

Design and Control of a Unique Hydrogen Fuel Cell Plug-In Hybrid Electric Vehicle

by
Michael Giannikouris

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Applied Science
in
Mechanical Engineering

Waterloo, Ontario, Canada, 2013

© Michael Giannikouris 2013

Author's Declaration

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

The University of Waterloo Alternative Fuels Team (UWAFT) is a student team that designs and builds vehicles with advanced powertrains. UWAFT uses alternatives to fossil fuels because of their lower environmental impacts and the finite nature of oil resources. UWAFT participated in the EcoCAR Advanced Vehicle Technology Competition (AVTC) from 2008 to 2011. The team designed and built a Hydrogen Fuel Cell Plug-In Hybrid Electric Vehicle (FC-PHEV) and placed 3rd out of 16 universities from across North America.

UWAFT design projects offer students a unique opportunity to advance and augment their core engineering knowledge with hands-on learning in a project-based environment. The design of thermal management systems for powertrain components is a case study for design engineering which requires solving open ended problems, and is a topic that is of growing importance in undergraduate engineering courses. Students participating in this design project learn to develop strategies to overcome uncertainty and to evaluate and execute designs that are not as straightforward as those in a textbook. Electrical and control system projects require students to introduce considerations for reliability and robustness into their design processes that typically only focus on performance and function, and to make decisions that balance these considerations in an environment where these criteria impact the successful outcome of the project. The consequences of a failure or unreliable design also have serious safety implications, particularly in the implementation of powertrain controls. Students integrate safety into every step of control system design, using tools to identify and link together component failures and vehicle faults, to design detection and mitigation strategies for safety-critical failures, and to validate these strategies in real-time simulations.

Student teams have the opportunity to offer a rich learning environment for undergraduate engineering students. The design projects and resources that they provide can significantly advance student knowledge, experience, and skills in a way that complements the technical knowledge gained in the classroom. Finding ways to provide these experiences to more undergraduate students, either outside or within existing core courses, has the potential to enhance the value of program graduates.

Acknowledgements

I owe a lot to the many people who helped me during my years of graduate studies at the University of Waterloo. First, I want to thank Dr. Roydon Fraser, who encouraged me to go “hang around” the UWAFt garage when I first started my degree. Walking into that garage started something that has changed my life and helped me to find work that I am passionate about. I have always found motivation in his enthusiasm and confidence in his encouragement. I also want to thank Dr. Michael Fowler for the support that he tirelessly provides to help UWAFt students achieve their goals. He is constantly looking for ways to help the team and we benefitted greatly from the knowledge and resources that he provided. Dr. Fraser and Dr. Fowler are completely dedicated to student learning and student teams, and I cannot overstate how incredible it is to have the support of people who are always ready to listen, to give honest advice, and go to bat for you when you need it.

I am also grateful for all of the support that we received from EcoCAR organizers and sponsors. Kristen and her team at Argonne National Laboratories worked very hard to organize all of EcoCAR’s events, workshops, and competitions. They make it possible for so many students to get this unique and amazing opportunity that often defines their future careers, and I feel very fortunate to have been a part of it. Engineers at General Motors, including Dan Mephram and Paul Shoytush, gave our team a tremendous amount of technical support. Without their unwavering dedication and support throughout all of our technical missteps we would never have been able to complete our vehicle. I want to thank them for the countless hours they put in to help us succeed.

Finally, I want to thank all of the UWAFt team members, past and present, that I’ve had the privilege of working with over the past several years. Alex and Hung were great friends and mentors during my first year with the team, and I want to thank them for everything they taught me. I also want to thank Chris, Carl, Mark, Josh, Mike, Jason, Gurhari, Dan, Eric, and everyone else that worked with UWAFt during EcoCAR. I am proud and very lucky to have been part of such an incredible team.

Dedication

To Allyson, for her love and support in this and all my endeavours.

Table of Contents

Author’s Declaration.....	ii
Abstract.....	iii
Acknowledgements.....	iv
Dedication.....	v
List of Figures.....	ix
List of Tables.....	xii
Chapter 1 Introduction.....	1
1.1 Thesis Objective.....	3
1.2 Thesis Outline.....	4
Chapter 2 Literature Review and Background.....	7
2.1 Motivation for Alternative Fuels.....	7
2.2 Electrification of the Automotive Powertrain.....	10
2.2.1 Gasoline-Electric Hybrids.....	10
2.2.2 Battery Electric Vehicles (BEVs).....	16
2.2.3 Fuel Cells.....	17
2.3 Fuel Cells in Automotive Applications.....	17
2.3.1 Fuel Cell Technologies.....	17
2.3.2 Current State of the Automotive Fuel Cell Powertrain.....	21
2.4 Evaluating the Emissions and Energy Consumption of Electrified Powertrains.....	23
2.4.1 Fuel Economy Evaluation for Conventional Powertrains.....	24
2.4.2 The Impact of Electrified Powertrains.....	26
2.4.3 Plug-In Powertrains, the Utility Factor, and New Testing Standards.....	28
2.5 Project-Based Learning.....	32
Chapter 3 The EcoCAR Hydrogen Fuel Cell Plug-in Hybrid Electric Vehicle.....	38
3.1 Powertrain Architecture.....	39
3.1.1 Fuel Cell System.....	41
3.1.2 Hydrogen Storage System.....	43
3.1.3 Front Electric Traction System.....	43

3.1.4 Rear Electric Traction System	44
3.1.5 Lithium-Ion Energy Storage System.....	45
3.1.6 DC-DC Power Converter.....	46
3.1.7 Battery Charger	47
3.1.8 Auxiliary Components.....	48
3.2 Powertrain Control System.....	49
3.2.1 Control System Hardware.....	49
3.2.2 Control System Software	50
Chapter 4 Thermal Management System Design	54
4.1 System Design Process	55
4.2 The Low-Temperature Front Cooling Loop	57
4.3 Component Arrangement.....	58
4.4 Specify Conditions	58
4.5 Component Pressure Drop Equations	61
4.6 Loop Temperature Equations.....	62
4.7 System Pressure Drop and Pump Selection	64
4.8 Sensitivity Analysis	66
4.9 Final Design Summary	67
4.10 Thermal System Design in a Project-Based Learning Environment	68
Chapter 5 Powertrain Control System Design.....	71
5.1 High Voltage Architecture.....	72
5.1.1 Bus Voltage Ratio.....	74
5.1.2 Battery Pack Voltage Modification.....	75
5.1.3 Rear Motor Re-Location.....	78
5.1.4 Control System Implications.....	79
5.2 Powertrain Modes of Operation.....	80
5.2.1 High Voltage Bus Pre-Charge	81
5.2.2 Fuel Cell Stack Start	81
5.2.3 Propulsion – Excess High-Side Power.....	82
5.2.4 Propulsion – Excess Low-Side Power	82
5.2.5 Regenerative Braking	83
5.3 Power Control Algorithm	83

5.4 Torque Control Algorithm	88
5.5 Control System Design in Project-Based Learning Environment	96
Chapter 6 Safety Control System Design	98
6.1 Frameworks for Safety Analysis.....	98
6.2 Fault Identification in the Torque Control System.....	101
6.2.1 Top-Down Safety Analysis: FTA	103
6.2.2 Bottom-Up Safety Analysis: FMEA	106
6.2.3 Top-Down vs. Bottom-Up	108
6.3 Fault Detection and Mitigation: Implementation and Validation Using HIL	108
6.4 Validation Using Hardware-in-the-Loop Testing	112
6.5 When the Analysis Fails: Mishandling the ETS Torque Enable Signal	116
6.6 Design for Safety in a Project-Based Learning Environment	121
Chapter 7 On-Road Vehicle Testing.....	122
7.1 Acceleration.....	123
7.2 Braking	126
7.3 Energy Efficiency	126
7.4 Range.....	130
7.5 Petroleum Energy Use and Greenhouse Gas Emissions	131
Chapter 8 Conclusions and Recommendations.....	132
8.1 Conclusions	132
8.2 Recommendations	136
References.....	138
Appendix A Tabulated Comparison of Fuel Cell Technologies.....	141
Appendix B Underbody View of Final Vehicle Showing Powertrain Integration	142
Appendix C Front Cooling Loop Thermal-Fluid Calculations.....	143
Appendix D EcoCAR Powertrain High Voltage Schematic.....	153
Appendix E Power Control Algorithm Top-Level Overview.....	154
Appendix F DC-DC Component Control Block for Driving Mode	155
Appendix G FMEA for Accelerator Pedal Sensor.....	156

List of Figures

Figure 1: Energy use in Canada in petajoules, 2008, breakdown by sector, generated using data from Natural Resources Canada’s Energy Use Handbook [6]	8
Figure 2: GHG emissions in Canada in megatonnes of CO ₂ equivalent, 2008, breakdown by sector, generated using data from Natural Resources Canada’s Energy Use Handbook [6].....	9
Figure 3: Series hybrid powertrain	12
Figure 4: Parallel hybrid powertrain	13
Figure 5: Series-parallel hybrid powertrain	15
Figure 6: A fuel cell powertrain with DC-DC converter on the fuel cell side of the bus	21
Figure 7: UDDS drive schedule driven by UWAFT EcoCAR at EPA National Vehicle & Fuel Emissions Laboratory.....	25
Figure 8: Utility factor curve generated using EcoCAR competition formula given in [18] .	30
Figure 9: UWAFT EcoCAR year 2 powertrain architecture	39
Figure 10: Fuel cell system, front view	42
Figure 11: Fuel cell system, rear view	42
Figure 12: Hydrogen storage system in its delivery crate.....	43
Figure 13: Front electric motor	44
Figure 14: Rear electric motor integrated into the rear sub-frame assembly.....	45
Figure 15: UWAFT designed and built lithium-ion battery pack, final CAD assembly with team-designed DC-DC (left) and installed in vehicle (right).....	46
Figure 16: Brusa BDC412 DC-DC converter [42]	47
Figure 17: Brusa NLG513-Sx battery charger.....	48
Figure 18: Vehicle-wide CAN bus topology	50
Figure 19: UWAFT EcoCAR supervisory control schematic	51
Figure 20: Vehicle thermal management system.....	54
Figure 21: Thermal management system design process	57
Figure 22: Coolant inlet temperature for each loop component as a function of flow rate	63

Figure 23: System pressure drop and pump curves	64
Figure 24: Loop temperature plot for PMD, PIM, DU arrangement	66
Figure 25: Transition from the year 2 to year 3 architecture	73
Figure 26: DC-DC modes of operation.....	74
Figure 27: UWAFT EcoCAR source polarization potential, 330 V battery configuration	76
Figure 28: UWAFT EcoCAR source polarization potential, 250 V battery configuration	78
Figure 29: Current flow within the powertrain.....	80
Figure 30: Calculation flow for the power command algorithm	84
Figure 31: Fuel cell power desired.....	85
Figure 32: Fuel cell internal current requirement during regenerative braking.....	85
Figure 33: Power control algorithm DC-DC current calculation.....	86
Figure 34: Fuel cell power request calculation.....	87
Figure 35: Torque control algorithm.....	90
Figure 36: Over current condition on DC bus leading to DCDC shutdown	91
Figure 37: Torque limit and shift algorithm	93
Figure 38: Torque shifting algorithm allows “limp home” on rear motor.....	95
Figure 39: Control path for torque actuation in the EcoCAR vehicle	102
Figure 40: Fault tree for unintended torque delivery	104
Figure 41: Fault tree for supervisory controller erroneous torque request	105
Figure 42: Accelerator pedal sensor block for EcoCAR vehicle	110
Figure 43: Accelerator pedal out of range checks	110
Figure 44: Accelerator pedal redundancy check.....	111
Figure 45: Traction_Allowed gate using APS_OK signal.....	111
Figure 46: Simulator hardware connected to the supervisory controller and PC interface ..	112
Figure 47: PC interface for the EcoCAR HIL setup	113
Figure 48: Simulation of a short to power for accelerator pedal signal 2.....	114
Figure 49: Results of the brake pedal wins logic executing on the simulator	116
Figure 50: Powertrain voltage plots leading to ETS failure event.....	117
Figure 51: Effect of floating the ETS torque enable signal on bus high side voltage	119

Figure 52: Speed, power, and torque for 0-60 mph acceleration and braking test	125
Figure 53: E&EC drive cycle, schedule A.....	127
Figure 54: Vehicle underbody view showing powertrain integration	142
Figure 55: EcoCAR powertrain high voltage schematic	153
Figure 56: Power control algorithm top-level view	154
Figure 57: DC-DC component control block for "driving" mode	155

List of Tables

Table 1: Fuel energy density figures for EcoCAR competition energy storage [18], [19].....	17
Table 2: EcoCAR powertrain component specifications.....	40
Table 3: Thermal management specifications for front loop components [45][46][47]	58
Table 4: Physical properties for water and ethylene glycol (50%) at 70°C [48]	59
Table 5: Properties of pipe system at common diameters	60
Table 6: Properties of air at various temperatures [48].....	60
Table 7: Minimum coolant flow rate at various radiator air flow rates	67
Table 8: Front loop final performance	67
Table 9: Vehicle performance metrics evaluated during on-road testing.....	122
Table 10: Competition fuel data	123
Table 11: Emissions and Energy Consumption Event - Drive Schedule Statistics	127
Table 12: Emissions and Energy Consumption Event – Vehicle Energy Use	128
Table 13: Vehicle usable energy reductions in year 3	130
Table 14: Comparison of fuel cell technologies [54].....	141
Table 15: 50/50 ethylene glycol/water properties at 70°C [48].....	143
Table 16: Pipe properties, friction factor from White [55].....	143
Table 17: Low temperature cooling loop thermal and fluid specifications	143
Table 18: Pressure drop equations	144
Table 19: Low temperature cooling loop pressure drop calculations	145
Table 20: Heat transfer equations	146
Table 21: Low temperature cooling loop temperature calculations for loop order FPIM, DU, PMD, radiator air flow rate 4857 [CFM]	147
Table 22: Low temperature cooling loop temperature calculations for loop order PMD, FPIM, DU, radiator air flow rate 694 [CFM].....	148
Table 23: Low temperature cooling loop temperature calculations for loop order PMD, FPIM, DU, radiator air flow rate 2082 [CFM].....	149

Table 24: Low temperature cooling loop temperature calculations for loop order PMD, FPIM, DU, radiator air flow rate 3469 [CFM].....	150
Table 25: Low temperature cooling loop temperature calculations for loop order PMD, FPIM, DU, radiator air flow rate 4857 [CFM].....	151
Table 26: Low temperature cooling loop temperature calculations for loop order PMD, FPIM, DU, radiator air flow rate 6245 [CFM].....	152
Table 27: FMEA of the accelerator pedal sensor.....	156

Chapter 1

Introduction

Automotive powertrains are evolving at a rapid rate. While the efficiency and emissions of conventional gasoline and diesel powered vehicles continue to improve, consumers are now starting to see departures from the conventional combustion vehicle, with gasoline-electric hybrids and fully electric vehicles being offered by many original equipment manufacturers (OEMs). Hydrogen vehicles are also available in limited markets. Barriers to commercialization of hydrogen vehicles currently include their higher relative cost, level of reliability and refinement, public acceptance of hydrogen as a fuel, and the limited availability of hydrogen fuelling infrastructure [1]. Nevertheless, some original equipment manufacturers (OEMs) will have a commercial fuel cell vehicle in the market in 2015.

OEMs are continuing to develop new ways to meet ever more stringent fuel economy standards while maintaining the overall experience that consumers have come to expect from car ownership. Vehicle performance, range, interior space, purchasing cost, maintenance costs, and long-term reliability are all considerations that factor heavily into a buyer's decision making process. It can be challenging to meet buyers' expectations in these areas when developing a vehicle powered by alternative fuels.

The University of Waterloo Alternative Fuels Team (UWAFT) is a student organization that designs and builds vehicles powered by alternative fuels. The team aims to give engineering students practical experience in automotive engineering, with the unique opportunity to work with cutting-edge powertrain technology. The team builds its vehicles with the aim to demonstrate how current technology can provide alternative fuels powertrains that still meet the demands of automotive consumers. The team pushes the boundaries of current technology (sometimes developing its own in the process) to demonstrate to consumers the exciting new vehicles that await them in the coming years. Students are also able to learn about the business and consumer aspects of the automotive world, performing fundraising, outreach, and educational activities to raise the profile of emerging powertrains and help consumers better understand their workings, benefits, and limitations.

To enable these goals, the team has a history of participation in Advanced Vehicle Technology Competitions (AVTCs) run by Argonne National Laboratories (ANL) in the United States. These competitions involve a select group of universities from across North America, providing a framework that enables students to learn about, work with, and contribute to the latest powertrain technology. AVTCs are typically sponsored by a major OEM that provides each team with a vehicle to work on, some powertrain components, and engineering support. The competition also attracts several tier 1 and tier 2 suppliers that provide components and support at no cost or at a reduced price. With this support teams select their own powertrain architectures and spend three years designing, building, and testing their concepts. Yearly competitions are held where the teams and their vehicles gather to participate in events that measure each vehicle's performance, fuel economy, environmental impact, and consumer acceptability. This is not an easy task, as teams must engineer their vehicles to not only exceed the original fuel economy and environmental ratings, but to meet and exceed the original vehicle's safety, crash-worthiness, and consumer features.

This unique student team experience provides an excellent opportunity to examine an alternative, more hands-on and project-driven method of engineering education, called project-based learning. Project-based education, or experiential learning, has been discussed by researchers and implemented by institutions looking to either replace or augment the widely-used lecture-based format of post-secondary professional studies, particularly in professions like engineering and medicine [2]. In project-based learning, education centers around the execution of projects or case studies, and through their execution students learn the desired technical concepts [3]. This approach has been suggested as a way to expose students to important engineering concepts required by industry, such as teamwork, resource management, and dealing with uncertainty in real-world problems [4]; essentially the ability to apply classroom theory in the project setting where real-world engineering takes place.

1.1 Thesis Objective

This thesis describes the author's two years of project-based education working as a lead member of the University of Waterloo Alternative Fuels Team (UWAFT). Using case studies from the author's work on the UWAFT EcoCAR competition vehicle, this thesis documents the model-based design and construction of an advanced hydrogen powered plug-in hybrid vehicle (PHEV). This thesis also demonstrates how engineering principles are learned in a project-based environment and explores the benefits and drawbacks of this type of experiential education.

The University of Waterloo Alternative Fuels Team (UWAFT) is an example of a pure project-based learning environment. Many students typically participate in UWAFT through the ME481 senior design course, however there are several students who become involved from their first year and who participate without specific academic credit. At UWAFT projects are the focal point of each student's education. Individual students direct their own learning (with guidance from faculty and industry professionals) in order to achieve the goals of the project. This thesis will examine projects executed in that environment and how they offer teaching opportunities that complement the traditional lecture-based model.

This thesis is an examination of the inter-related design projects involved in building an electric powertrain fuel cell vehicle, with some of the author's select projects outlined in detail. The thesis highlights technical and non-technical problems that arise in a project-based learning environment, how they are solved, and how these experiences are valuable contributions to engineering education. References are made to the undergraduate Mechanical Engineering curriculum at the University of Waterloo (of which the author is a graduate) when discussing how project-based learning compares with and sometimes draws from the traditional lecture-based format. The aim is to identify elements of the curriculum that enable the team to meet the technical requirements of the project, and to highlight unique elements of project-based learning that are not found in the classroom. This thesis can be used by future engineering educators and student design teams to better prepare students for the technical and non-technical challenges faced by today's graduating engineers, and to

extract the most value from experiential learning opportunities in both student teams and lecture-based courses.

1.2 Thesis Outline

The thesis begins with an examination of the motivations driving research into alternative fuels, specifically the limited availability of oil resources [5], the environmental impacts of their use, and the role of the transportation sector [6] as contributor to urban air pollution. Next is a review of the literature on conventional and alternative automotive powertrains including hydrogen powered fuel cells. This is followed by a review of the current state of emissions and energy consumption evaluation procedures for conventional and emerging powertrains, including “dual-fuel” hybrid vehicles such as UWAF’s EcoCAR vehicle. This review provides a technical background for understanding the design work that is described in the body of the thesis. Finally, a review of recent project-based and experiential learning literature, with a focus on modern engineering education, sets the stage for the exploration of the project-based learning environment at UWAF.

With a technical background in electrified vehicle powertrains, the UWAF EcoCAR vehicle is introduced. The vehicle architecture is explained and the individual powertrain components are described. UWAF’s project-based learning environment is also introduced at this point, along with the technical design projects that are the basis for discussion of project-based learning. Specifically, these are the design of the thermal management systems, the powertrain control system, and the vehicle safety control system. Each of these projects highlight some of the key features of project-based technical learning, including the analysis of safety, assessment of risk, considerations for reliability and robustness, and dealing with uncertainty. These are all considerations that go well beyond the solutions in the back of a textbook and yet are ever-present in any real-world engineering design task that involves an open-ended problem. Non-technical elements, such as management of time, materials, cost, teams, and tasks are also discussed. While engineering graduates may not encounter all of these directly upon graduation, all will certainly encounter some of them throughout their careers.

