H7B)

"This makes it work and I'm not sure why"
T3=T_3
T4=T_4

APPENDIX 3: MOTOR COOLING EES FILE

{Solver for the cooling of the Parker GVM210-150P motor
Points:
1: Before radiator/after motor
2: After radiator/through pump/before motor}

{---}
"Ambient and Heat Load Settings"
{---}
"T_air=40 [C]"
Airspeed=25 [mph]
A_radiator=A_L7B
UA_rad=UA_rad_L7B
Q_dot_heating= 84.1*((1/.959)-1)*1000 "From manufacturer data and using recommended sizing formula"

{---}
"Physical Parameters"
{---}
P_atm=101.325 [kPa]
g=9.81 [m/sec^2]

cp_air_in=specheat(air_ha,T=T_air,P=P_atm)
cp_air_out=specheat(air_ha,T=T_airout,P=P_atm)

mu_1=viscosity(EG,C=50[%],T=T_1)
mu_2=viscosity(EG,C=50[%],T=T_2)

k_1=Conductivity(EG,C=50[%],T=T_1)
k_2=Conductivity(EG,C=50[%],T=T_2)

Pr_1=Prandtl(EG,C=50[%],T=T_1)
Pr_2=Prandtl(EG,C=50[%],T=T_2)

cp_EG_1=specheat(EG,C=50[%],T=T_1)
cp_EG_2=specheat(EG,C=50[%],T=T_2)

rho_1=density(EG,C=50[%],T=T_1)
rho_2=density(EG,C=50[%],T=T_2)
Re_2=rho_2*Velocity_Fluid*ID_pipe/mu_2

{---}
"Heat Transfer"
{---}
"Convective heat transfer to air through radiator"
P=(T_airout-T_air)/(T_1-T_air)
R=(T_1-T_2)/(T_airout-T_air)
F=LMTD_CF('crossflow_both_unmixed',P,R)
T_LMTD=((T_1-T_airout)-(T_2-T_air))/LN((T_1-T_airout)/(T_2-T_air))
Q_dot_rad=-F*UA_rad*T_LMTD

"Average of inlet and outlet properties"
cp_air_rad=(cp_air_in+cp_air_out)/2
cp_EG_rad=(cp_EG_1+cp_EG_2)/2

"Mass flow equation for heat transfer to air"
mdot_air=Airspeed*convert(mph,m/s)*A_radiator*density(air_ha,T=T_air,P=P_atm)

Q_dot_rad=m_dot*cp_EG_rad*(T_2-T_1)
Q_dot_rad=mdot_air*cp_air_rad*(T_air-T_airout)

"Motor side to complete the loop"
Q_dot_rad=-Q_dot_heating

{--}
"Fluid Dynamics"
{--}

"General Parameters"
ID_pipe=0.010922 [m]
Velocity_Fluid=LPM*convert(L,m^3)*convert(sec,min)/(pi*ID_pipe^2/4)
m_dot=LPM*rho_1*convert(Liter,m^3)*convert(sec,min)

"Major Losses in piping"
Length_Total=6 [m] "Rough estimate of back to front to back of car"
roughness=0.000070 [m] "high value for rubber hose"
1/(ff^.5)=-2*log10(((roughness/ID_pipe)/3.7)^1.11+(2.51/(Re_2*ff^.5))) "Equation 8.35a (Colebrook Equation)"
Losses_major=ff*(Length_Total/ID_pipe)*(Velocity_Fluid^2/(2*g)) "[meters]"

"Motor Losses"
Motor_Head_bar = 0.0014 [bar/(L/min)^2] *LPM^2 + 0.008 [bar/(L/min)] *LPM "Curve fit from manufacturer data, [bar]"
Losses_motor=Motor_Head_bar*convert(bar,Pa)/(g*rho_1) "Convert to head in meters for adding to the others"

"Radiator Loss"
Pressure_Drop_Rad=1.6 [psi] "Manufacturer says this is maximum"
Losses_radiator=Pressure_Drop_Rad*convert(psi,Pa)/(rho_1*g)

"System Total"
Losses_Total=Losses_major+Losses_motor+Losses_radiator
PressureHead=Losses_Total*g*rho_1
PressureHead_psqi=PressureHead*convert(Pa,psi)

LPM=6.78513986 [L/min] - 0.126564977 [(L/min)/psi] *PressureHead_psqi + 0.00411429196 [(L/min)/psi^2]
*PressureHead_psqi^2 - 0.0000737106643 [(L/min)/psi^3] *PressureHead_psqi^3 + 6.28642191E-07 [(L/min)/psi^4]
*PressureHead_psqi^4 - 2.08333333E-09 [(L/min)/psi^5] *PressureHead_psqi^5
"Data for Shurflo 8000-543-138"

Note: the calculation of Trucool radiator performance is included at the end of this code but was

removed here since that section is identical to what is at the end of the Battery Cooling EES File

in Appendix 1.