These ‘outcomes’ of an educational program are consistent with the Canadian Engineering Accreditation Board (CEAB) ‘outcome-based’ educational criteria [7]. These criteria, to quote from the CEAB accreditation guidelines, are:

- *“A knowledge base for engineering: Demonstrated competence in university level mathematics, natural sciences, engineering fundamentals, and specialized engineering knowledge appropriate to the program.*
- *Problem analysis: An ability to use appropriate knowledge and skills to identify, formulate, analyze, and solve complex engineering problems in order to reach substantiated conclusions.*
- *Investigation: An ability to conduct investigations of complex problems by methods that include appropriate experiments, analysis and interpretation of data, and synthesis of information in order to reach valid conclusions.*
- *Design: An ability to design solutions for complex, open-ended engineering problems and to design systems, components or processes that meet specified needs with appropriate attention to health and safety risks, applicable standards, economic, environmental, cultural and societal considerations.*
- *Use of engineering tools: An ability to create, select, apply, adapt, and extend appropriate techniques, resources, and modern engineering tools to a range of engineering activities, from simple to complex, with an understanding of the associated limitations.*
- *Individual and team work: An ability to work effectively as a member and leader in teams, preferably in a multi-disciplinary setting.*
- *Communication skills: An ability to communicate complex engineering concepts within the profession and with society at large. Such abilities include reading, writing, speaking and listening, and the ability to comprehend and write effective reports and design documentation, and to give and effectively respond to clear instructions.*

- *Professionalism: An understanding of the roles and responsibilities of the professional engineer in society, especially the primary role of protection of the public and the public interest.*
- *Impact of engineering on society and the environment: An ability to analyse social and environmental aspects of engineering activities. Such abilities include an understanding of the interactions that engineering has with the economic, social, health, safety, legal, and cultural aspects of society; the uncertainties in the prediction of such interactions; and the concepts of sustainable design and development and environmental stewardship.*
- *Ethics and equity: An ability to apply professional ethics, accountability, and equity.*
- *Economics and project management: An ability to appropriately incorporate economics and business practices including project, risk and change management into the practice of engineering, and to understand their limitations.*
- *Life-long learning: An ability to identify and to address their own educational needs in a changing world, sufficiently to maintain their competence and contribute to the advancement of knowledge.”*

This thesis ends with a review of the on-road evaluation of the UWAFTEcoCAR from the final year of competition. The vehicle’s performance in each category is discussed with reference to the design tasks, demonstrating how project-based learning culminated in the successful construction of a working vehicle.

Chapter 2

Literature Review and Background

2.1 Motivation for Alternative Fuels

There are many reasons motivating the current research into alternative fuels for automotive powertrains. Oil, obtained using various methods ranging from drilling to the oil sands, is becoming increasingly expensive as resources are depleted. As of 2011 the world's proven (likely recoverable) oil reserves have been estimated by various sources as 1,474 billion barrels [5], 1,523 billion barrels [8], and 1,654 billion barrels [9]. Oil consumption in that year was between 29 billion barrels [9] and 32 billion barrels [5]. That is, as of 2011 there was enough recoverable oil known to exist to enable 50 years of global oil consumption at 2011 rates. The question of how long we can really continue to rely on oil for energy is, of course, much more complex. Economic growth, especially in developing countries like India and China, will drive increasing oil consumption rates, and new oil fields and new oil recovery technologies will almost certainly be discovered. More importantly, oil continues to be more difficult and more costly to exploit as easier reserves are depleted and we turn to "unconventional oil" like the Alberta oil sands or disruptive technologies such as 'fracking'. Finally, the nature of estimating oil reserves leaves a great deal of uncertainty that makes prediction of future economic feasibility extremely difficult. The question of when oil-based energy will become economically infeasible is far beyond the scope of this thesis. While no one can be certain, the general conclusion appears to be that we must replace oil with other energy sources over the next century if we are to meet our future energy needs [10].

With this goal in mind we must examine where energy is consumed in our economy so that we can identify the most effective strategies to reduce our petroleum dependence. In Canada, the transportation sector accounted for just under 31% of the country's total energy consumption in 2008 [6]. Further analysis of the data, shown in Figure 1, shows that road transportation alone accounted for 23.4% of the country's total energy consumption, with the overwhelming majority of that energy being derived from gasoline and diesel fuel. In fact,

petroleum products account for 32% of Canada’s total energy consumption as of 2010, followed closely by natural gas at 22% [11]. Globally, it has been estimated that 14% of all energy consumption is attributable to road transportation [12].

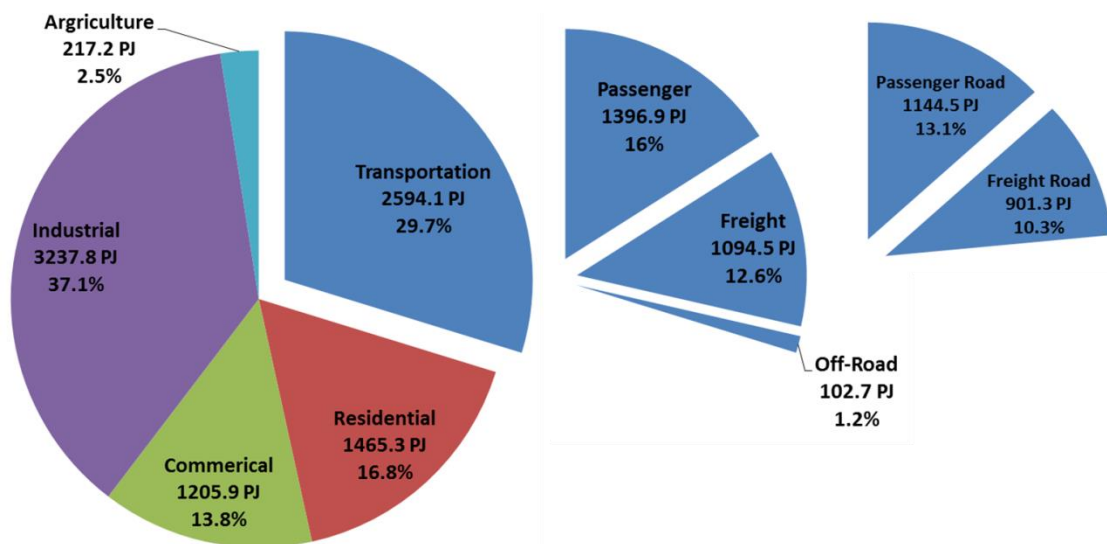


Figure 1: Energy use in Canada in petajoules, 2008, breakdown by sector, generated using data from Natural Resources Canada’s Energy Use Handbook [6]

Curbing the emission of the products of combustion is another motivator in the search for alternative energy sources. Greenhouse gases such as CO₂ have been targeted as contributors to climate change, and tailpipe emissions resulting from internal combustion engines additionally contribute to smog and declining urban air quality. Additionally, human populations (especially in developing nations) are steadily concentrating in higher density urban environments and a higher percentage of the population owns vehicles [13] [14]. Urban environments have been shown to possess lower emissions per capita, due in part to the higher density of homes and businesses and the mass availability of public transit, both encouraging a reduction in automobile use [15]. However, while lower emissions per capita (and hence lower emissions overall) makes a positive argument for urbanization, the airborne concentration of pollutants and hence per capita exposure is higher in cities, leading to (sometimes severe) air quality concerns.

In Canada, the road transportation sector accounts for a higher portion of nationwide greenhouse gas emissions than it does nationwide energy consumption. Figure 2 shows that road transportation is responsible for 28.7% of all greenhouse gas emissions in Canada as of 2008 [6]. In the United States the figures are similar with 32.6% of all GHG emissions attributed to the transportation sector as a whole [16].

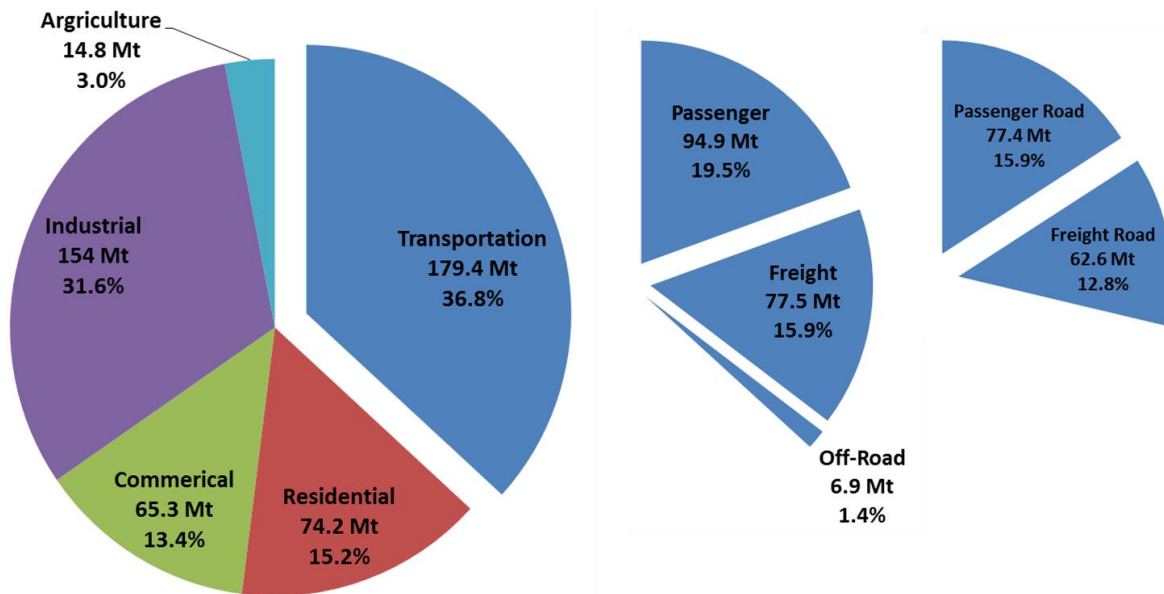


Figure 2: GHG emissions in Canada in megatonnes of CO₂ equivalent, 2008, breakdown by sector, generated using data from Natural Resources Canada’s Energy Use Handbook [6]

Additional motivation comes in the form of energy security. For many countries the lack of domestic oil reserves and the consequent need to import this vital energy source from potentially volatile regions raises the issue of having an independent, secure energy supply. These factors are driving significant investments in research for advanced automotive powertrains that use alternative fuels.

Compared to the industrial sector (the largest contributor to Canada’s energy requirements), transportation involves a relatively small set of technologies that evolve at a relatively rapid rate. For these reasons the transportation sector is an excellent target for advanced technologies that will provide alternatives to petroleum-based energy.

2.2 Electrification of the Automotive Powertrain

The efficiency of the internal combustion engine (ICE) is limited by the principle of thermodynamic efficiency and cannot be significantly improved, though manufacturers are still finding incremental optimizations in engine and powertrain design. The internal combustion engine running on gasoline or diesel fuels cannot be sustained into the future as petroleum supplies become scarcer and/or more expensive. The greenhouse gases emitted by these engines are also harmful. Using electricity for automotive propulsion is being explored with increasing interest as a solution to these problems.

Powertrains using some form of electricity for propulsion are rapidly becoming available to consumers at dealerships around the world. There are a wide variety of electrified powertrain architectures, each with specific advantages and disadvantages that make them best-suited for specific situations. With such variety of fundamental powertrains now available, the need to understand the specific application for a vehicle has become critical to maximizing the benefits offered by these new technologies. This section will examine the various electrified architectures that are currently employed in automotive powertrains, providing an understanding of their operation, their benefits, their drawbacks, and how they compare.

2.2.1 Gasoline-Electric Hybrids

Gasoline-electric hybrids combine a traditional gasoline engine with the ability to generate, store, and utilize electrical energy for propulsion. Batteries are often used to store electricity (though an ultra-capacitor can sometime be used as well), and one or more electric motors and generators are used to charge the batteries and drive the wheels. Many arrangements between these components and the ICE exist, and various “degrees of hybridization” can be achieved by varying the degree of electrification or the amount that the electrical part of the powertrain contributes to the overall propulsion of the vehicle.

Common arrangements of a hybrid powertrain have typically been classified as series, parallel, and series-parallel [17]. A series hybrid, as its name implies, has only a single pathway for propulsion energy. The usual combustion engine is solely connected to a generator, which creates electricity. That electricity flows into an electric motor for

propulsion, and/or into a battery/capacitor for storage. This appears counter-intuitive, given that this arrangement adds two energy conversion steps to the conventional ICE-only powertrain (mechanical-to-electrical at the generator, and electrical-to-mechanical at the motor). The value in this arrangement is realized by the de-coupling of the vehicle load and engine load. The combustion engine is relatively inefficient when operating in the low power region. During low speed city driving steady-state driving, deceleration, and idling, the vehicle demand is not high enough for the engine to operate at peak efficiency, and fuel is consequently wasted. With the addition of an energy storage buffer such as a battery, the engine can operate at a higher power output than would otherwise be permitted by the instantaneous vehicle power requirement, thus using fuel more efficiently. The additional energy conversion steps at the electric generator and motor are more than justified by the increase in the engine efficiency. The electric motor can also act as a generator during deceleration, recovering some of the vehicle's kinetic energy that is normally lost to friction braking. This is known as regenerative braking. When the vehicle is idle the engine can be shut off altogether without concern for the time required to re-start it, since the electric motor can run off battery energy for a short time. This saves significant fuel that is otherwise completely wasted. This is commonly referred to as 'start-stop' technology. An additional benefit of the series arrangement is that the electric motor is the sole source of traction power. Since electric motors can produce maximum torque at zero speed, the series hybrid provides superior acceleration performance compared to the combustion engine, which must spool up to several thousand RPM before maximum torque is produced. Finally, the size of the combustion engine can be reduced, as it no longer needs to supply all of the power required for acceleration. The smaller engine will be more efficient at the lower power levels of a load-leveling energy source, and it will be easier to integrate into a vehicle that is now sharing physical space with several new components.




Figure 3: Series hybrid powertrain

There are some drawbacks to the series hybrid arrangement. The series hybrid gains efficiency over its non-hybrid counterpart when the vehicle operates outside the engine's optimal operating region, but the extra energy conversion steps that enable this gain, which are unavoidable given the single energy pathway, yield a net loss when the vehicle load falls inside the engine's efficient operating region, where a conventional vehicle would thus operate more efficiently without the extra energy conversion steps (not to mention without the additional weight of the motor, generator, and batteries). Also, while the engine size can be reduced / optimized for the lower engine power demands of a series hybrid, the generator and motor must be large enough to support all of the vehicle's power requirements, resulting in significant cost.

A parallel hybrid uses a motor and batteries to augment the combustion engine, but the arrangement of these components is drastically different than in the series version. In a parallel hybrid there are multiple energy pathways leading from the power plants to the wheels. The engine and electric motor are connected in parallel; this connection can happen at multiple points along the powertrain, for example at the input or output of the transmission, or at the engine's crankshaft as a belt-driven accessory. The battery is connected to the electric motor, in the same manner that the fuel tank is connected to the engine. The two power plants, connected in parallel, can interact in several ways. The engine and the electric motor can add their output torques to drive the wheels together. Either the engine or the motor can drive the wheels on their own, with the other either disconnected via

a clutch or spinning freely at the same speed. The engine can also drive the motor in this arrangement, which will then act as a generator and charge the battery.

The parallel hybrid arrangement solves some of the disadvantages of the series hybrid. The electrical components do not solely fulfill the traction demands of the vehicle, and so they can be reduced in size, weight, and cost. The arrangement also permits the engine to drive the wheels, to charge the battery, or both. Unlike the series hybrid, if the vehicle load falls within the optimal efficiency region of the engine, the engine may simply drive the wheels directly, skipping the additional energy conversion steps that are required for a series hybrid.

Regenerative braking and engine stop-start are also possible. Due to the mechanical connection between the electric motor, the engine, and the wheels, however, the parallel hybrid cannot operate the engine speed independently of the motor and vehicle speed. Also, the parallel hybrid typically does not have the instant-torque acceleration of the series arrangement, due to the smaller electric motor and the reliance on the engine for some traction power. In some vehicles, the “parallel” connection between the engine and the electric motor is achieved by placing one on the front axle and one on the rear. The road becomes a link between them, and the engine can charge the battery by putting torque on the front axle, which is consumed by the electric motor on the rear axle. This is known as “parallel through the road”.

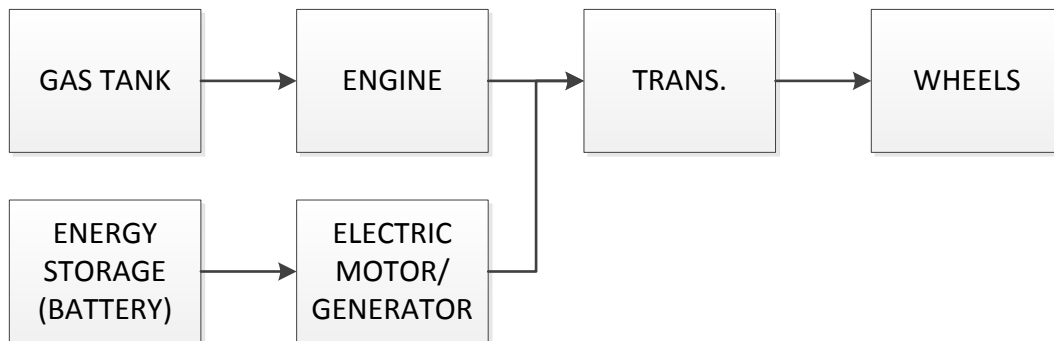


Figure 4: Parallel hybrid powertrain

Parallel hybrids are further divided into full, mild, strong, and various other descriptions of the “degree of hybridization”. The combustion and electrical power plants operate in parallel to split the load, and the degree of this split varies depending on the objectives of the design.

Mild hybrids have a smaller motor and battery, which means that their ability to both capture braking energy and drive without the use of engine power is limited. Strong or full hybrid use larger components, adding cost and weight but enabling greater energy recovery and more engine-off operation, both yielding greater system efficiencies.

It is clear that series and parallel hybrids each have distinct advantages and disadvantages. Some of these advantages (regenerative braking, stop-start) are common to both arrangements, as are some of the disadvantages (cost, complexity, weight). In other areas, such as efficiency during steady-state driving or efficiency during transient driving, these arrangements are polar opposites. The logical solution is to combine the benefits of both arrangements, which has been done in the form of a series-parallel hybrid.

The series-parallel hybrid connects the engine, two electric motors/generators, and the wheels together in a planetary gear set. Depending on where the power is split within the transmission, this type of hybrid can be classified as input-split, output-split, or compound-split. By controlling the speed of the electric motors the ratio of the traction power supplied by the internal combustion engine (ICE) and electric components can be varied infinitely. This is the parallel part of the architecture. The planetary gear set also allows the engine to operate independently of the wheels, giving the de-coupling capabilities of a series hybrid as well. Thus, the series-parallel is the most flexible and efficient hybrid arrangement, allowing for all the benefits of regenerative braking, optimized / downsized engine and motor/generator components, engine start-stop, and optimized energy pathway selection. This comes at the cost of increased complexity in the design of the series-parallel device and the additional electric motor that completes the triplicate of components controlling the planetary gear set. The series-parallel hybrid is the most common arrangement in commercially available vehicles, and the most popular of those is probably the Hybrid Synergy Drive in the Toyota Prius. General Motors also participated in the development and sale of the Two-Mode Hybrid.

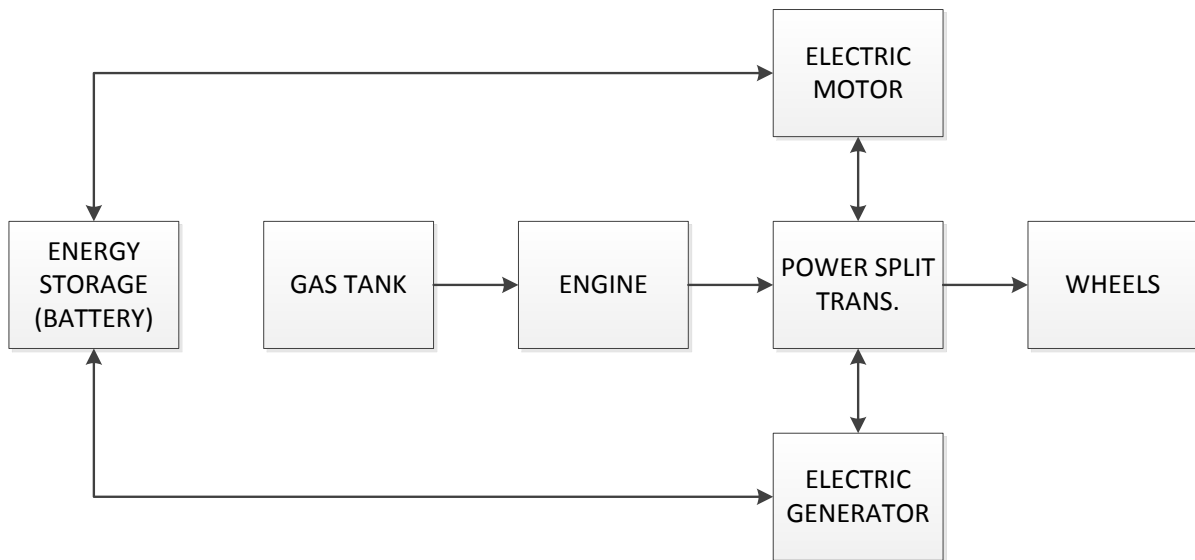


Figure 5: Series-parallel hybrid powertrain

It is important to note that all of the hybrids discussed thus far, despite using both gasoline and electricity, are actually single fuel vehicles. All on-board electricity, including the kinetic energy captured during regenerative braking, is originally generated by the combustion engine.

A final category of gasoline-electric hybrids are plug-in hybrid electric vehicles (PHEVs). The PHEV can theoretically be based on any of the aforementioned hybrid architectures. It adds a battery charger that can source electricity from the grid to charge the battery, resulting in a true dual-fuel vehicle. The PHEV retains the advantages and disadvantages of the hybrid architecture that it is based on, with the additional advantage that grid electricity can, in some circumstances, be generated more efficiently and with less environmental impact than electricity generated by a hybrid powertrain. To extract the maximum benefit from a PHEV it is typically based on a full hybrid that has the battery capacity to store enough grid-sourced energy to provide a meaningful vehicle range (typically at least 40 to 60 km), and that has a large enough electric motor to propel the vehicle without the aid of the engine.

2.2.2 Battery Electric Vehicles (BEVs)

If combustion engine vehicles lie at one extreme of the electrified powertrain spectrum (in that the electrification factor is zero), then battery electric vehicles lie on the opposite extreme, with the various types of hybrids discussed in section 2.2.1 lying in-between. Pure electric vehicles forego the combustion engine entirely. They rely solely on batteries charged from the grid for energy storage and on electric motor(s) for propulsion. In concept they are simple designs with few components. A few notable examples of pure electric vehicles are the Nissan Leaf, Mitsubishi i-MiEV, Ford Focus Electric, Chevrolet Spark Electric, and the Tesla Model S sedan.

Battery electric vehicles are touted to have the advantage of absolutely zero emissions. While this is true at the vehicle level (point of use), emissions are simply shifted upstream to the electrical generating stations. In regions where the grid is supplied with a large percentage of “clean” electricity (e.g. hydro, solar, wind, nuclear) this shift results in an overall reduction in greenhouse gas emissions. In other regions that rely on coal, the shift can actually have a net negative impact, as discussed in section 2.4.

Electric motors are a mature technology, and the recent interest in electrified vehicles has spawned a wide range of highly capable electric motors that are engineered for automotive applications and that offer performance comparable to a combustion engine. Battery technology, however, has not yet managed to meet the energy storage capabilities of fossil fuels. Lithium-ion (Li-ion) battery chemistries are the popular choice for automotive applications, due to their high energy and power density. Table 1 summarizes fuel energy density figures for hydrogen, gasoline, and one type of Li⁺ cell. While energy storage system packaging is not accounted for, Table 1 makes it clear that battery technology falls significantly short of the energy storage capability of fossil fuels. Thus, while electric vehicles have a significant advantage in energy conversion efficiency, they typically weigh more than conventional vehicles and have a significantly shorter range.

Table 1: Fuel energy density figures for EcoCAR competition energy storage [18], [19]

FUEL	DENSITY (kg·m⁻³)	MASS ENERGY DENSITY (kWh·kg⁻¹)	VOLUMETRIC ENERGY DENSITY (kWh·m⁻³)
Hydrogen (atm)	0.081	33.3	2.7
Hydrogen (70 MPa)	39.2	33.3	1305
Li+ (A123 LiFePO ₄)	1818	0.135	245
E10	753	11.6	8742

2.2.3 Fuel Cells

Fuel cell vehicles are a type electric vehicle that stores energy in on-board hydrogen tanks. Hydrogen is fed into the fuel cell anode and air is fed to the fuel cell cathode. The hydrogen fuel and oxygen in the air are converted within the fuel cell into electricity and water. There are no greenhouse gas emissions at the vehicle level. The electricity generated by the fuel cell is used to power electric motor(s) for propulsion. A battery is often used in a fuel cell powertrain to provide the same de-coupling benefits (reduced power plant size, operation in optimal efficiency region, regenerative braking) that are found in gasoline-electric hybrids. The battery is also needed to handle the rapid response times necessitated by the vehicle, as this is currently difficult for fuel cell technology [20]. Fuel cell powertrains are discussed in more depth in section 2.3.

2.3 Fuel Cells in Automotive Applications

2.3.1 Fuel Cell Technologies

Fuel cell vehicles are fully electric vehicles that use a hydrogen fuel cell as the main power plant. The development and commercialization of fuel cell powertrains is driven by several inherent benefits of the technology. Fuel cells, emitting only water vapour as a by-product of the energy conversion process, have zero “tailpipe” emissions. Having no emissions at the vehicle level is highly beneficial in urban, densely populated environments where air quality is a major concern. Fuel cells are also able to achieve much higher energy conversion efficiencies as compared to internal combustion engines. Today, large scale production of

hydrogen fuel is typically done by reforming natural gas. However, fuel cells have the potential to run on hydrogen generated from renewable energy sources. In this scenario hydrogen gas is generated by renewable energy sources such as solar, wind, or tidal power through electrolysis of water. Hydrogen gas is then used as an “energy vector” with a much higher mass energy density than would otherwise be achieved from the best battery technology, and is re-converted into electricity by an on-board fuel cell system.

While there are a number of fuel cell technologies in existence, automotive applications have typically focussed on the polymer electrolyte membrane (PEM) fuel cell. Table 14 in Appendix A summarizes key attributes, such as operating temperature, power density, and efficiency of the common fuel cell technologies; polymer electrolyte membrane (PEM), alkaline fuel cell (AFC), phosphoric acid fuel cell (PAFC), molten carbonate fuel cell (MCFC), and solid oxide fuel cell (SOFC). The tabulated format quickly illustrates the attractiveness of PEM technology in automotive applications. PEM fuel cells operate at relatively low temperatures, comparable to internal combustion engines, meaning that existing automotive thermal management systems can be adapted to them. By comparison, higher temperature fuel cells such as MCFC and SOFC would introduce significant challenges in materials selection and packaging of components near the stack, and are better suited to stationary applications where their higher temperatures (thus higher quality thermal waste energy) can be used as an advantage in combined heat and power cycles (CHP). The PEM cathode reaction is tolerant of non-oxygen gases in the atmosphere on the cathode side, allowing a PEM fuel cell to breathe the surrounding air instead of carrying compressed and purified oxygen. AFCs, on the other hand, are sensitive to CO₂ in our atmosphere and are thus more suitable for space applications (where they have been common for decades). PEM fuel cells also have quick start-up times, good transient response, and relatively high efficiency and power density, all critical in automotive powertrains where weight must be minimized and operating conditions change constantly.

PEM fuel cells convert hydrogen and oxygen into electricity via complementary oxidation and reduction reactions. The anode and cathode are physically separated by a membrane. Hydrogen fuel is fed to the anode of the cell, where a catalyst (usually a precious metal, such

as platinum or platinum alloy) initiates the separation of H₂ molecules into two H⁺ ions and two electrons. Oxygen is made available at the cathode, usually via a compressor which simply directs ambient (filtered) air into the stack, much like a turbocharger in a combustion engine. The electrons generated at the anode create the electrical current used to power the load (e.g. electric motor). The H⁺ ions (protons) travel through the membrane to the cathode (hence the same proton exchange membrane), where they combine with available oxygen and the electrons flowing through the circuit to produce H₂O. Due to mass transfer effects, the hydrogen concentration at the anode must typically exceed the requirements of the chemical reaction. This excess hydrogen can be recycled by the fuel supply system and re-supplied to the anode. Air and water (usually a mix of liquid water and water vapour) exit the cathode and are simply exhausted into the atmosphere.

The efficiency and cell potential of a hydrogen fuel cell is derived from an application of the first and second laws of thermodynamics to the fuel cell. The maximum thermodynamic or “reversible” efficiency and the corresponding maximum “reversible” cell voltage are described as:

$$\eta = \frac{\Delta g(T, P)}{\Delta h(T, P)} \quad \text{Equation 1}$$

$$E_r = \frac{-\Delta g}{nF} \quad \text{Equation 2}$$

For the hydrogen fuel cell reaction with a liquid water product, given Enthalpy of Formation H = -285.8 kJ·mol⁻¹ and Gibbs Free Energy G = -237.2 kJ·mol⁻¹, the ideal efficiency for the 2H₂ + O₂ = 2H₂O_(l) reaction is calculated below.

$$\eta = \frac{-237.22 \text{ kJ/mol}}{-285.8 \text{ kJ/mol}} = 83\% \quad \text{Equation 3}$$

$$E_r = \frac{-(-237.22)}{(2)(96487)} = 1.23 \text{ V} \quad \text{Equation 4}$$

This maximum theoretical cell voltage is not achievable in practice, however. Operation of the fuel cell leads to various losses that reduce the achievable cell voltage, and hence the achievable output power and overall efficiency. The main sources of loss in an operating fuel cell are the activation, ohmic, and concentration losses. Respectively, these losses dominate at low, medium, and high currents. A plot of voltage vs. current for the EcoCAR vehicle's fuel cell can be found in Figure 28 of section 5.1.2. This plot is generally referred to as a "polarization curve", and the behaviour it demonstrates has a significant impact on the design and operation of a fuel cell vehicle.

The activation polarization refers to the losses associated with the charge transfer reaction that results in moving electrons from the electrode. Its effect is most pronounced at low currents, where it is the dominant loss type. Ohmic polarization results from the resistance to charge transfer within the fuel cell. The voltage loss due to ohmic polarization generally increases linearly with current, throughout the operating curve, according to Ohms law ($V = IR$), where the resistance to charge transfer (both negative electron charge transfer through current collection plates and positive H^+ ion transfer through the membrane) all contribute to voltage loss in the cell. During fuel cell operation it is ohmic polarization that causes the "voltage sag" of the fuel cell system as current is increased. Concentration polarization becomes dominant at very high currents, as a result of the limitations of mass transfer within the fuel cell. It is the concentration polarization that quickly limits the overall reaction rate and thus governs the maximum current that the fuel cell can produce. The limited ability to move reactants within the fuel cell ultimately limits the reaction rate and hence the achievable output current.

The polarization curve of a fuel cell has a significant influence on the electrical design and control of a fuel cell powertrain. It has already been established that fuel cells are typically paired with an electrical energy storage system (battery) in automotive powertrains. These batteries operate on similar electrochemical energy conversion principles, and hence demonstrate polarization effects similar to a fuel cell. During vehicle operation the fuel cell and battery power requirements will change significantly and rapidly, leading to large swings in voltage for the two energy sources. If these two components were simply connected

together the energy transfer between them would be uncontrolled and potentially damaging, with current flowing from the higher potential to the lower through very little resistance. This leads to the requirement of a DC-DC power converter, placed between the two sources, to bring control to this arrangement. The role of the DC-DC power converter and how it is involved in the control of a fuel cell powertrain is discussed extensively in section Chapter 5.

2.3.2 Current State of the Automotive Fuel Cell Powertrain

A typical fuel cell vehicle powertrain is shown in Figure 6. The fuel cell receives hydrogen from on-board storage tanks, and feeds it into the anode. Air from the vehicle's surroundings is forced into the cathode by a compressor. The resulting electricity is fed through a power conditioning device, a DC/DC converter, which permits voltage matching with the hybrid battery. The DC-DC converter can also be placed at the battery instead of the fuel cell to provide voltage matching. The fuel cell supplies current to this main bus to drive the electric traction motor(s) and vehicle auxiliaries (e.g. air conditioning), together referred to as the vehicle load. The fuel cell can also be commanded to supply additional current, above the vehicle load, to charge the battery. This is desirable if the extra current, with consideration for the efficiency impact of charging and discharging the battery, causes the fuel cell system to operate in a more efficient region. This is identical to the concept of load de-coupling that is employed in ICE-based hybrid vehicles, as discussed in section 2.2.1.

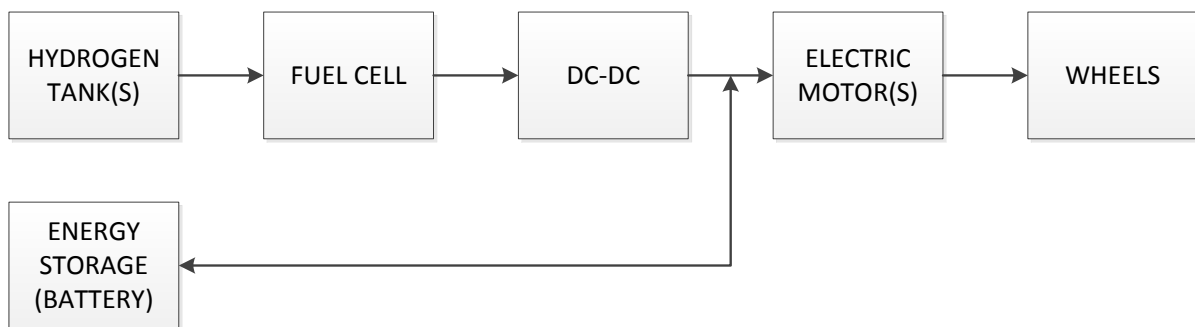


Figure 6: A fuel cell powertrain with DC-DC converter on the fuel cell side of the bus

Being radically different than current vehicle architectures, fuel cell vehicles face significant barriers to commercialization. Some of the most significant barriers are the in-vehicle powertrain technologies. The materials required to construct a fuel cell are relatively

expensive. The catalyst at the anode is traditionally comprised of platinum alloy; reducing the required mass of this precious metal or substituting it altogether is the subject of much research [21] [22]. Ford has estimated that 40% of the cost of a fuel cell system comes from the platinum used in catalyst layers [23]. The PEM that enables cell operation is also actively researched to mitigate side reactions and humidification requirements [23]. Power electronics are needed throughout the powertrain to control the traction motor(s), manage voltage levels, and charge the high voltage battery; all of these components are new to the automotive environment and are undergoing significant research and development (R&D) to reduce size, weight, and cost.

Another major obstacle to fuel cell vehicle commercialization is the lack of fuelling infrastructure. Electrified plug-in powertrains have been able to leverage the ubiquity of household electricity, workplaces, and public spaces requiring only the installation of a charging station to connect to the grid. Conversely there are practically no fuelling options for hydrogen vehicles today (except in limited markets such as California and the greater Vancouver region), aside from a very small number of stations maintained mostly by operators of small hydrogen vehicle fleets. Without hydrogen fuelling stations it will be difficult to market hydrogen vehicles, and without hydrogen vehicles there is little incentive to deploy a network of hydrogen fuelling stations. Close collaboration between vehicle manufacturers and potential fuelling companies is needed to resolve this issue, or vehicle manufacturers must invest in fuelling infrastructure.

A final key barrier to fuel cell vehicle commercialization is consumer acceptance. Currently interior space must be sacrificed, even in hybrid and plug-in hybrid vehicles, in order to accommodate the additional powertrain components. Even with the huge storage pressures of the latest hydrogen fuel tanks (up to 10,000 psi or 70 MPa), a significant volume is required to integrate them with the vehicle. This also directly impacts range. While hydrogen vehicles do not pay the same weight penalty for energy storage as do electric vehicles, the volume penalty is significant even at high storage pressures (as shown in Table 1 of Chapter 3). Current generation hydrogen vehicles can achieve greater range than pure electrics, for example the Honda FCX Clarity can run up to 380 km [24], compared with 120 km for the

similarly classed Nissan Leaf [25]. While hydrogen powertrains are also more efficient, Ford has estimated that a 70 MPa hydrogen storage tank requires 5 times the volume of a gasoline tank to achieve the same vehicle range [23]. However this is still below the range of most passenger cars and fueling infrastructure remains non-existent.

Consumers must also trust the safety of hydrogen vehicles. While compressed hydrogen gas can be dangerous, it is often forgotten that gasoline is also dangerous, and years of familiarity and design have mitigated the risk. Hydrogen is significantly easier to ignite, having a minimum ignition energy (MIE) of about 0.02 mJ, as compared to approximately 0.1 mJ for typical hydrocarbons [26]. Hydrogen is flammable in air at concentrations ranging from 4% to 74% again compared with 1.4% to 7.6% for gasoline [27]. Unlike gasoline, hydrogen leaks disperse quickly due to its buoyancy in air, reducing the likelihood of the gas pooling to a concentration that can be ignited. Consumer education on the characteristics of hydrogen gas and how they compare to gasoline, education on the safety systems designed into fuel cell vehicles, and proven safety records stemming from on-road trials and government testing specific to hydrogen fuels can alleviate consumer safety concerns. One study in Germany has shown that the general public actually has relatively little concern for the safety of hydrogen. The key requirement of acceptance for those consumers is that the hydrogen be sustainably sourced using environmentally-friendly methods [28].

2.4 Evaluating the Emissions and Energy Consumption of Electrified Powertrains

It was previously stated that the key drivers of alternative fuels and advanced powertrains are the reduction of atmospheric emissions, increased energy efficiency, and a shift to renewable sources of energy, all contributing to the overall goal of a sustainable transportation infrastructure. How we evaluate these key metrics is then of great concern to both consumers buying these new vehicles, governments and regulatory bodies that set the evaluation methods and standards for such metrics (and sometimes subsidize vehicles who meet certain targets, or penalize others that miss them), and industry organizations developing these new technologies. The objective of this section is to review the current state of fuel economy and

emissions testing (with a focus on North America), to identify the strengths and weaknesses of current approaches, and how approaches are changing (and still require change) to respond to advanced powertrains driven by alternative fuels. It is notable that fuel economy and emissions were once simply correlated, due to the use of a common fuel source (gasoline) with common production (liquid oil wells) and consumption (internal combustion engine) methods. Today we have electrics, fuel cells, natural gas, and other automotive fuels that complicate this correlation.

2.4.1 Fuel Economy Evaluation for Conventional Powertrains

Fuel economy testing for conventional vehicles is complex, owing to the great number of variables that can affect the fuel consumption of a particular vehicle. The “drive cycle”, or vehicle speed as a function of time, is one of the key factors affecting fuel economy. It is well known that high acceleration rates, long idling times, and varying speed (e.g. stop-and-go) all contribute to poor fuel economy. A vehicle speed trace captures these factors, and standardizing such “drive cycles” across all vehicle fuel economy testing ensures consistent conditions. Cycles are driven in a controlled environment, usually inside a laboratory on a dynamometer that has been programmed to simulate the loads that the vehicle would experience on the road (air resistance and rolling resistance, which vary as a function of vehicle speed). The methodology, then, takes great care to ensure that the testing conditions across all vehicles are consistent and will reflect the fuel economy of the vehicle under the same conditions in the real world.

One of the concerns with standardized fuel economy testing is that the “standard” cycles are not representative of the typical driving style. Thus, while they may be able to simulate the fuel consumption of the vehicle, they are not good at simulating the way people drive their vehicles. Standard cycles such as the Urban Dynamometer Drive Schedule (UDDS) (shown in Figure 7), the HWFET (Highway Fuel Economy Driving Schedule), and the US06 Driving Schedule, are examples of standard drive cycles used by the EPA to test vehicles [29]. These cycles are commonly criticized for being unrepresentative of the way that consumers drive their vehicles, having lower than average speeds and accelerations. Standard drive cycles

based purely on a vehicle speed trace also do not account for the effect of auxiliary loads (e.g. HVAC system) and the effect of “cold-start”, where the engine is less efficient if it is warming from ambient temperature. In northern climates, requiring a longer warm-up period before the engine reaches peak efficiency has a significant impact on fuel economy [30].

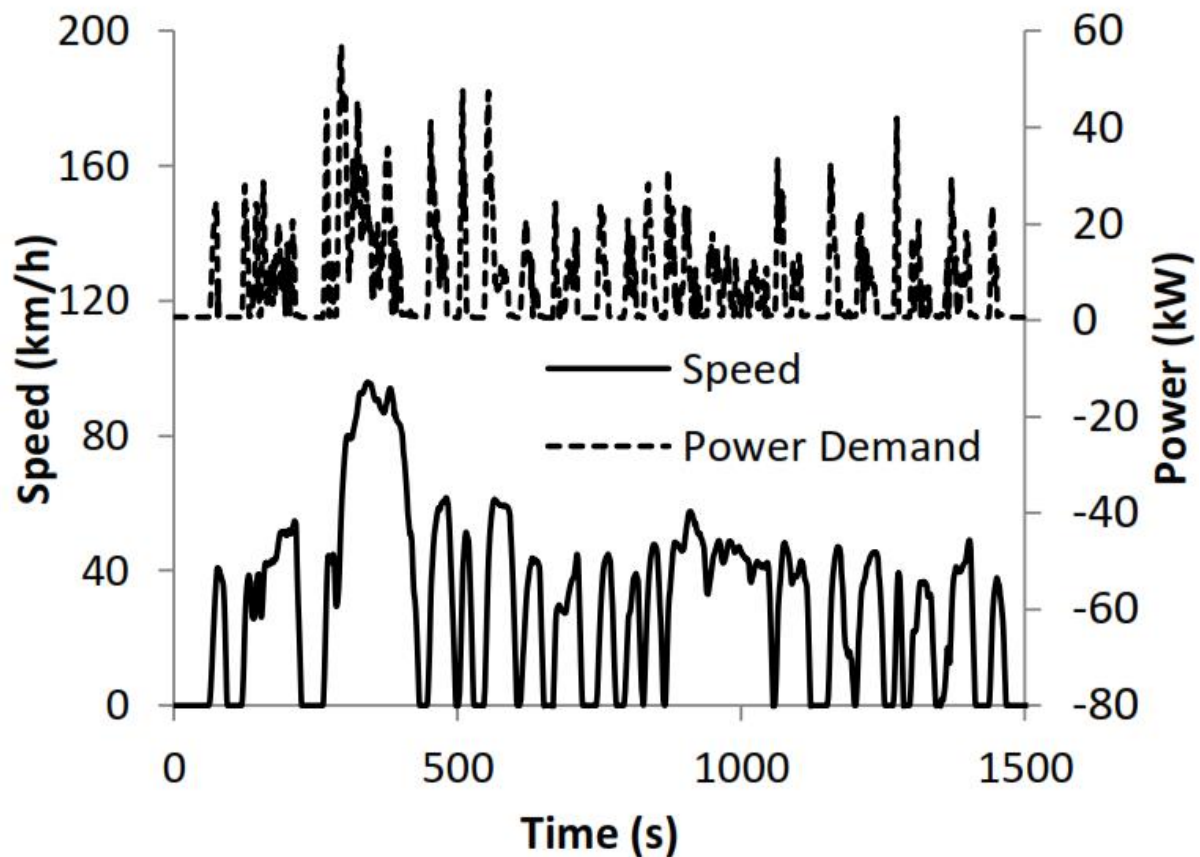


Figure 7: UDDS drive schedule driven by UWAFET EcoCAR at EPA National Vehicle & Fuel Emissions Laboratory

To address these concerns, starting with the 2008 model year the EPA instituted a new 5-cycle test procedure. The traditional city cycle is represented by the Federal Test Procedure (FTP), which runs the UDDS cycle from a cold start, then runs the first 505 seconds immediately afterwards (a hot start) to capture the effects of both cold and hot starting. Highway driving is still represented by the HWFET cycle. Attempting to account for the relatively higher aggressiveness of the typical driver, the US06 schedule has been added to

the test procedure, which features higher average speeds and accelerations. The Speed Correction Driving Schedule (SC03) is also used to account the impact of air conditioning use with hot (35°C) ambient temperatures. Finally, the effects of cold temperatures are evaluated by running the FTP schedule at -7 °C [31].