APPENDIX 4: INVERTER COOLING EES FILE

{Solver for the cooling of the Sevcon GEN4S8 inverter
Points:
1: Before radiator/after inverter
2: After radiator/through pump/before inverter}

{---}
"Ambient and Heat Load Settings"
{---}
"T_air=40 [C]"
Airspeed=25 [mph]
A_radiator=A_B7B
UA_rad=UA_rad_B7B
Q_dot_heating=1642 "From manufacturer data"

{---}
"Physical Parameters"
{---}
P_atm=101.325 [kPa]
g=9.81 [m/sec^2]

cp_air_in=specheat(air_ha,T=T_air,P=P_atm)
cp_air_out=specheat(air_ha,T=T_airout,P=P_atm)

mu_1=viscosity(EG,C=50[%],T=T_1)
mu_2=viscosity(EG,C=50[%],T=T_2)

k_1=Conductivity(EG,C=50[%],T=T_1)
k_2=Conductivity(EG,C=50[%],T=T_2)

Pr_1=Prandtl(EG,C=50[%],T=T_1)
Pr_2=Prandtl(EG,C=50[%],T=T_2)

cp_EG_1=specheat(EG,C=50[%],T=T_1)
cp_EG_2=specheat(EG,C=50[%],T=T_2)

rho_1=density(EG,C=50[%],T=T_1)
rho_2=density(EG,C=50[%],T=T_2)
Re_2=rho_2*Velocity_Fluid*ID_pipe/mu_2

{---}
"Heat Transfer"
{---}
"Convective heat transfer to air through radiator"
P=(T_airout-T_air)/(T_1-T_air)
R=(T_1-T_2)/(T_airout-T_air)
F=LMTD_CF('crossflow_both_unmixed',P,R)
T_LMTD=((T_1-T_airout)-(T_2-T_air))/LN((T_1-T_airout)/(T_2-T_air))
Q_dot_rad=-F*UA_rad*T_LMTD

"Average of inlet and outlet properties"
cp_air_rad=(cp_air_in+cp_air_out)/2
cp_EG_rad=(cp_EG_1+cp_EG_2)/2

"Mass flow equation for heat transfer to air"
mdot_air=Airspeed*convert(mph,m/s)*A_radiator*density(air_ha,T=T_air,P=P_atm)

Q_dot_rad=m_dot*cp_EG_rad*(T_2-T_1)
Q_dot_rad=mdot_air*cp_air_rad*(T_air-T_airout)

"Inverter side to complete the loop"
Q_dot_rad=-Q_dot_heating

{--}
"Fluid Dynamics"
{--}

"General Parameters"
ID_pipe=0.010922 [m]
Velocity_Fluid=LPM*convert(L,m^3)*convert(sec,min)/(pi*ID_pipe^2/4)
m_dot=LPM*rho_1*convert(Liter,m^3)*convert(sec,min)

"Major Losses in piping"
Length_Total=8 [m] "Rough estimate of back to front to back of car"
roughness=0.000070 [m] "high value for rubber hose"
1/(ff^.5)=-2*log10(((roughness/ID_pipe)/3.7)^1.11+(2.51/(Re_2*ff^.5))) "Equation 8.35a (Colebrook Equation)"
Losses_major=ff*(Length_Total/ID_pipe)*(Velocity_Fluid^2/(2*g)) "[meters]"

"Inverter Losses"
Inverter_Head_bar = 0.0083 [bar/(L/min)^2] *LPM^2 - 0.0366 [bar/(L/min)] *LPM +0.2146 [bar] "Curve fit from manufacturer data, [bar]"
Losses_Inverter=Inverter_Head_bar*convert(bar,Pa)/(g*rho_1) "Convert to head in meters for adding to the others"

"Radiator Loss"
Pressure_Drop_Rad=1.6 [psi] "Manufacturer says this is maximum"
Losses_radiator=Pressure_Drop_Rad*convert(psi,Pa)/(rho_1*g)

"System Total"
Losses_Total=Losses_major+Losses_Inverter+Losses_radiator
PressureHead=Losses_Total*g*rho_1
PressureHead_psqi=PressureHead*convert(Pa,psi)

LPM=6.78513986 [L/min] - 0.126564977 [(L/min)/psi] *PressureHead_psqi + 0.00411429196 [(L/min)/psi^2] *PressureHead_psqi^2 - 0.0000737106643 [(L/min)/psi^3] *PressureHead_psqi^3 + 6.28642191E-07 [(L/min)/psi^4] *PressureHead_psqi^4 - 2.08333333E-09 [(L/min)/psi^5] *PressureHead_psqi^5
"Data for Shurflo 8000-543-138"

Note: the calculation of Trucool radiator performance is included at the end of this code but was removed here since that section is identical to what is at the end of the Battery Cooling EES File in Appendix 1.