The change to the more realistic 5-cycle testing predictably lowered the fuel economy ratings of tested vehicles, bringing them closer to real world numbers. Even with the shift to a more representative test procedure, the EPA is careful to note that “your mileage may vary” due to the effect of acceleration, idling, speed, ambient temperature, trip duration, terrain (e.g. grade), cargo, and auxiliary loads [32]. In the end fuel economy is highly personal, and the EPA standardized testing can at best represent the average that the population as a whole can expect to get from each vehicle.

2.4.2 The Impact of Electrified Powertrains

While alternative fuels like natural gas and propane have been in use in the transportation sector for decades, their market penetration has been minimal and typically limited to commercial use. Hybrid-electric technology is the first great shift in powertrain design since the combustion engine, and has caused a re-think of conventional methods for evaluating fuel economy and emissions.

For the purposes of fuel consumption measurement, plain hybrids without “plug-in” technology are well-served by the existing evaluation methods, at least as well as conventional vehicles. While electric energy is consumed by the powertrain, it is generated solely by the gasoline engine on-board, whether directly via a generator or in-directly through the recovery of kinetic energy in regenerative braking. As a result, only a single fuel source powers the vehicle and existing measurement methods capture this fuel use.

An interesting consequence of hybrid technology, however, is the inversion of conventional wisdom on the relative efficiency of city and highway driving. As discussed in section 2.2.1, hybrids gain efficiency in transient driving scenarios by using engine load decoupling and regenerative braking to both optimize energy conversion efficiency and re-capture otherwise lost kinetic energy. This is well-represented by “stop-and-go” city driving cycles. On the

highway, however, a hybrid powertrain operates in a relatively steady state, thus engine load decoupling and regenerative braking offer less benefit, and in some cases could reduce efficiency due to the weight of the additional components relative to the non-hybrid version. As a result, fuel economy ratings for hybrid vehicles sometimes indicate that the vehicle is equally or more efficient in city driving than in highway, a reversal of the trend for conventional powertrains. For example, the 2013 Toyota Camry (base 2.5 L) is rated for 25 MPG city / 35 MPG highway, while its hybrid counterpart is rated for 43 MPG city / 39 MPG highway, according to published EPA fuel economy numbers [33]. While still more fuel efficient than the conventional powertrain, the gap closes significantly in highway driving, reducing the hybrid benefit and extending the payback period for its incremental cost.

More recently, plug-in hybrid electric vehicles (PHEVs) have become available, resulting in electric “dual-fuel” vehicles that are powered both by gasoline converted through a combustion engine and electricity sourced from the grid. Plug-in hybrids, essentially hybrids with an on-board battery charger, offer the capability to use externally-sourced electricity (e.g. from the grid), with the potential benefit of displacing petroleum energy products (depending on the source of the electricity, of course). The addition of a second energy source created a few problems for the existing fuel economy test methods, however.

First, a simple matter of inconsistent units existed, where electricity use is conventionally measured in kilowatt-hours and gasoline is usually measured by volume as litres or gallons. Consumers are used to thinking about their fuel consumption volumetrically, and so “equivalent litres” of electricity has become a common conversion.

Second, the instantaneous efficiency and emissions of a plug-in powertrain depend on the ratio of the externally-sourced electricity and gasoline being used at that instant. In addition, there is only a relatively small amount of externally-sourced electricity stored on-board due to weight and volume restrictions, resulting in a limited operating range where this energy can be used. This leads to two distinct operating modes for a plug-in vehicle, and significant consequences for fuel economy testing.

2.4.3 Plug-In Powertrains, the Utility Factor, and New Testing Standards

The two distinct operating modes of a PHEV are charge-depleting (CD) and charge-sustaining (CS). In charge-depleting operation the externally-sourced electricity is drawn from the battery, and is either used alone (EV-only charge-depleting operation) or in combination with energy from the engine (blended charge-depleting operation). While short-term operating conditions may see the battery SOC maintained or even increase (e.g. regenerative braking event), the long-term trend characteristic of charge-depleting operation is a steady consumption of the externally-sourced energy.

Eventually (and in a relatively short period of time) the externally-sourced energy is nearly depleted. The PHEV then switches to charge-sustaining operation, where the battery SOC is maintained about a certain set-point. In this mode of operation the PHEV reverts to the functionality of an HEV, where the engine provides the sole source of energy and the electrification benefits are limited to engine load decoupling and regenerative braking. Again, short-term operating conditions will cause the SOC to fluctuate, but the long-term trend characteristic of charge-sustaining operation is an SOC maintained about a set-point, using (on average) zero externally-sourced electricity. This can be seen in a plot of data collected from the UWAFI vehicle during the EcoCAR competition, in Figure 53 of section 7.3.

With a focus on “fuel” economy, the hybrid electric vehicle fuel economy test standard, SAE J1711, was updated in 2010 to formally address the challenge of measuring and reporting (two very different concepts in a dual-fuel vehicle, as will be shown) the fuel economy of PHEVs [31]. While J1711 is a long and complex document, its core feature pertaining to fuel economy measurement is the simple recognition that the new charge-depleting operation mode exhibited by PHEVs has a unique fuel economy and a limited range. The 5-cycle tests remain the core driving schedules for fuel economy testing, but instead of a single test procedure (comprised of one run through each of the five cycles) there are two; the Charge-Sustaining Test (CST) and the Full-Charge Test (FCT).

The CST and FCT are generally differentiated by the preparation of the vehicle. HEVs exhibit only charge-sustaining operation and therefore complete only the CST, which is comprised of the standard five driving schedules performed with a vehicle that has been carefully prepared according to the standard in an effort to ensure consistency across all vehicles tested. PHEVs have the additional requirement to complete the FCT. The FCT specifies the additional preparation of ensuring that the rechargeable battery has a full charge before each driving schedule is run. The energy required to perform this charge must be measured accurately, as the efficiency of the vehicle's battery charger is included in the overall electrical efficiency of the powertrain; that is to say that electrical energy usage of a PHEV is measured at the input to the vehicle, not at the battery. For each driving schedule, the vehicle drives the cycle in a repeating loop until the charge depleting operation mode is complete and at least one full cycle of charge-sustaining operation is observed. This enables the FCT to capture both the fuel economy of the charge-depleting mode, and the duration that charge-depleting mode can be expected to last under each of the five schedules. With these measurements, PHEV fuel economy can be reported.

In J1711 fuel economy reporting is not as straightforward as simply providing the fuel economy measurement. In the case of PHEVs, there are two fuel economy ratings, the charge-depleting and the charge-sustaining. Giving consumers both numbers could be confusing, so a way to combine them into a single, familiar number was required. For this the standard introduces the concept of the utility factor (UF).

The UF is a weighting factor that can be used to combine the CD and CS fuel economy and emissions values into a single value, as described in Appendix A of J1711. While a UF can be derived to serve several purposes, the basic purpose is to have a number that represents the fraction of total driving distance that can be done in CD operation. For a given CD range, the UF curve shows the percentage of total vehicle distance driven in the sample data that would be accomplished in CD mode. A sample UF curve generated from the EcoCAR formula given in [18] is shown in Figure 8 to illustrate the point. If a vehicle has a CD range of 50 kilometers, according to Figure 8 over 50% of the sample population's vehicle kilometers would be driven in CD mode. Thus, it could be said that the average fuel economy

of the population driving that vehicle would be 25% of the CD measurements and 75% of the CS measurements. That is exactly what J1711 prescribes; the measurements from each of the 5 driving schedules, from the CST and FCT tests, are combined using the UF, where the measurements from CD operation are weighted by UF, and measurements from CS operation are weighted by $(1 - UF)$. This method enables a single “utility-factor weighted fuel economy” to be calculated for a PHEV.



Figure 8: Utility factor curve generated using EcoCAR competition formula given in [18]

While this single weighted value appears simple (a desirable attribute for communicating to consumers), significant assumptions are in-built that would not be immediately obvious to the uninitiated. The statistical data for the UF is collected from the National Highway Transportation Survey (NHTS), and is used in SAE J2841 to calculate UFs for weighting CD and CS tests. It has already been established that, even for conventional vehicles, fuel economy is highly dependent on an individual’s driving style, and fuel economy testing can only capture an average over a range of styles / driving schedules. Hybrid technology introduced new variations to the effect of driving style on fuel economy, via regenerative braking and idle prevention (engine stop-start). PHEVs now add a new source of variability; the distances travelled by an individual between charging opportunities. For a vehicle with a CD of 10 miles, a driver with a daily commute to work of 8 miles each way, with charging

available in-between, would have an individual UF of 100% (all miles are travelled in CD mode). Another driver with the same vehicle may drive 15 miles each way, with no charging in-between, yielding a UF of only 33%. These drivers could exhibit the exact same driving style, even drive the same schedule, and receive very different fuel economies. The exact same process can be derived to average the AC electrical energy consumption over the vehicle miles of a population, yielding very different values depending on the intervals between charging. This is the pitfall of UF-weighted fuel economy. It uses the statistics of the population to provide a single average value for the fuel and electrical energy consumption of a PHEV deployed in that population, leading to a large “your mileage may vary” disclaimer for the individual. When only the fuel economy rating is considered (without the AC electric energy consumption), the fuel economy becomes even more misleading, rewarding the displacement of gasoline without consideration of the impact of the electrical consumption. The EPA reporting procedures have evolved to consider this, and fuel economy stickers for PHEVs report the CD and CS operation modes separately, allowing the consumer to consider the impact of these ratings on their individual driving habits. The UF-weighted fuel economy and its impact on the EcoCAR evaluation process is discussed in Chapter 7.

The concept of UF-weighting is also equally applicable to evaluating vehicle emissions, both upstream and tailpipe. Emissions data collected during the FCT and CST test procedures are weighted according to the method describes above. The ability to run on grid-sourced electricity complicates evaluating the environmental impact (in EcoCAR measured via GHG and PEU produced per kilometer) of a vehicle because electricity is produced by a wide range of methods that vary geographically. For example, a plug-in hybrid or full electric vehicle charged in Canada (where much electricity is produced using hydro-electric, nuclear, and natural gas) would have a significantly different environmental impact than a vehicle charged in China (where coal plants are a primary electricity producer). The fact that these upstream emissions are the sole source of GHG and PEU for plug-in powertrains, and that they vary so widely, means that unlike combustion engines it is not straightforward to evaluate their environmental impact.

2.5 Project-Based Learning

Researchers have found that employers are looking for more than technical competence from engineering graduates [4]. Engineering is a complex profession, involving multi-disciplinary teamwork, competing demands from various stakeholders, resource management, uncertainty, and incomplete data [4]. The field of engineering evolves quickly, and practicing engineers must keep skills up-to-date to be effective. It has been indicated that the traditional educational model does not provide the knowledge integration, practical application, teamwork and communication, or social, environmental, legal, and economic exposure that practicing engineers require [4]. Educators and education researchers have raised the concern that the current engineering education model, where knowledge is “transmitted” from teacher to student in a lecture format and tested in examinations, is unable to provide students with the aforementioned experiences. Alternative models of education, based around project- and/or problem-based learning, have been proposed as a solution.

The traditional engineering education model, derived from experience in the Mechanical Engineering curriculum, is well-described by Perrenet [2]. In the traditional model students are directed to learn core technical concepts (e.g. math, physics, chemistry) in a lecture format. Students internalize this knowledge by practicing on closed-ended problems from the lecturer or a text book (assignments) and in directed tutorial sessions. Directed laboratory exercises are used to give students an opportunity to apply or observe the lecture material in a real-world situation. Student progress is largely assessed through written examinations, in addition to grading performance on the aforementioned assignments and labs.

Nonaka proposed a framework for considering the ways in which tacit and explicit knowledge is converted through the four modes of socialization, internalization, externalization, and combination [34]. The traditional learning model emphasizes internalization of explicit knowledge taught by the lecturer. Students practice lecture material in assignment problems, converting explicit to tacit knowledge which becomes personal to the student and applicable in broader situations.

The traditional format has the benefit of a controlled learning environment, where each student receives all of the desired material in the sequence that the lecturer has deemed necessary or preferable. This sequencing is considered to be important in engineering education, where the body of knowledge is hierarchical [2]. For example, machine design courses build on concepts learned in materials science, mechanics of deformable solids, and dynamics. However, this format only emphasizes technical competence and does not encourage or develop the other skills considered essential for working engineers.

Problem and project-based learning methods reverse the traditional learning structure. In project-based learning, projects become the primary focal points of engineering courses and the vehicles for learning technical concepts [3]. These methods are distinguished from the traditional model by the focus of the learning experience (problems/projects as opposed to technical concepts) and by the changed roles of teachers and students. The traditional model, as previously discussed, is an experience directed by the teacher. Lecture material, assignments, tutorials, and lab exercises are all delivered in the format and sequence desired by the teacher, directing students to obtain the desired knowledge in the desired manner. In problem and project based learning, the students become the directors of the learning. Again, Perrenet provides a good overview [2]. Teachers will introduce a problem or project that requires a solution. Students must identify the concepts and knowledge that are needed to solve the problem, seek out appropriate resources, and apply the knowledge to solve the problem. In this model learning is student-directed, and teachers become facilitators of that learning. Teachers pose and answer questions to help students arrive at solutions on their own instead of giving the answers away, and they ensure that the desired knowledge is encountered by each student.

Problem and project based learning are closely related concepts in the literature, and there are varied attempts to distinguish them. Savery indicates that in project based learning the teacher generates the problem to be solved and sets goal(s) in the form of an end product and specifications [35]. Students are asked to execute a project to meet the goal(s) set by the teacher. Learning is achieved through completing a procedure to generate the end product, and through the problem-solving that occurs along the way. Savery states that project-based

learning is distinguished from problem-based learning by the stronger role of the teacher in the learning process. Where teachers define the problem and the goal(s) in project-based learning, in problem-based learning the identification of the problem and goal-setting could be up to the student, which Savery states is a valuable skill that project-based learning does not provide.

Perrenet et. al. have a different view on problem vs. project based learning [2]. They consider projects to be complex real-world entities that take long periods of time and which result in a concrete product. They define project based learning to be focused on the application of knowledge, and problem-based learning to be concerned with the acquisition of knowledge. They also identify that projects require time and resource management capabilities, as well as the ability to differentiate roles and tasks. In direct contrast to Savery, Perrenet et. al. consider project-based learning to be more self-directed and less controllable than problem-based learning.

Despite differences in definition and execution, project based learning has some common and important themes. It is learning through the execution of projects, where students acquire knowledge as it is needed to solve problems inherent in the project. It can be self-directed, where students seek out, acquire, and apply knowledge on their own. It is a process facilitated by teachers, who guide students towards the desired knowledge by answering student questions and posing their own. It is learning within a specific context. It is a process that exposes students to skills that are outside of the core technical knowledge taught in classrooms, skills like time and resource management, role and task differentiation (a part of teamwork), and self-directed learning. Its focus on complex and multi-faceted projects provides an opportunity for subject matter knowledge integration. The concept of solving an “ill-defined” problem is discussed in the problem-based learning literature, and it seems that retaining this element in project-based learning can ensure that Savery’s problem-definition skills are exercised.

Project based learning can address some of the concerns surrounding engineering education that were raised at the beginning of this section. However, where project based learning has

an advantage over the traditional format outside of the acquisition of core technical skills, there are some potential disadvantages discussed in the literature. Hung et. al. pose that project-based learning, with a focus on a particular problem, may encourage too much depth at the expense of breadth [36]. In the less controlled environment of project-based learning it can be difficult for educators to ensure that the required concepts are learned by each student. The non-technical elements of projects, having significant time implications, can also make it challenging for students to devote enough time to the technical fundamentals. Linked to these concerns is the idea that the higher-order or top-level thinking emphasis in project-based learning will de-prioritize and reduce the acquisition of core technical knowledge [36]. Resistance from students and teachers is also discussed. Students are sometimes concerned about their ability to be self-directed and to learn the required concepts in a project-based environment, and teachers would be concerned with how to ensure delivery of all required material in the less controlled environment of project-based learning. Perrenet adds that the hierarchical nature of engineering knowledge makes a case for the more structured approach of traditional lecturing, and concurs that students may fail to construct the “right knowledge” in a project-based environment [2].

Some literature holds medicine as an example of a professional institution where education through problem-based learning has been widely and successfully adopted [4] [37]. However, while problem-based learning has not been shown to hinder success in that field, its benefit over the traditional lecture style has been questioned [38].

Experiential learning theory is closely related to problem and project based learning. Kolb provides an overview of the background on experiential learning, drawing from the works of Lewin, Dewey, and Piaget [39]. The Lewinian model for learning emphasizes having concrete experiences, where individuals can reflect on those experiences, extract concepts, and test them in new situations in a continuing feedback loop. Dewey’s learning model also exists as a loop, where observation of an experience, combined with knowledge obtained from last experiences, observations, or advice, is combined to form a judgement of the consequences of subsequent action. Piaget considers that individuals develop from a concrete to abstract view of their environment, and from an active to reflective mode of learning. A

continuing balance of accommodation and assimilation of experiences is required for learning and growth. From these works, Kolb poses that the pairs of concrete experience and abstract conceptualization, and active experimentation and reflective observation, form four abilities of learning from experience that are in conflict. Kolb then goes on to describe four learning styles resulting from combining these opposing learning modes: diverging, assimilating, converging, and accommodating [40]. Assimilating (abstract conceptualization and reflective observation, the collection, organization, and generation of ideas and concepts) and converging (abstract conceptualization and active experimentation) are closely related and tend to represent the learning styles of scientists and engineers, respectively. Both require abstract conceptualization to learn new information, but the converging learning style is focused on practical application of that information, as opposed to reflecting on and expanding upon it.

Bringing these learning models together, the University of Waterloo Alternative Fuels Team certainly qualifies as experiential and project-based learning, which often forms part of an engineering capstone project [41]. Students are given a design project to complete a piece of the vehicle design. They must manage their time and resources, as well as deal with significant uncertainty and incomplete data to complete the design on time and to the required specifications. The projects, as part of a larger vehicle design effort, tend to require a degree of teamwork across disciplines. The complexity of these projects usually demands an integration of skills and knowledge from various subjects taught in the engineering curriculum. Foremost, these projects are based around a real vehicle and offer enormous opportunities for practical, hands-on experience. Students engaging in this new environment move through the experiential learning processes described by Kolb. Students that are exposed to new experiences and who take the opportunity to reflect on those experiences gain new insights, knowledge, and skills that can be applied to new experiences. The self-directed nature of student teams means that UWAFST students also influence the learning environment; higher levels of student engagement lead to stronger learning, more knowledge generation, and a more enriched environment from which others may learn.

The design tasks discussed in this thesis exemplify experiential and project-based education principles. Each one is discussed not only on its technical merits but also how they reflect an experiential learning experience, and how they fit within the larger vehicle design experience.

Chapter 3

The EcoCAR Hydrogen Fuel Cell Plug-in Hybrid Electric Vehicle

UWAFT's entry into the EcoCAR competition is a Hydrogen Fuel Cell Plug-in Hybrid Electric Vehicle (FC-PHEV). The powertrain is contained in a 2009 Sport Utility Vehicle donated by competition sponsor General Motors, and is comprised of a hydrogen fuel cell, gaseous hydrogen storage system, lithium-ion battery pack, AC-DC battery charger, bi-directional DC-DC power converter, and two AC electric motors with DC-AC power inverters. Several supporting electronics are also required, including a high voltage buck converter to generate 12 V for auxiliary systems, cooling pumps, fans, and radiators, power distribution and switching boxes for high and low voltage systems, and various electronic control units (ECUs).