REFERENCES

[1]   Argonne National Laboratories, "AVTC Competitions," 2016. [Online]. Available: http://avtcseries.org/competitions/.

[2]   Acura, "Acura RLX," [Online]. Available: http://www.acura.com/RLXLanding.aspx?model=RLX#~pIEPJqtmokbOBT. [Accessed 17 April 2016].

[3]   S. Collie, "Acura MDX gets sport hybrid system, gets a new beak," gizmag, [Online]. Available: http://www.gizmag.com/acura-mdx-sport-hybrid-nyias/42538/. [Accessed 17 April 2016].

[4]   Argonne National Laboratories, *Non-Year Specific Rules Rev G,* 2016.

[5]   GM Authority, "GM 8-Speed 8L45 M5N Hydra-Matic Automatic Transmission," [Online]. Available: http://gmauthority.com/blog/gm/gm-transmissions/m5n/.

[6]   GM Authority, "GM 8-Speed 8L90 M5U Hydra-Matic Automatic Transmission," [Online]. Available: http://gmauthority.com/blog/gm/gm-transmissions/m5u/.

[7]   Fastenal, "Torque-Tension Relationship for Metric Fasteners," [Online]. Available: https://www.fastenal.com/content/feds/pdf/Torque-Tension%20Chart%20for%20Metric%20Fasteners.pdf. [Accessed 17 April 2016].

[8]   Online Metals, "1018 Mild (low-carbon) steel," [Online]. Available: https://www.onlinemetals.com/productguides/alloycat.cfm?alloy=1018. [Accessed 17 April 2016].

[9]   Leeco Steel, "A514 Steel Plate," [Online]. Available: http://www.leecosteel.com/a514-steel-plate.html. [Accessed 2016 17 April].

[10]   Engineering Toolbox, "Rolling friction and rolling resistance," [Online]. Available: http://www.engineeringtoolbox.com/rolling-friction-resistance-d_1303.html. [Accessed 15 April 2016].

[11]   Everything Explained Today, "Hans B. Pacejka Explained," [Online]. Available: http://everything.explained.today/Hans_B._Pacejka/. [Accessed 16 April 2016].

[12]   MathWorks, "Rolling Resistance," [Online]. Available: http://www.mathworks.com/help/physmod/sdl/ref/rollingresistance.html#btk5yw8-1. [Accessed 2016 17 April].

[13]   T. L. A. I. F. D. D. Bergan, Fundamentals of Heat and Mass Transfer 7th Ed., Jefferson City: John Wiley & Sons, 2011.

[14]   MatWeb, "Aluminum, Al," [Online]. Available: http://www.matweb.com/search/DataSheet.aspx?MatGUID=0cd1edf33ac145ee93a0aa6f c666c0e0&ckck=1. [Accessed 16 April 2016].

[15]   Matweb, "Copper, Cu; Cold Drawn," [Online]. Available: http://www.matweb.com/search/DataSheet.aspx?MatGUID=1980eb23287a4408adc404d d39293942. [Accessed 16 April 2016].

[16]   Beck's Auto Center, "Transmission," [Online]. Available: http://becksautocenter.com/transmission/. [Accessed 4 April 2016].

[17]   NXP, "MPC8548E: PowerQUICC III Processor with DDR2, PCI, PCI Express®, Serial RapidIO, SerDes, 1 GB Ethernet, Security," NXP, [Online]. Available: http://www.nxp.com/products/microcontrollers-and-processors/power-architecture-processors/powerquicc-processors/powerquicc-iii-85xx/powerquicc-iii-processor-with-ddr2-pci-pci-express-serial-rapidio-serdes-1-gb-ethernet-security:MPC8548E. [Accessed 16 April 2016].

[18]   ETAS, "ES910 – Prototyping and Interface Module," [Online]. Available: http://www.etas.com/en/products/es910_rapid_prototyping_module-technical_data.php. [Accessed 16 April 2016].