This complex and ambitious powertrain evolved significantly over the three years of the EcoCAR competition, adapting to the changing circumstances and uncertainty that are commonplace in a student team environment. This includes but is certainly not limited to the availability of team members, development of the required skills (which are typically learned via self-study by team members as they work towards their goals), and the painful realization that, despite the best efforts of everyone involved, that some design elements of the vehicle are not going to progress to the required level of performance and safety in time for the competition. This last issue was of particular importance in the third year of EcoCAR, where the team's custom high voltage DC-DC power converter could not be completed in time. This led to a significant re-design of the powertrain architecture, and in the end to a successful performance by the team, but this re-design did impact the final efficiency and performance of the vehicle, which will be discussed in future sections.

This chapter will describe the UWAFT EcoCAR vehicle, with a focus on the powertrain architecture and how it evolved into its final arrangement. The basic theory of operation for the powertrain and a high-level description of the vehicle's control system hardware and

software are also presented. UWAF T’s project-based environment and the design projects that will be studied in further detail are also presented.

3.1 Powertrain Architecture

The final UWAF T EcoCAR powertrain architecture evolved significantly over the three years of the competition. In particular, a significant set of changes in the final months of the competition were required to get the vehicle operational. The design decisions that led to the final architecture shown in Figure 9 are discussed in Chapter 5. The integration of all powertrain components can be seen in the underbody vehicle view provided in Figure 54 of Appendix B.

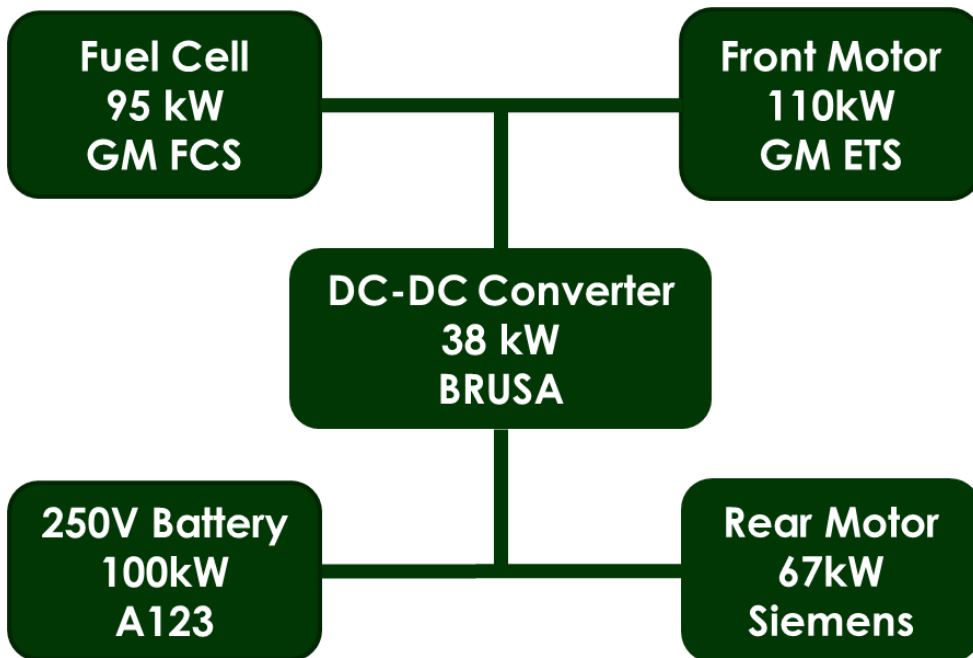


Figure 9: UWAF T EcoCAR year 2 powertrain architecture

All the essential components for a fuel cell powertrain are present, namely the fuel cell stack and supporting balance of plant components (collectively the fuel cell system or FCS), hydrogen storage, batteries, a DC-DC power converter to control power transfer, and two electric motors for propulsion. Key specifications for each powertrain component are listed in Table 2.

Table 2: EcoCAR powertrain component specifications

COMPONENT	SPECIFICATIONS	
General Motors Hydrogen Fuel Cell System (FCS)	Voltage Range:	240 – 400 V
	Maximum Output Power:	95 kW (375 A @ 250 V)
	Voltage Range (operating):	240 – 400 V
Quantum Technologies Hydrogen Storage Tanks (HSS)	Maximum Pressure:	70 MPa (10,000 psi)
	Tank Capacity:	4.2 kg
	Tank Volume:	114 L
Brusa BDC412 DC-DC Power Converter (DC-DC)	Voltage Range:	125 – 425 V (high side) 100 – 400 V (low side)
	Peak Current:	150A (37.5 kW @ 250 V)
	Continuous Current:	100A (25 kW @ 250 V)
General Motors Permanent Magnet AC Electric Motor + General Motors Power Inverter (ETS)	Voltage Range:	240 – 450 V
	Peak Input Power:	130 kW (325 A @ 400 V)
	Continuous Input Power:	95 kW (235 A @ 400 V)
	Peak Output Power:	110 kW
	Continuous Output Power:	80 kW
	Peak Output Torque:	350 Nm
	Peak Output Speed:	12,000 RPM
	Gear Reduction:	9.76 (fixed)
Ballard 312V67 AC Induction Motor + Rinehart PM100 Power Inverter (RTS)	Voltage Range:	260 – 360 V
	Peak Input Power:	80 kW (318 A @ 250 V) (est.)
	Continuous Input Power:	40 kW (159 A @ 250 V) (est.)
	Peak Output Power:	67 kW
	Continuous Output Power:	33 kW
	Peak Output Torque:	190 NM
	Peak Output Speed:	1250 RPM
	Gear Reduction:	11.58 (fixed)

3 x A123 25S2P Lithium-Ion Battery Energy Storage System (ESS)	Voltage Range:	234 – 253 V (open-circuit)
	Peak Discharge Current:	400 A
	Continuous Discharge Current:	120 A
	Peak Charge Current:	200 A
	Continuous Charge Current:	40 A
	Capacity:	39 Ah
	Energy:	9.68 kWh
Brusa NLG513-Sx Battery Charger	Input Voltage Range (AC):	90 – 264 V
	Peak Output Power:	3.3 kW

In total, the fuel cell and battery can generate up to 195 kW (260 hp) of electrical power. This power is made available to two electric motors, together capable of consuming up to 210 kW (280 hp) of electrical power and generating up to 540 Nm of combined rotor torque, or over 5600 Nm of axle torque. A total of 150 kWh of energy is stored on-board. Each component and its role in the operation of the EcoCAR vehicle are described in the following sections.

3.1.1 Fuel Cell System

The fuel cell system (FCS) is a General Motors (GM) 4th generation design. This system was produced in limited quantities to power a fleet of approximately 100 Fuel Cell Chevrolet Equinoxes for project ‘Driveway’. The fuel cell system is comprised of two hydrogen fuel cell stacks as well as supporting balance of plant equipment such as the hydrogen injectors that supply the anode, high-speed air compressor that supplies the cathode, a coolant pump for the stacks, and a power management and distribution module (PMD) that houses contactors, bus bars, fuses, control circuitry, a 12 V buck converter to supply the vehicle’s low voltage systems, and various switched high voltage outputs that provide power to high voltage auxiliary components such as the coolant pre-heater and fuel cell coolant pump. The fuel cell system is shown with key parts highlighted in Figure 10 and Figure 11. The supervisory controller interacts with the FCS controller over CAN (see Figure 18).

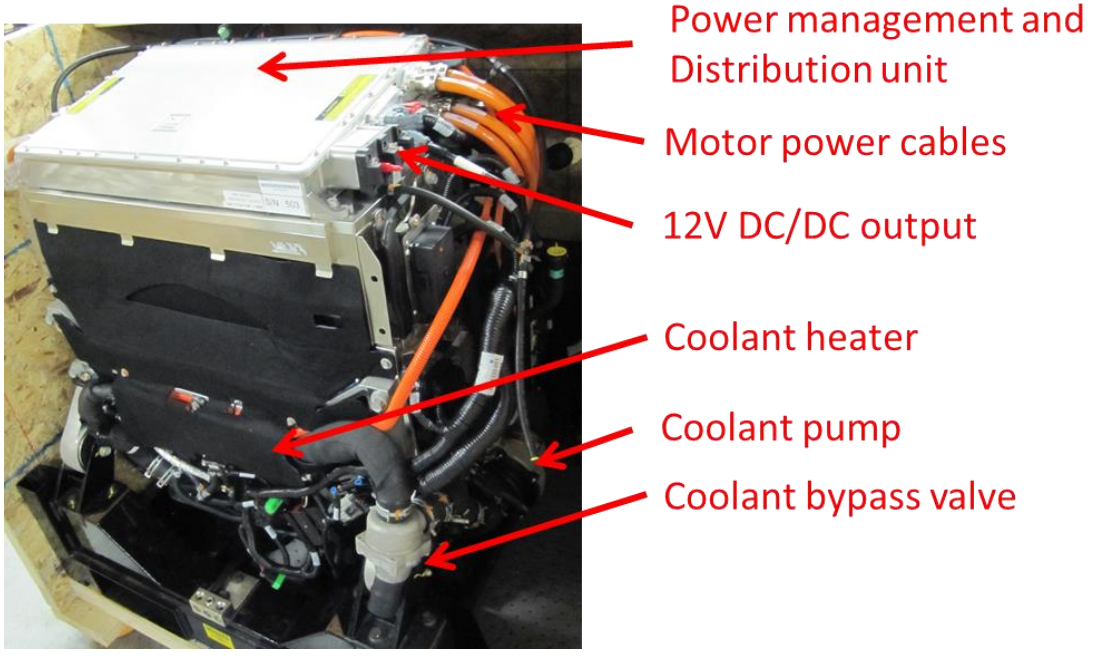


Figure 10: Fuel cell system, front view




Figure 11: Fuel cell system, rear view

3.1.2 Hydrogen Storage System

The hydrogen storage system (HSS) is also sourced from the GM Fuel Cell Equinox, a requirement of the donation agreement under which the fuel cell stack was procured. The HSS is comprised of three composite tanks capable of storing hydrogen at pressures of up to 70 MPa (10,000 psi). The HSS also includes a hydrogen controller and the associated sensors and actuators needed to monitor the status of the hydrogen tanks (tank and line pressures and temperatures) and control the tank's various valves. The HSS is pictured in Figure 12. The supervisory control system interacts with the hydrogen controller over CAN (see Figure 18).




Figure 12: Hydrogen storage system in its delivery crate

3.1.3 Front Electric Traction System

The front electric traction system (ETS) is the third and final component sourced from the GM Fuel Cell Equinox. It is mechanically integrated with the FCS, as indicated in Figure 11. The ETS is comprised of a permanent magnet alternating current (PMAC) electric motor connected to an AC motor controller that is powered from the high side of the high voltage DC bus. The supervisory controller interacts with the ETS over CAN (see Figure 18) to command torque, set operational limits on speed and current, and to select the motor state.

The digital control signal “Torque Enable” is a secondary enable signal used as an emergency stop. The front traction system is pictured in Figure 13.




Figure 13: Front electric motor

3.1.4 Rear Electric Traction System

The rear electric traction system (RTS) is assembled from a Ballard 312V67 AC induction motor and a Rinehart PM100 AC motor controller. The AC induction motor has an integrated differential, is lubricated by an internal electric oil pump, and is water-cooled. It is mounted directly on the rear axle and connects to the rear wheels directly through half-shafts. The Rinehart motor controller is connected to the low side of the high voltage DC bus and outputs 3-phase AC power to drive the induction motor. The Rinehart is commanded by the supervisory controller via CAN (see Figure 18). Control signals are sent to command motor torque, to set limits on motor speed and current, and to set the state of the traction system (on/off, forward/reverse). The rear motor and inverter are shown in Figure 14.




Figure 14: Rear electric motor integrated into the rear sub-frame assembly

3.1.5 Lithium-Ion Energy Storage System

The energy storage system (ESS) is a lithium-ion battery pack, designed and built by UWAF. The author was particularly responsible for the overall design, construction, and integration of the pack, and specifically contributed the final casing design, final structural support finite element analysis, the high voltage wiring design, the cooling system design (including custom-manufactured thermal plates), and the vehicle power and control system integration. The pack contains three 25S2P lithium-ion battery modules supplied by A123 Systems. As the name implies, each module contains 2 parallel strings of 25 series lithium-ion cells, for a nominal voltage of 83 V. In series, the three modules provide a pack with a nominal voltage of 250 V and a capacity of 9.68 kWh. The pack is water-cooled via custom cooling plates sandwiched between the modules. The battery pack is managed by an A123-supplied controller. The battery controller is mainly responsible for the operational safety of the pack. It monitors cell temperatures, voltages, and currents, and will automatically disconnect the pack from the main DC bus in the event of a critical fault such as cell

over/under-voltage, cell over-temperature, etc. The battery controller actuates the contactors that connect the pack terminals to the main DC bus to ensure that the battery can be safely isolated in the event of a fault. The battery controller interfaces with the vehicle control system via CAN (see Figure 18). The battery pack (before removal of the fourth module), is shown in detailed CAD imaging in Figure 15.




Figure 15: UWAFT designed and built lithium-ion battery pack, final CAD assembly with team-designed DC-DC (left) and installed in vehicle (right)

3.1.6 DC-DC Power Converter

In a fuel cell powertrain the DC-DC converter connects the two sides of the high voltage bus, providing a key, but limited, power transfer link to fulfill several critical functions that both hybridize the powertrain and enable it to satisfy several key requirements. The Brusa BDC412 bidirectional DC/DC converter acts as the bridge between the high and low sides of the DC bus, and performs the critical task of controlling the voltage (and hence power flow) on both sides of the high voltage system. The supervisory controls commands the DC-DC converter over CAN (see Figure 18). The BDC412 can convert up to 200 A of current bi-directionally (from low-to-high or high-to-low). A key limitation of the BDC412 is that while it can convert current bi-directionally, the high and low sides of the bus cannot switch over. That is, the voltage on the high side of the BDC412 cannot fall below the voltage on the low side; if such a condition occurs irreparable damage to the converter can result. The converter has internal control logic to hold off this condition if bus conditions are driving towards it (e.g. driving additional current from low-to-high in an attempt to hold up the high side

voltage) however these measures can only handle brief transient events, so the vehicle bus must be designed and controlled to prevent a “voltage cross-over” from occurring. This requirement plays a key role in the hybrid control system discussed in Chapter 5. The Brusa BDC412 DC-DC power converter is shown installed in the vehicle in Figure 16.




Figure 16: Brusa BDC412 DC-DC converter

The DC-DC power converter was originally intended to be custom-designed by UWAFI. When the design could not be completed on time, the Brusa off-the-shelf commercial solution was sourced. This introduced a few significant changes in the powertrain layout and consequently had significant implications on the powertrain power and torque control algorithms. The changes are discussed in Chapter 5.

3.1.7 Battery Charger

The battery charger enables the vehicle to harness externally-sourced grid electricity. The charger used in this vehicle is a Brusa NLG513-Sx air-cooled device. It can run from a 120 V and 15 A (standard household) circuit, effectively supplying about 1 kW of charging power at a Level 1 charge rate. It can also connect to a 208 V (industrial high voltage) or 240 V (residential high voltage) circuit, allowing for peak 3.3 kW charging power at a Level 2

charge rate. The charger is controlled directly by the battery management system over CAN (see Figure 18), which allows the battery controller to limit the charging rate as a function of pack voltage. The NLG513-Sx battery charger is shown in Figure 17.



Figure 17: Brusa NLG513-Sx battery charger

3.1.8 Auxiliary Components

There are a number of auxiliary components that play a key role in the operation of the fuel cell hybrid powertrain. The fuel cell system has numerous auxiliary components (known as the “balance of plant” or BOP) such as hydrogen fuel injectors, a high voltage coolant pump, and a high voltage coolant heater. A high-powered buck converter, connected to the high side of the DC bus, is also required to supply 12 V for vehicle auxiliaries (lights, radio, wipers, controller power, etc.) typically supplied by the alternator / lead acid battery system in a traditional vehicle. Various low-voltage cooling pumps and fans throughout the vehicle enable thermal management. Hydrogen sensors are located inside the cabin, at the fuel tanks, and in the engine bay to detect leaking hydrogen before the concentration becomes combustible.

3.2 Powertrain Control System

3.2.1 Control System Hardware

The vehicle's control system hardware has previously been discussed in [42], and is divided into two groups: supervisory and component control. Component controllers are developed by the manufacturer of each powertrain component to regulate its function (e.g. fuel cell controller regulates hydrogen and air delivery to the anode and cathode, respectively). This level of control is largely untouched by UWAFT, with a few exceptions where tunable parameters exist. Supervisory control is entirely developed by UWAFT to control and receive feedback from individual component controllers, to coordinate the overall operation of the vehicle, and to ensure the safety of vehicle occupants.

The supervisory control hardware consists of the main supervisor and a subordinate actuation controller. The main supervisor contains the vast majority of developed software and makes all powertrain operational decisions. The subordinate actuator acts as an extension of the main supervisor, receiving commands to drive various electrical circuits that are either not supported by the main supervisor or which are offloaded to the actuator unit for consistency. The actuator has very little decision making capability, and is basically an I/O extension for the main supervisor.

The vehicle control hardware layout is shown in detail in Figure 18. The vehicle relies heavily on controller area network (CAN) bus networks for communication between the component controllers and the supervisor. The vehicle's original high-speed GMLAN bus (hsGMLAN) and powertrain expansion bus (PTECAN) are retained for body systems control, power steering, and electronic brake control. A dedicated network is required for the interaction between the fuel cell controller, hydrogen storage controller, and the supervisor (FCPS CAN). The majority of the powertrain components (motor controllers, battery, some hydrogen safety sensors) exist on the H2CAN bus. A simple two-device CAN network links the rear motor controller and the supervisor due to interference that was experienced between the two motor controllers. There are also two data logging busses, FCPMi (dedicated to the fuel cell system) and LOG (EcoCAR's standardized data logging network).

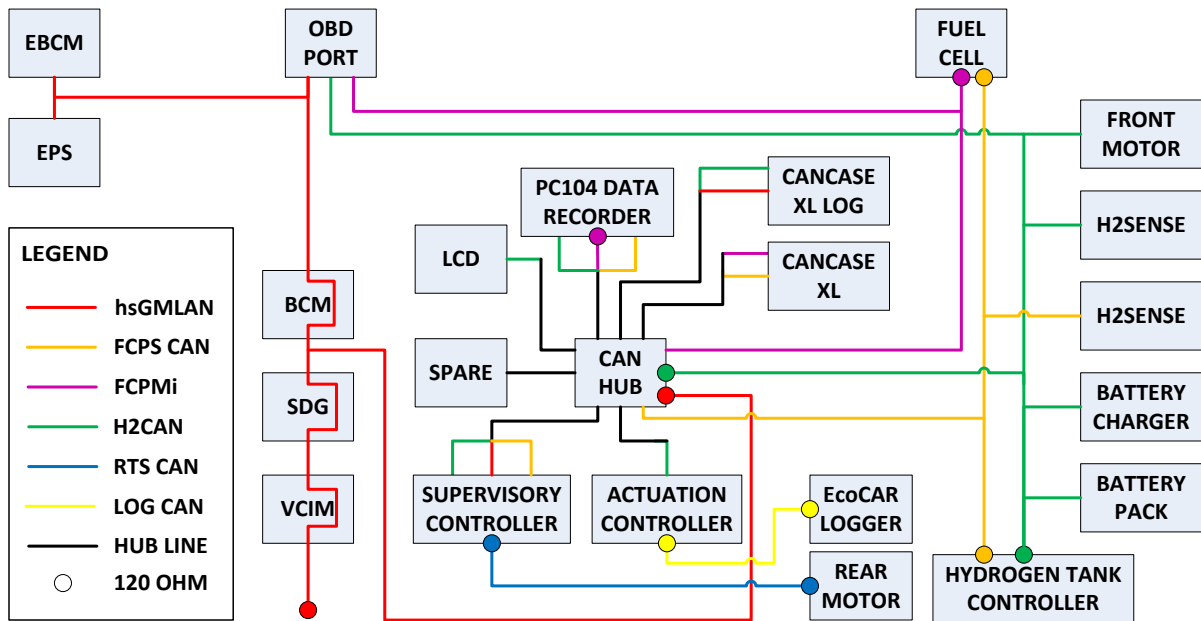


Figure 18: Vehicle-wide CAN bus topology

Some supervisory control is implemented using discrete low voltage signal lines. Both 5 V and 12 V circuits are used. 12 V digital logic signals are typically used to manage the power state of component controllers (i.e. wake up signals) and 5 V analog signals are typically read from onboard sensors (e.g. accelerator pedal position).

3.2.2 Control System Software

The supervisory control software has previously been discussed in [42], is divided into a hierarchy, sub-dividing responsibility for various aspects of supervisory control into units with defined interactions. This system architecture is a continuation of the work done in the first two years of the project by the previous team leader Alex Koch [43], who originally defined some of the subsystems and interactions between the control strategy and component control units. This work is extended to define four units at the top-level of the supervisory control code: Sensors, Diagnostics, Control Strategy, and Control Algorithms. Each unit is sub-divided into smaller control blocks, as shown in Figure 19, that handle specific functions within the unit. Those functions may be comprised of one or more algorithms that accomplish the required function. This division of control responsibility is designed to organize the control code logically, simplify development by breaking the system down into

manageable and distributable pieces, and to enhance safety by separating the safe actuation of powertrain control signals from the optimization routines that calculate their desired values. In this way, optimization of the control strategy (typically with the aim to improve energy conversion efficiency or drivability) can be carried out without concern for the effect on the safety of the powertrain and the passengers, which is the responsibility of the Control Algorithms unit.

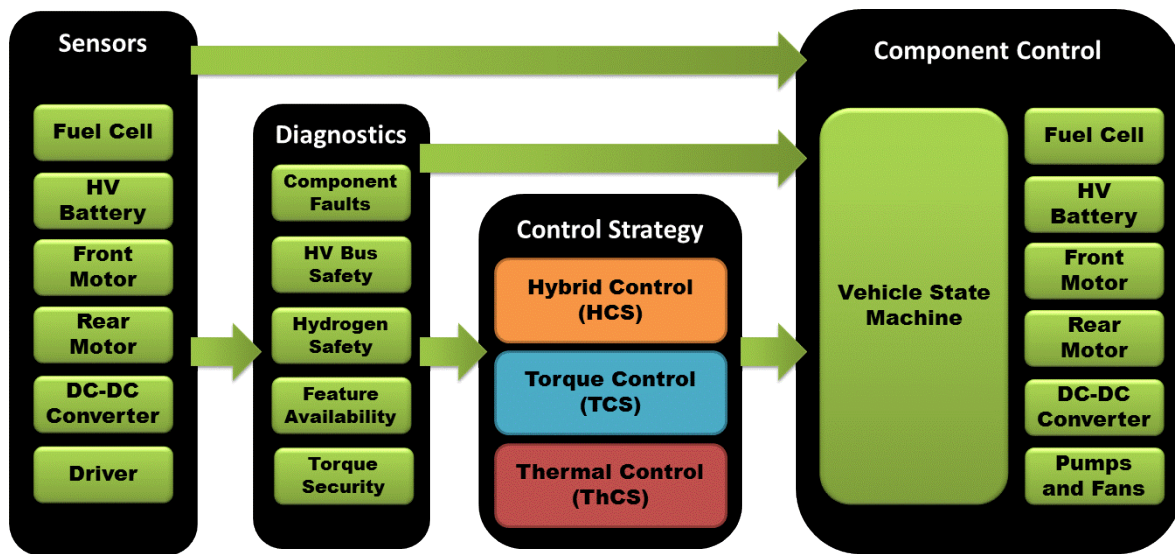


Figure 19: UWAFTEcoCAR supervisory control schematic

The sensors subsystem acquires CAN, digital, and analog signals and applies the appropriate scaling and offsets to convert them into engineering units. Sensor values are also filtered where necessary (e.g. to smooth a jittering signal). The diagnostic control subsystem is responsible for monitoring the status of the vehicle. It processes raw sensor data and creates flags (usually single bit variables) which are used by subsequent processes to determine if certain actions are allowed or if mitigating actions are necessary to maintain safety.

Powertrain component diagnostics are of concern for safety and performance. The diagnostic subsystem monitors the fuel cell, high voltage battery, high voltage DC-DC converter, and both motors for both expected responses to actuation commands and specific fault flags sent by component controllers. Component diagnostics are used both to determine when actions can be taken (e.g. when is it allowed to send a start-up signal to the component) and to take

remedial actions to maintain safety. Examples of remedial actions include emergency fuel cell shutdown in the event of a hydrogen leak, disabling one of the electric motors when CAN communication is lost, and limiting torque output when the DC-DC converter is overloaded (overcurrent condition).

High voltage bus safety in the diagnostic subsystem is mainly concerned with maximum and minimum voltage levels found on the high voltage busses. Voltage overloads can cause internal component shorts, and under-voltage conditions can have impacts ranging from reduced performance to catastrophic component damage. The diagnostic subsystem monitors multiple high voltage sensors and provides subsequent control units with status flags including “bus discharged” (confirms high voltage bus is dead after contactors open), and “high side under-voltage” (which tells the control strategy unit that there is not enough power on the high side to support the loads, and initiates a load reduction strategy).

Torque security monitors the accelerator and brake pedals, the gear selector, communication health between the supervisor and each electric motor, and the actual torque output as compared to the commanded torque for each electric motor. Its purpose is to ensure that torque is only released to the motors when it is safe and expected by the driver, and to ensure that the released torque is the amount that the driver expects. The accelerator pedal has redundant sensors that are monitored to ensure that they are in agreement (health check on the sensors while in range) and to ensure that they are in a valid range (e.g. identifies shorts to power and ground). The gear selector is monitored to ensure that it is in the appropriate position for torque release (Reverse or Drive, and not Park and not Neutral). The estimated torque output is monitored via CAN communication with each motor, and compared to the requested torque. A mismatch can be detected and appropriate mitigating action taken.

The control strategy unit contains the power split (hybrid control strategy), the torque split (traction control strategy), and the thermal control strategy. The power split strategy determines the desired fuel cell output power as a fraction of the overall instantaneous power demand, the torque split strategy computes the desired division of the overall driver

requested torque between the two electric motors, and the thermal control strategy determines the desired state of all pumps and fans as a function of component temperatures.

The component control unit receives sensor, diagnostic, and control strategy inputs, and determines what control actuation is necessary to carry out the desired control within safe limits. It is comprised of subsystems that represent the state of each individual powertrain component and of the vehicle as a whole. Each subsystem is typically executed using state flow diagrams, which provide the exacting control over the order and timing of events that is crucial for coordination between components.

Chapter 4

Thermal Management System Design

Powertrain thermal management is a significant challenge in electrified vehicles. A major cause for this challenge is the power electronics and fuel cells that are involved. These components, while more efficient than internal combustion engines, tend to have much lower temperature requirements. The resulting waste heat, although lower in quantity than that emitted by a combustion engine, exists at a much lower quality and is therefore harder to dissipate to the ambient. Radiators with heat dissipation ratings much higher than those typically found in a combustion engine vehicle, along with higher air flow, are needed to reject the waste energy. The vehicle thermal management system is shown schematically in Figure 20.

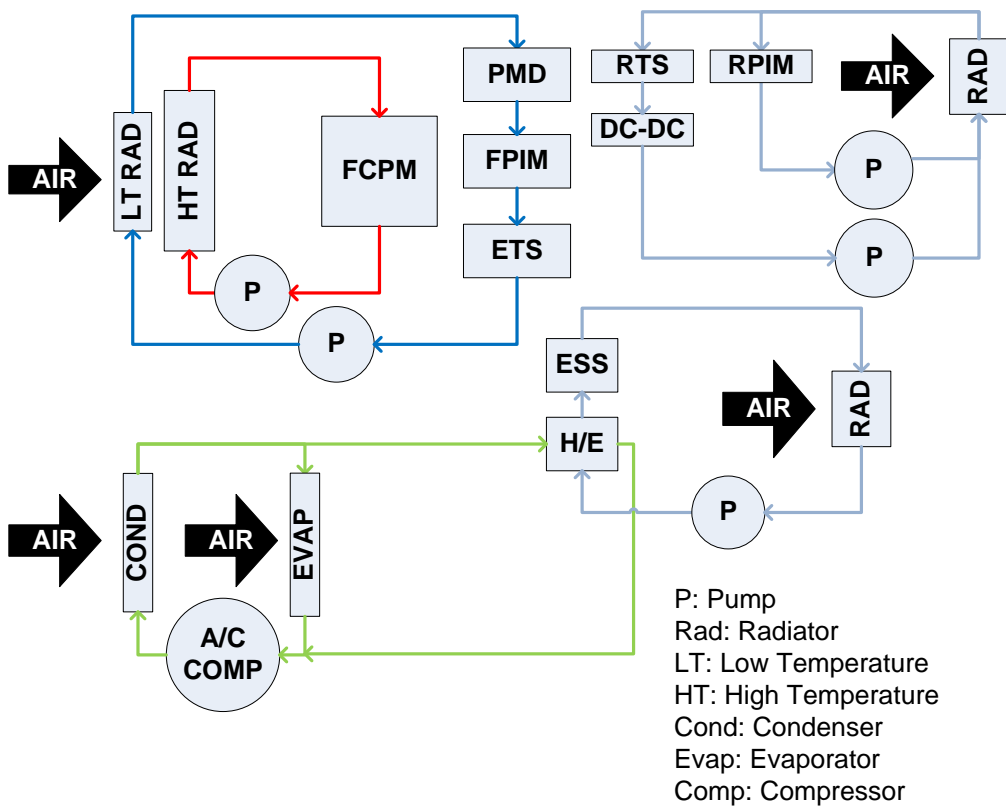


Figure 20: Vehicle thermal management system

Four separate cooling loops were designed to meet the thermal needs of all powertrain components. The need for several separate loops was driven by several factors:

1. The large number of powertrain components distributed throughout the vehicle, which if joined together into a single loop would require a great deal of pipe length incurring significant flow rate losses (or very high pumping power requirements).
2. The greatly varied temperature requirements of the various components. Joining “hot” components like the fuel cell stack with relatively “cooler” components like the battery pack in a single loop would require a very high flow rate.
3. It is easier to procure and integrate a larger number of small radiators than to try to use a single large radiator. The only location on the vehicle with enough space to house larger radiators is the engine bay, and it does not contain enough space to house a single or multiple radiators of sufficient size to dissipate all of the powertrain waste heat.

4.1 System Design Process

The thermal control system design project exemplifies several of the key attributes of project-based learning. First, the project is highly open-ended; there are no pre-determined set of components, arrangements, or design criteria to follow, only a set of thermal specifications for each component. The project involves incomplete data, with components often missing desired specifications like flow rate, and variables like total pipe length and equivalent pressure drop are difficult to determine exactly. There is significant uncertainty, for example in the quality of the analysis that results from inexact data and consequent assumptions. The integration of knowledge from several subjects is required, including various topics in fluid mechanics and heat transfer as well as the ability to collect, organize, analyze, and visualize data, and to make recommendations from said analyses in the face of the aforementioned uncertainty. Finally, students must interact with industry to select components that meet the needs of their design. A gap in terminology and knowledge tends to exist between the classroom theory that students are familiar with and the methods of industry.

Fourth year Mechanical Engineering undergraduate students tend to take on the thermal control system design project (either in whole or part) for their fourth year design course. The project directly relates to core subject matter from the curriculum, and therefore appears to be a straightforward task. In fact the ME381 design course taken in the third year uses a heat sink optimization project as a vehicle for teaching the design process, and closely resembles UWAFB thermal design projects. However, after two years of observing many students tackle this project it became apparent that many students were not prepared to execute this level of project. Students, (the author included) initially have great difficulty handling the large and unstructured scope of the project. The number of possible ways to combine components, and loop arrangements proved to be a significant barrier to progress. Even when the project was broken down into single cooling loops with defined powertrain components, the incomplete and uncertain nature of the available data, combined with a lack of a clear “recipe” for the task of solving for a set of inter-related (and initially unknown) variables, made it impossible for most students to find a suitable starting point.

In assisting students with the thermal design project, a design process for the vehicle cooling loops is developed to enable a structured and methodical approach. This process is applied to each of the four cooling loops, enabling rapid and accurate selection of pumps, fans, and radiators that meet the needs of the system. This design process is interesting for several reasons. First, it serves as a body of knowledge that can be used by the team in future thermal design projects. More importantly, development of the process itself is a valuable tool for students, helping them to deal with various types of uncertainty and to make the transition from analyzing engineering problems to designing in the real-world.

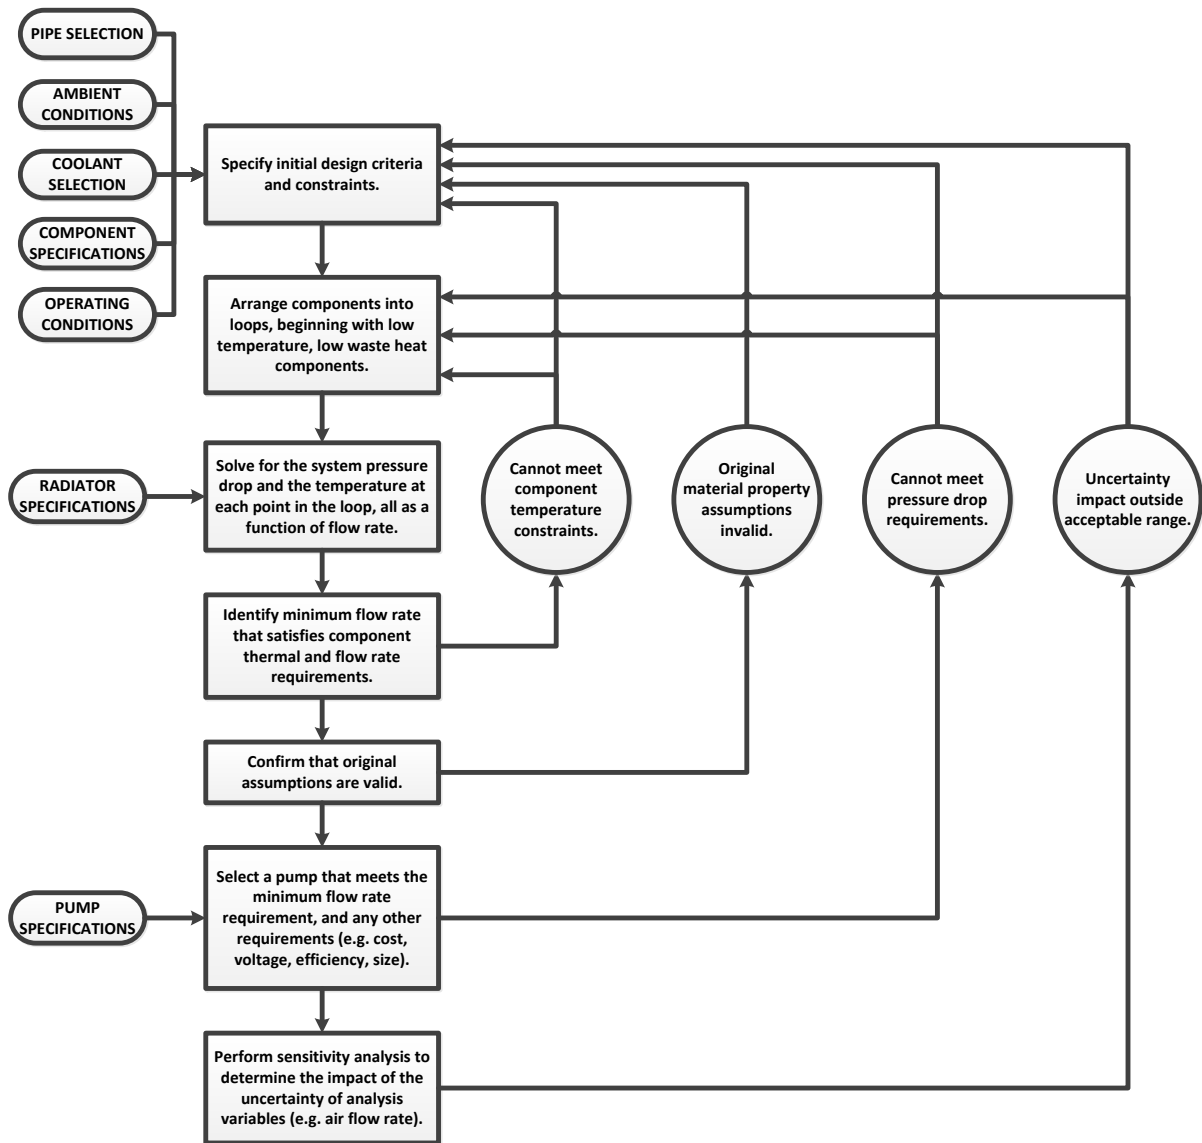


Figure 21: Thermal management system design process

4.2 The Low-Temperature Front Cooling Loop

The low-temperature front cooling loop is responsible for removing waste heat from the front motor power inverter (PIM), the front motor drive unit (DU), and the electronics inside the power management and distribution system (PMD). The thermal requirements for these components are outlined in Table 3. Subsequent sections will apply the thermal control

system design process to this design problem to demonstrate how it applies the principles of project-based learning and yields a beneficial learning experience.

Table 3: Thermal management specifications for front loop components [44][45][46]

COMPONENT	MAX. COOLANT INLET TEMPERATURE (°C)	MIN. COOLANT FLOW RATE (L·min ⁻¹)	COOLANT PRESSURE DROP AT MIN. FLOW RATE (Pa)	PEAK THERMAL POWER (W)
PIM	65	13	10,000	3,942
DU	69	13	15,000	11,158
PMD	68	5	15,000	2,000

4.3 Component Arrangement

Generally, the components should be arranged in order from smallest to largest in terms of both maximum temperature and dissipated thermal power. The goal is to find the arrangement that maximizes the coolant temperature entering the radiator, which in turn maximizes heat transfer to the ambient air. This will permit the smallest possible radiator to be used, where a smaller heat transfer coefficient is needed to dissipate a given amount waste heat.

Looking at Table 3, the temperature order is PIM, PMD, and DU. The thermal power order is PMD, PIM, and DU. The PIM and DU are physically co-located, so placing the PMD in-between requires significant additional piping. The initial attempt at loop arrangement was driven mostly by finding the shortest physical piping path, resulting in the order of PIM, DU, and then PMD.

4.4 Specify Conditions

The ambient conditions, cooling fluid, pipe properties, and component specifications must all be determined before the system's thermal-fluid equations can be solved. Coolant and pipe selection are arbitrary choices that should aim to maximize the effectiveness of the system and meet the requirements of the components and the environment. For example, water has a higher specific heat than the typical 50/50 mix of water and glycol, which enables the vehicle

to run cooler in the summer at the expense of the ability to run in sub-freezing temperatures. Additionally, some of the components in an electrified powertrain have strict requirements for the type of coolant and piping material that are used. The fuel cell stack, for example, is highly sensitive to conductive ions in the coolant. Certain materials used in conventional brazed radiators and metal piping can leach ions into the coolant, resulting in an increase in its electrical conductivity. If the coolant conductivity gets too high the isolation between the fuel cell high voltage system and the vehicle chassis is reduced to the point that the maximum permission leakage current is exceeded, and the vehicle must shut down for safety purposes. Fortunately, for the front loop components the material requirements are not strict, and traditional automotive brazed aluminium radiators and rubber heater hose are permitted. The ambient conditions and component thermal waste heat are a function of the environment that the vehicle will operate in, and the drive cycle(s) that will be encountered. Clearly, running the vehicle in a drag race at continuous wide open throttle generates more heat than a leisurely drive in a residential area. Additionally, the radiator will need a larger heat transfer coefficient to reject a given thermal power during summer in the Arizona desert than it would during winter in Toronto.

The first iteration of the design process aims to reject the peak thermal load of all components in the most challenging environment that the team expects the vehicle to operate. This may not be a frequent or even likely scenario, but as the worst case it is a good starting point for the analysis. If these loads cannot be dissipated, or if optimization of the system is desired, then average thermal dissipation rates can be generated from the team's Simulink powertrain simulations. The peak thermal loads are already listed in Table 3. The coolant specifications, pipe properties, and ambient conditions are listed in Table 4, Table 5, and Table 6 respectively.

Table 4: Physical properties for water and ethylene glycol (50%) at 70°C [47]

COOLANT	DENSITY ($\text{kg}\cdot\text{m}^{-3}$)	SPECIFIC HEAT ($\text{J}\cdot\text{kg}^{-1}\cdot\text{K}^{-1}$)	KINEMATIC VISCOSITY ($\text{m}^2\cdot\text{s}^{-1}$)
Water	977	4067	4.030E-07
50/50 Glycol	1029	3398	1.039E-06

Table 5: Properties of pipe system at common diameters

PIPE SIZE (in)	INNER DIAMETER (m)	ABSOLUTE ROUGHNESS (mm)	PIPE LENGTH (m)	NUMBER OF 90° BENDS
1/2"	0.0127	0.07	5	15
3/4"	0.01905	0.07	5	15
1"	0.0254	0.07	5	15

Table 6: Properties of air at various temperatures [47]

TEMPERATURE (°C)	DENSITY (kg·m ⁻³)	SPECIFIC HEAT (J·kg ⁻¹ ·K ⁻¹)	KINEMATIC VISCOSITY (m ² ·s ⁻¹)
30	1.1649	1006.5	1.604E-05
40	1.1275	1006.9	1.698E-05

The first iteration of the analysis assumes peak thermal loads for each component. The ambient air temperature is set to 40°C, which is representative of conditions at the General Motors test track in Arizona, where the year 2 competition is held. A pipe size of 3/4" is selected as it fits most of the components without the need for adapters and is commonly available. 50/50 glycol is selected as the initial cooling fluid, as it enables year-round operation of the vehicle and hence simplifies maintenance requirements. These assumptions can be varied on further iterations if required.

Of particular analytical importance is the assumption of the coolant temperature when specifying the density, viscosity, and specific heat. These properties are actually functions of temperature. The analysis can proceed by making an educated initial guess of the average system coolant temperature; however the loop coolant temperatures resulting from the analysis must be compared to the initial assumed value. If the properties of the assumed and solved fluid temperature are significantly different then the analysis must be repeated with new property values. Alternatively the coolant properties can be expressed as a function of temperature, and a computer program (e.g. MATLAB) can be used to iteratively solve the

thermal-fluid equations. Assuming initial values is valuable when learning to perform the analysis, however, because it allows the solution to be visualized concretely at each step.

It is important to note that in setting up the problem to this point many assumptions have been made regarding materials, sizes, and operating conditions. Making assumptions regarding the configuration and conditions of the design is one of the important ways of dealing with uncertainty that students can use to help move the design forward.

4.5 Component Pressure Drop Equations

The fluid mechanics part of the analysis is concerned with solving for the pressure drop in the system. The pressure drop, as a function of flow rate, is dependent on the coolant properties (density, viscosity) and the flow restriction of the components and interconnecting pipes. The coolant properties, at the initial temperature assumption, are already known. The flow restriction of the piping and components are a function of the flow rate. For the piping, the pressure drop is determined by conventional correlations for the loss coefficients in straight (Darcy-Weisbach) and curved (K-factor) pipe. These equations are provided in Table 18 of Appendix C.

Losses in powertrain components are typically provided as a single value at the minimum recommended flow rate; full pressure drop curves are hard to come by. In this case it is ideal to experimentally determine the pressure drop in each component, but this is not always possible at the initial design stage. Thus, an assumption is required to move the analysis forward. Since pressure drop is typically a quadratic function of flow rate it is possible to extrapolate an estimate of the losses in each component as a function of flow rate by fitting a quadratic curve through the single provided data point and through zero. This is clearly an estimate, but in the absence of better data it is the best estimate that can be made and helps students move the design forward in the face of uncertain or unavailable data. The loss coefficient “K” for each component is calculated as $\text{Pa} \cdot \text{LPM}^{-2}$ from the provided data point. The pressure drop curve from each component is then simply $P = K \cdot Q^2$.

The combination of the pressure drop equations for the components and piping results in the system pressure drop. Equations for the pressure drop within components and piping are provided in Table 18 of Appendix C. At this stage the radiator has not yet been selected, thus the total system pressure drop cannot be fully calculated. Suitable choices for the radiator are derived from the thermal analysis of section 4.6. The resulting radiator(s) are then used to complete the system pressure drop calculation.

4.6 Loop Temperature Equations

The objective of the loop temperature calculation is to find a radiator that can reject enough heat to ensure that the coolant temperature at each component that is below the maximum rated value. The temperature at each point in the loop is a function of the ambient conditions (air temperature), the coolant properties (density, specific heat capacity), the component operating conditions (waste heat load), and the radiator performance (heat transfer coefficient) that is a function of coolant flow rate. The ambient conditions, coolant properties (at the initial assumed temperature), and component waste heats are known at this stage. The radiator heat transfer coefficient “U” is obtained from manufacturer data sheets, and is a strong function of flow rate.

The temperature of the coolant entering the radiator is derived from the radiator heat transfer equation. The temperature of the coolant at subsequent points in the loop is expressed from the heat gain of the coolant as it passed through each component. These equations are provided in Table 20 of Appendix C.

Each equation has two unknowns, the inlet and outlet coolant temperatures, which link them together in order along the loop. Thus the equations are easily solved in an iterative fashion to arrive at the set of coolant temperatures at each point in the loop that result as a function of the third unknown, coolant flow rate.

Radiator HX-260 is selected for the analysis. This radiator is manufactured by the Dana Manufacturing Corporation and is the only unit that fit within the physical constraints available in the engine bay. The heat transfer rate of the radiator is a function of the coolant flow rate, the air flow rate, and the temperature of the coolant and air as they enter the

radiator. The coolant flow rate and temperature are solved as a part of the analysis. The air temperature is simply that of the environment, which is also initially selected. The air flow rate, however, is extremely difficult to determine. Even with performance data for the radiator fan, the complex geometry within the engine bay makes it challenging to determine the resistance to air flow. One starting point is to use the air flow rate from a reference design of similar construction, which has an air flow rate of $2.4 \text{ kg}\cdot\text{s}^{-1}$, or 4510 CFM at $40 \text{ }^\circ\text{C}$. The HX-260 manufacturer provides cooling performance data at 4587 CFM, which will be used as a starting point for the analysis.

With a radiator air flow rate, the thermal equations can be solved. The equations and tabulated calculations are provided in Appendix C, and the resulting coolant temperature at each point in the loop is plotted as a function of flow rate is shown in Figure 22.

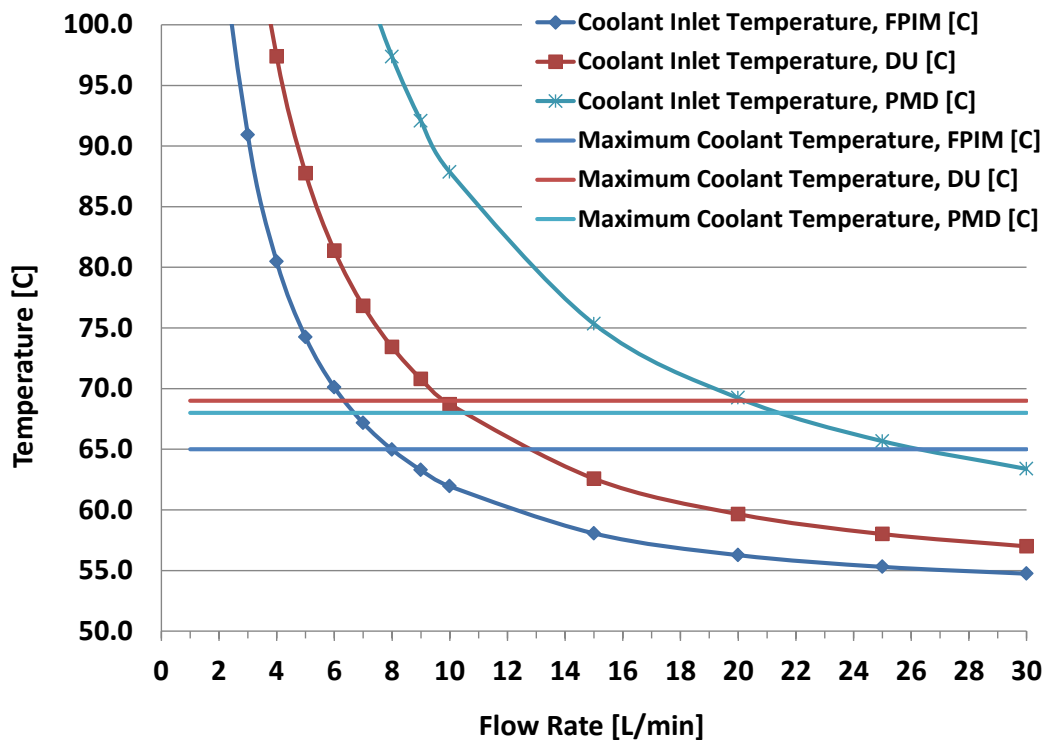


Figure 22: Coolant inlet temperature for each loop component as a function of flow rate

The output of this analysis is the minimum flow rate required to reject the total waste heat of the operating condition to the ambient environment, while keeping the coolant temperature at

each component within specifications. The total waste heat and ambient environment are built into the analysis as conditions to be met. The maximum coolant temperature for each component is plotted in Figure 22 to illustrate how the flow rate is selected. The flow rate must be such that the temperature at each component is at or below the maximum value. Shaded areas on the plot highlight flow rates that are insufficient to meet the temperature requirements of one or more components. The minimum flow rate in this loop arrangement and under these conditions is 21.5 LPM.

4.7 System Pressure Drop and Pump Selection

With a known radiator and minimum required system flow rate, it is possible to plot the system pressure drop curve and select an appropriate pump. Given the HX-260 specifications and the previous pressure drop calculations outlined in section 4.5, the system pressure drop curve is calculated and shown in Figure 23. Also shown are the pump curves for two commonly available automotive electric water pumps, the Bosch PAD12V and PCA12V.

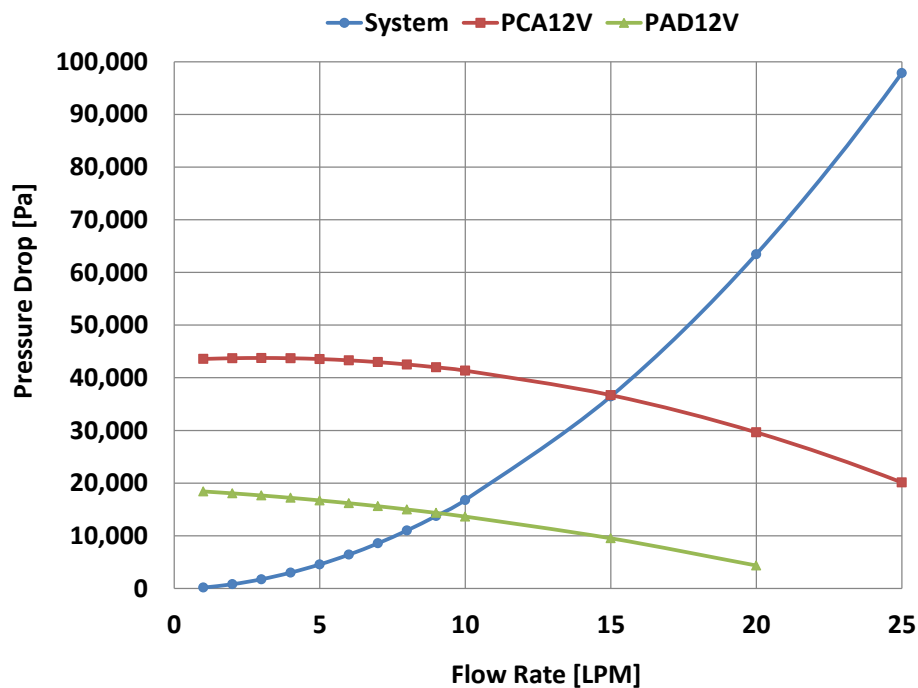


Figure 23: System pressure drop and pump curves

It is evident that the required flow rate of 21.5 LPM cannot be achieved by either pump. This leads to iteration on some of the initial assumptions regarding configuration and conditions; several solutions are possible. A simple solution may be to retain all of the underlying assumptions of the analysis to this point, and simply connect additional pumps in series to boost the flow rate to the required value. Another solution lies in refining the thermal dissipation requirements of each component through simulations of the powertrain on expected drive cycles. Another possibility is to re-examine the loop configurations to see if a more optimal arrangement is possible.

Further examination of the loop temperature calculations shows that the PMD drives the minimum required flow rate by a significant margin. The DU before it generates by far the greatest amount of waste heat and results in a significant temperature rise in the coolant. The flow rate must consequently be high enough to limit this temperature rise to meet the PMD's temperature requirements.

By re-arranging the loop so that the PMD precedes the PIM and DU, the DU's temperature rise does not need to be limited. In fact a higher temperature is desired to maximize the effectiveness of the radiator. The loop temperature plot for this new arrangement is shown in Figure 24.

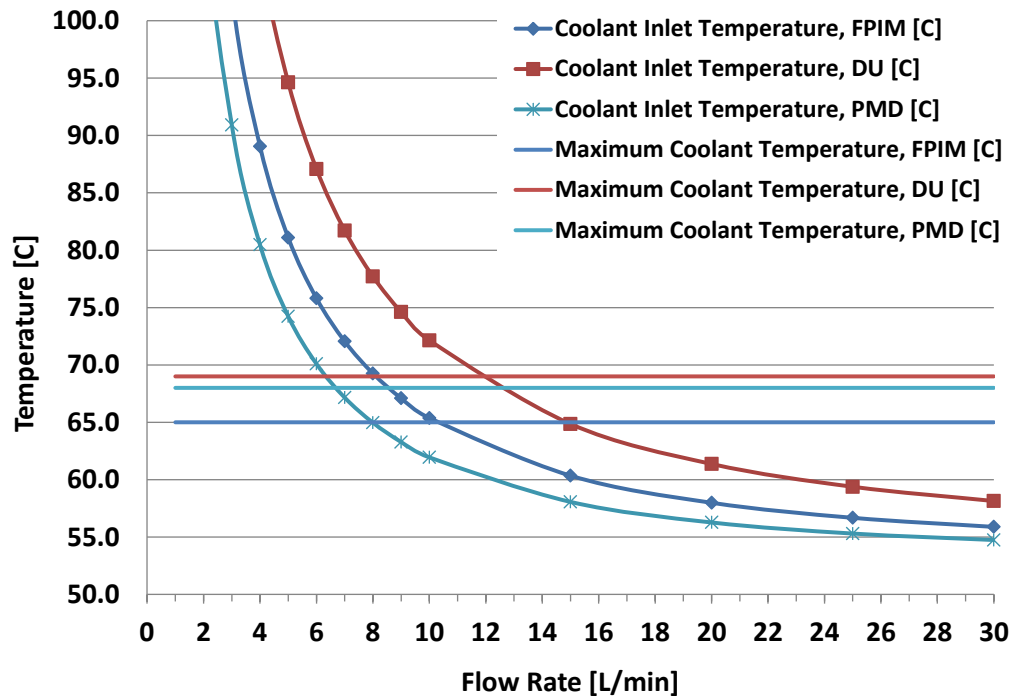


Figure 24: Loop temperature plot for PMD, PIM, DU arrangement

The new minimum flow rate for the loop is driven by the DU at 11.7 LPM. The relatively small waste heat load of the PMD makes it more ideally suited at the beginning of the loop, as it does not cause a significant temperature rise for subsequent components. With the components arranged as such, the PCA12V pump can be selected to operate at 15 LPM and meet the system requirements.

4.8 Sensitivity Analysis

Several variables had to be estimated to bring the analysis to this point. Most notably, the radiator air flow rate cannot be determined given the available information. The reference design flow rate of 4510 CFM is one way to estimate the performance of the cooling system. To better understand the impact of air flow rate on the performance of the system a range of values should be run through the analysis. Thermal calculations for alternative air flow rates provided by the radiator manufacturer are tabulated in Table 20 of Appendix C. The coolant flow rates required to meet the initial design conditions are summarized in Table 7 for each air flow rate. For the PCA12V pump, which can achieve 15 LPM, an air flow rate of about

2000 CFM is required. This is less than half of that required in the reference design. Further analysis or experimentation is required to determine if this air flow rate can be achieved.

Table 7: Minimum coolant flow rate at various radiator air flow rates

AIR FLOW RATE (CFM)	MIN. REQUIRED COOLANT FLOW RATE (LPM)
694	∞
2082	14.7
3469	12.4
4857	11.7
6245	11.2

4.9 Final Design Summary

With all of the system components now specified, the actual performance of the loop can be calculated (Table 25); the results are summarized in Table 8 for the original radiator air flow rate of 4857 CFM.

Table 8: Front loop final performance

FLOW RATE (LPM)	T_{c,i,R} (°C)	T_{c,i,PMD} (°C)	T_{c,i,PIM} (°C)	T_{c,i,DU} (°C)	PRESSURE DROP (Pa)
11.7	85.4	60.2	63.2	69.0	36,872

The front loop components are able to operate within their respective maximum temperature limits under full operating load within 40°C ambient conditions. The pressure drop within some of the components, the air flow rate over the radiator, and the properties of the coolant are assumptions that require validation. Other assumptions regarding configuration and conditions of the design include piping (size, layout, and length), ambient temperature, and the thermal load of each component. The design assumes a Dana Corporation HX-260 radiator that fits within the physical space constraints of the engine bay. Given these

assumptions, a Bosch PCA12V pump meets the flow rate requirements of the system so long as a minimum 2000 CFM air flow rate is achieved at the radiator.

It is evident that, given the available information, many assumptions were needed to advance the design to this point. These assumptions, so long as they are clearly stated and justified, and valid and necessary steps in the initial stages of a real-world design, yet they are difficult to for students to make when first starting a project at UWAFI. This often leads to procrastination for much of the term and a lost opportunity for students to advance their skills. Using this process encourages students to make and justify the assumptions needed to move the design forward, and gives them the confidence to address such problems in the future. Once the design is complete, students can then go back and iterate on multiple design configurations, operating conditions, and assumptions, and to visualize how all this affects system performance. For the author, this led to a better and more intuitive grasp of the material than had been previously gained from classroom instruction or labs.

4.10 Thermal System Design in a Project-Based Learning Environment

Uncertainty is common in engineering design projects, yet in the author's experience student design projects tend to have well-defined inputs and outputs, and exist as the design-based equivalents of analysis-based textbook problems. Designing the thermal management system for the UWAFI EcoCAR is reminiscent of the third year ME380 design course in the University of Waterloo undergraduate curriculum, which teaches the engineering design process through a project to develop an optimal heat sink design. ME380 does do a good job of teaching a process that includes conceptualization, prototyping, analysis and optimization, experimentation, and refinement. ME380 also requires students to integrate knowledge from fluid mechanics, heat transfer, and computational analysis, with the objective of generating optimal values for a set of design parameters that yields minimal mass and volume and maximum thermal dissipation. The equations needed to compute these parameters are simultaneously taught in the curriculum. Therefore, the ME380 design process focuses on the iteration and refinement of a few design variables using a known analytical process to achieve a specific set of goals; essentially iterative analysis.

The design experience at UWAFI provides a similar experience but with additional challenges faced by practicing engineers. Additionally, students are given a much higher-level design goal, such as “to design a cooling system for the powertrain components”. Students are then responsible for scoping the problem, which includes identifying and gathering the information, perhaps from previous projects, that is needed to perform the analysis. While support is always available from other students and faculty the onus is on the student to take the initiative, as is typical in project-based learning [2] [3]. This contrasts with ME380, where the scope of the problem is already defined. The ability to scope a design problem, that is to understand the goals that the design needs to achieve, has been identified as a key differentiator between students and practicing engineers [48]. Thus this form of uncertainty, not knowing the objectives of the design, is addressed in UWAFI projects.

After the problem is scoped students must generate a strategy to execute the design. In the thermal design example students must identify the various elements of the design, including materials, sizes, environmental conditions, and thermal loads. Students must then synthesize system configurations and construct an analytical process to evaluate them. Students are already familiar with all of the variables and equations needed to complete the design, and if the task was to execute any one of those analyses (e.g. pressure loss calculation or heat transfer calculation) students generally have no issue tackling the problem. However, UWAFI design projects like the thermal management system challenge students to combine these various simple analyses into a greater and more complex analytical framework to achieve a higher-level design goal. This form of uncertainty comes from not knowing how the variables and equations combine to yield the final design. The process described in this Chapter is one way to help students learn to organize and execute a complex design.

Finally, even with a well-understood design scope and a good formulation of the analytical process, there is often significant uncertainty in the value of variables being used or in how they are obtained. One example is the pressure drop within powertrain components. Due to a lack of detailed data and without having the components available for testing, some simplifying assumptions about the relationship between pressure drop and flow rate are made to permit the analysis to proceed. Other sources of this type of uncertainty exist in the length

of pipes, the ambient conditions that the vehicle will need to perform in, and the average and peak thermal loads that the cooling system will experience.

Thermal and fluid analyses are well-taught at Waterloo and UWAFST students typically have a good understanding of how to perform them, but stumble when they are unable to compute the “exact” answer that would usually come from a textbook problem. For the individual, executing a thermal design project at UWAFST helps students learn to deal with various types of uncertainty by asking questions about the scope of the problem (much like practicing engineers must do), by making (justified and documented) assumptions where information is lacking, by using “what-if” scenarios and sensitivity analyses to understand the impact of the assumptions that have been made.

The design experience described in this Chapter does more than yield a set of specifications, it yields a process that the team can carry forward as the design matures. For the team, having the design process laid out in this manner makes it easy to iterate the analysis to find the best solution to a single set of conditions and to make changes as conditions evolve. In the dynamic student team environment where members, requirements, and resources are constantly changing, process like this enable a knowledge transfer that can improve the team’s ability to execute designs over the years within and between competitions.

Chapter 5

Powertrain Control System Design

Controlling the unique powertrain described in Chapter 3 requires a detailed knowledge of the capabilities and limitations of the individual components, the range of operational conditions that the vehicle will experience, and how these factors combine to dictate the required flow of energy throughout the powertrain. This chapter will describe the electrical and control systems that were implemented to enable a safe, reliable, and predictable flow of power between vehicle components.

The high voltage layout of the powertrain is reviewed and the equations of power flow are combined with operational requirements to yield the algorithms for the various powertrain operational modes. The powertrain moves through these distinct modes of operation during starting up and as condition change during a drive. These modes are described in terms of the power flow equations, and their implementation in software is discussed.

Next, the algorithms governing control of the major powertrain components (fuel cell, battery, DC-DC, front and rear motor) are developed. These algorithms implement the control rules needed to achieve the various modes of operation by controlling the power flow through the high voltage system, following the power flow equations and enforcing the component limitations as needed to ensure reliability and safety. The two algorithms that will be discussed are the power control algorithm, which manages the power requests to the fuel cell, battery, and DC-DC, and the torque control algorithm, which manages the torque requests to the front and rear motors.

Control system design for the EcoCAR vehicle requires knowledge and skills outside of the mechanical engineering curriculum, drawing on both unfamiliar core engineering concepts (e.g. electrical circuit design and protection, software design, algorithm development, system validation, safety analysis) and application-specific knowledge (e.g. powertrain component control). To execute this project a great deal of self-directed learning is required in areas such as electrical circuit protection (wire and fuse selection for high and low voltage systems),

processing digital and analog I/O in an embedded microcontroller, implementing modular and scalable software architectures, writing Simulink control diagrams for auto-code generation, and developing and validating control rules at the software, hardware, and vehicle level. These skills are not taught in the classroom, but are learned as needed to enable completion of the project design tasks; a key feature of project-based education. Having hands-on opportunities to identify and learn new skills provides significant motivation and offers the chance to see the tangible results of that learning. Learning new skills in a project-based manner also helps turn what could be abstract concepts or theory (when taught in a classroom setting) into internal knowledge by providing a variety of experiences that can be used to engage the subject matter, in a more engaging manner than textbook problems.

Real-world engineering problems also often require trade-offs between efficiency, performance, and reliability. In this Chapter the powertrain architecture, power control algorithm, and torque control algorithms demonstrate how these considerations impacted the final implementation of the powertrain control system. Learning engineering concepts through projects like these allows students to gain experience making these trade-offs during the design process. In several cases the control system must impose restrictions on the performance of the powertrain (power output) in order to maintain a high level of reliability. These trade-offs have a negative impact the performance of the vehicle during competition events (see Chapter 7 for a review of the vehicle's on-road performance), but they ensure that the vehicle is able successfully complete them. This emphasis on safety and reliability is typical of real-world engineering projects.

5.1 High Voltage Architecture

It was previously stated that the team's custom DC-DC power converter could not be completed on time. The shift to a commercially available and (in terms of technical specifications) less capable Brusa DC-DC power converter necessitated several changes to the powertrain layout. The focus of this section is to outline the impacts of the new DC-DC converter, describe how those impacts were mitigated through a re-design of the powertrain

layout, and identify how the control system would be modified to support the new arrangement.

The original powertrain architecture used all of the same components, but the particulars of each component and how they were arranged within the powertrain were different. Figure 25 shows the powertrain architecture (top left) and the final architecture for comparison (bottom right). A complete high voltage system schematic can be found in Appendix D.

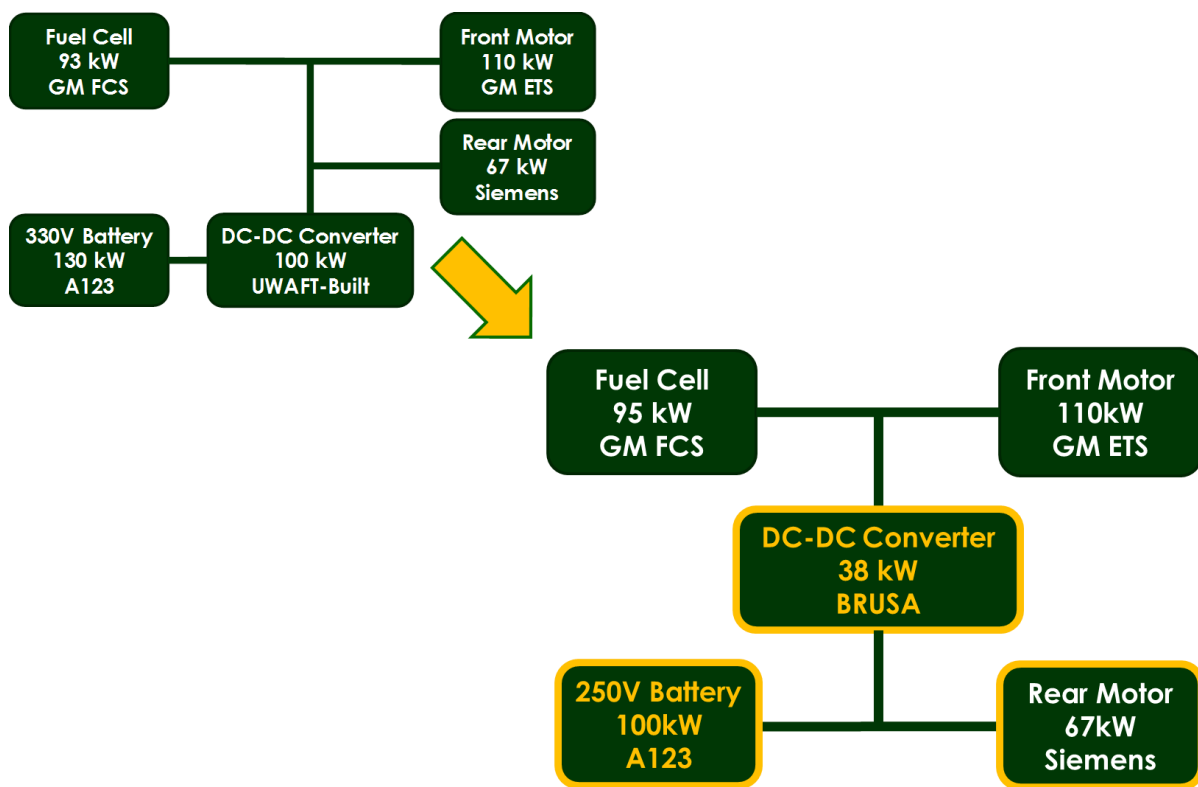


Figure 25: Transition from the year 2 to year 3 architecture

The fuel cell and both motors were connected to the high side of the high voltage bus. Only the battery pack, which contained four A123 25S2P modules, was connected to the low side of the bus. The DC-DC power converter bridged the two sides. The DC-DC power converter was to be designed by UWAFT to be bi-directional buck-boost. That is, it could move current in both directions (bi-directional) and the high and low side could be swapped at any time (buck-boost). The buck-boost mode switching requirement is further discussed in

section 5.1.2, and arises due to the overlapping operating voltage regions of the fuel cell and battery pack, caused by the polarization effect discussed in section 2.3.1.

5.1.1 Bus Voltage Ratio

The voltage ratio between the two high voltage sides of the DC-DC converter is defined as:

$$R = \frac{V_{BatterySide}}{V_{FuelCellSide}} \quad \text{Equation 5}$$

This ratio is a critical metric for a fuel cell vehicle. Due to electro-chemical processes within the fuel cell and the battery, the original architecture allowed for the condition where either the fuel cell or the battery could be the “high side” of the high voltage system during operation, thus the voltage ratio could be either less than or greater than 1. The cross-over point where the two sides are at equal voltages ($R = 1$) could occur at a wide variety of operating points given the dependence of polarization on temperature, state of charge, and other factors. To operate the fuel cell through the full range of power outputs required a DC-DC converter that could transfer current in both directions and which could switch between bucking and boosting in each direction, as shown in Figure 26.

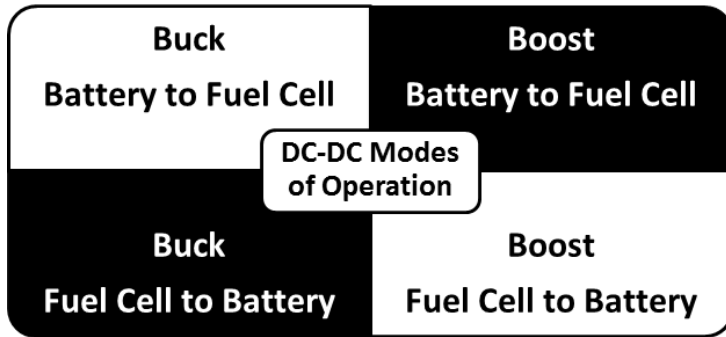


Figure 26: DC-DC modes of operation

The Brusa DC-DC converter requires that the high side and low sides of the high voltage system are fixed. Thus, the ratio R must always be either greater than 1 or less than 1 for the architecture. Crossing the point $R = 1$ causes uncontrolled current flow and can damage the converter. Physically, then, the Brusa BDC412 can only boost from the low side to the high side, and buck from the high side to the low side. Thus, the vehicle can be physically

configured to operate in only two of the four quadrants shown in Figure 26, depending on which of the battery and the fuel cell has the higher operating voltage.

The Brusa DC-DC, while bi-directional, does not have the ability to change buck-boost modes. The high and low sides of the converter are fixed. The following sections describe how the powertrain changed to accommodate this restriction while still maintaining the original capability of the vehicle.

5.1.2 Battery Pack Voltage Modification

The open circuit voltage of the fuel cell and the battery are different. Further, due to polarization the voltage of the fuel cell and the battery decrease as current increases. The shape of the polarization curve is a function of the electro-chemical device and the operating conditions (e.g. temperature).

The original UWAFTEcoCAR architecture used a 4 module, 330 V battery pack. Figure 27 shows the region of potential battery operating voltage over the range of battery discharge power levels. This region is derived from A123 test data of individual Li-ion cells. Measured fuel cell voltages corresponding to fuel cell discharge power during a drive are also plotted. It is evident that in this configuration the fuel cell output voltage easily drops below the battery voltage at power levels as small as 10 kW fuel cell output. It is also evident that the cross-over point is not easily predicted due to the large overlapping area of the two polarization regions.

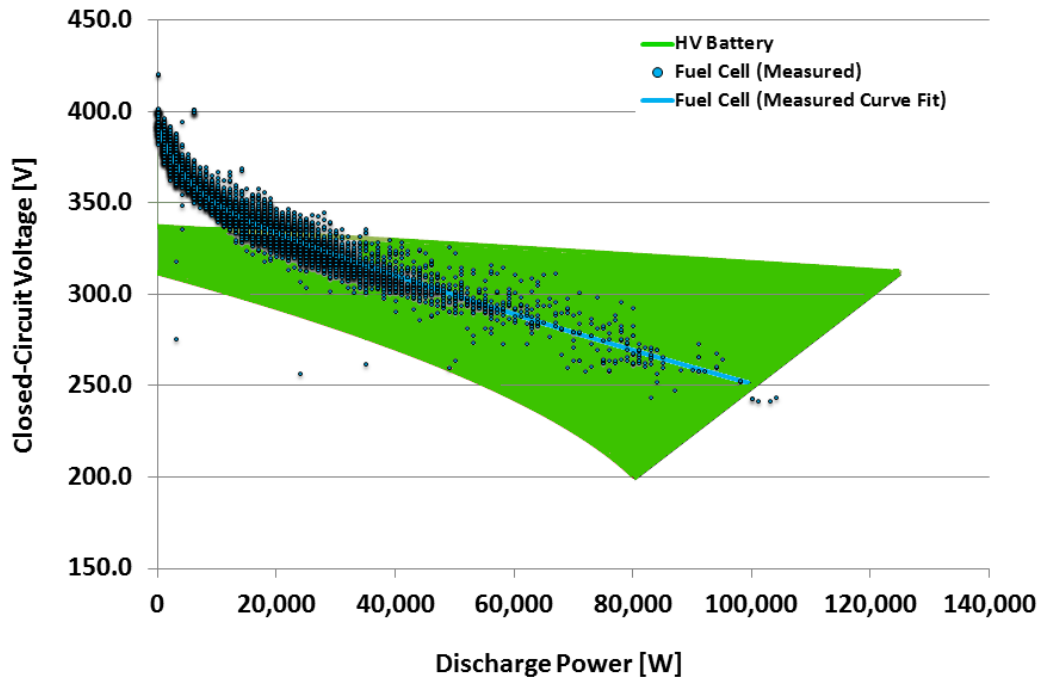


Figure 27: UWAFT EcoCAR source polarization potential, 330 V battery configuration

To accommodate the two-mode operation of the Brusa BDC412, this voltage crossover must be prevented. Three options were considered and have been previously discussed in [42].

1. Increase the battery voltage by installing a 5-module battery pack, so that the battery would be the high side of the high voltage bus at all times during operation.
2. Decrease the battery voltage by installing a 3-module battery pack, so that the battery would be the low side of the high voltage bus at all times during operation.
3. Retain the 4-module battery pack, adjust the control to operate the powertrain such that $R < 1$ at all times (i.e. control fuel cell such that its voltage is always greater than that of the battery).

Option 1 required finding room to add a fifth battery module, plus integrating the additional wiring and re-designing the entire support structure and skin of the battery pack to support it. There was not enough time to support this option and so it was discarded. This option required no changes to the control system, but the mechanical challenges were too great.

Option 3 was the simplest mechanically, requiring no structural modifications. However, after much discussion internally and with GM experts the team concluded that the significant controls system complexity needed to ensure safe operation of the powertrain, the risk of damage to the converter due to inadvertent violation of the voltage ratio restrictions, the severely reduced power availability from the fuel cell leading to serious operational restrictions, all made this option infeasible. In short, this option was mechanically trivial but had significant control system implications.

Option 2 was selected, sitting between the other two in terms of mechanical and control system modification effort. Removing a battery module was, mechanically, a relatively simple task. The only impact was a lower battery voltage, and hence a lower peak battery power output. This is more than compensated for by the gain in fuel cell power output, however. The battery and fuel cell polarization curves for the 250 V battery configuration are shown in Figure 28. Only a small possibility of violating the voltage ratio restriction exists, at low battery power outputs and simultaneously high fuel cell power outputs. Compared to the operational restrictions of Option 3, this is a relatively rare and easy situation to manage. By implementing option 2, only the top-right and bottom-left modes of operation (black background) described in Figure 26 allow the voltage ratio $R < 1$.

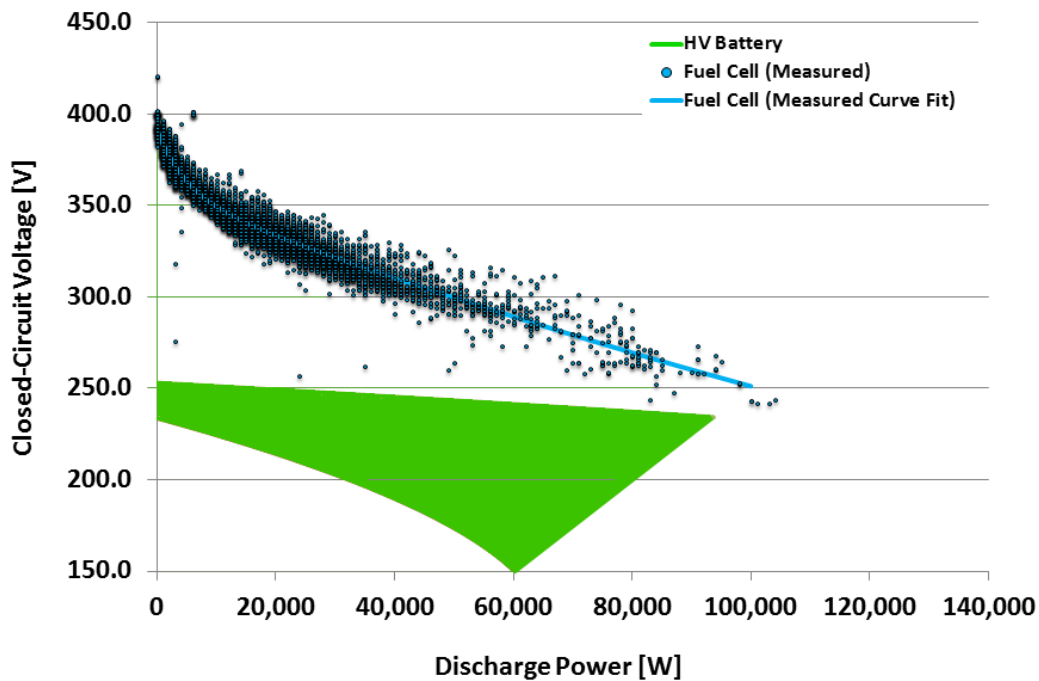


Figure 28: UWAFT EcoCAR source polarization potential, 250 V battery configuration

5.1.3 Rear Motor Re-Location

The smaller 38 kW DC-DC converter, if used in the original architecture, would have restricted motor access to roughly a third of the available battery power output, significantly limiting acceleration performance and practically negating the sought performance benefit of the rear motor altogether. To increase the useful battery power, the rear motor power inverter was re-located to the low side of the DC-DC converter. Thus, battery power can be used in full to operate the rear motor, without being limited by the DC-DC converter. In fact, given the component specifications listed in

Table 2 the rear motor could be saturated at its maximum power input and there would still be nearly enough battery power in reserve to also saturate the DC-DC converter in boosting current to the high side of the bus. This change did not incur any significant drawbacks or re-designs, only a few changes in how the power split control strategy calculated the net power production on each side of the high voltage bus.

5.1.4 Control System Implications

The control system was impacted in several ways by this new component arrangement. The power and torque control algorithms, which aim to balance the production and consumption of power in the vehicle, had to be modified to account for the new energy paths. The DC-DC component control block also had to be completely changed. The original team-designed DC-DC converter was meant to be current-controlled, with the voltages on each side of the bus being automatically managed by the converter to meet the magnitude and direction of the current being requested. The DC-DC component control block was also originally required to command the DC-DC mode switch when the fuel cell and battery voltages crossed over. In the new architecture, the Brusa BDC412 is a voltage-controlled device, and its primary control point is the high side bus voltage. Current will be converted automatically by the DC-DC to meet the voltage commanded on the high side. However, the operation of the power control algorithm relies on balancing electrical current (or power), not voltage, and Figure 28 shows that it is not easy to predict the fuel cell voltage required to obtain a given power output due to variable polarization characteristics. This is primarily why the original DC-DC converter was designed to be current controlled. A closed-loop feedback control that varies the high side the voltage request until the desired fuel cell power is reached was a potential solution, but this complex control would require much time and testing to ensure its speed and stability. A faster and more robust solution was to use the built-in current limit commands for the BDC412, in combination with the primary high side voltage command, to generate the desired fuel cell output current. This is discussed in further detail in section 5.3. The final architecture is very similar to a split-axle hybrid, having an independent propulsion system on each axle, but with two distinctions. Instead of a combustion engine there is a fuel cell and electric motor on the front axle, complimenting the battery and electric motor on the rear axle. Additionally, the two powertrains are electrically coupled via the DC-DC converter, instead of purely through the road.

5.2 Powertrain Modes of Operation

The placement of both energy sources and sinks on both sides of the DC-DC power converter yields a powertrain energy balance equation that is unique in current fuel cell architectures.

The possible current pathways within the vehicle are shown in Figure 29.

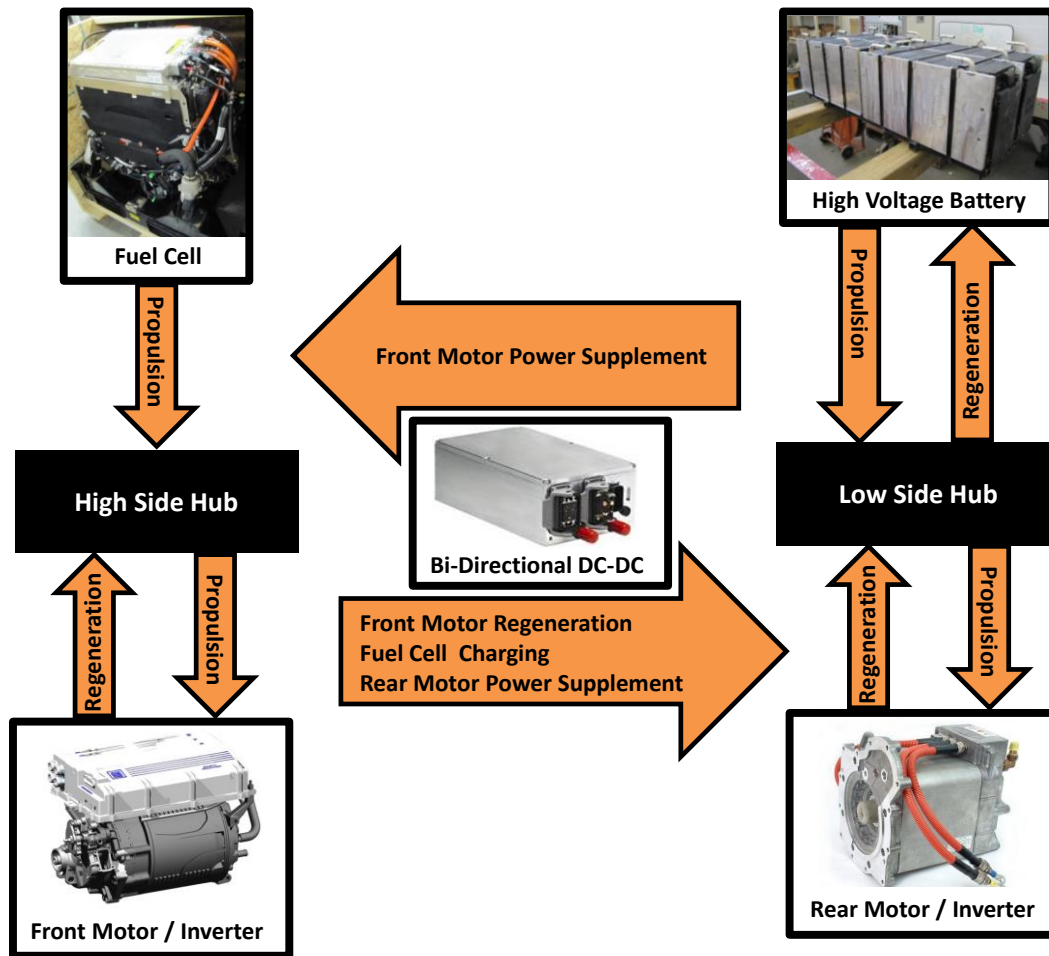


Figure 29: Current flow within the powertrain

From the energy flow diagram, the powertrain high voltage energy balance is developed. As a convention, a production of electrical energy is considered positive current, and consumption is negative current.

$$P_{ESS} + P_{FCPM} + P_{ETS} + P_{RTS} + P_{12V_DCDC} - P_{Loss,DCDC} = 0 \quad \text{Equation 6}$$

This expression simply demonstrates that all power produced at one point in the powertrain must be consumed at another point. Additionally, the DC-DC power converter is responsible for moving power between the high and low sides of the high voltage system, and this power transfer is expressed in terms of the difference in power generated / consumed on each side of the high voltage bus. For convention, the DC-DC transferred power is positive when boosting from the low side to the high side, and negative when bucking from the high side to the low side.

$$P_{DCDC} = (P_{ESS} + P_{RTS}) - (P_{FCPM} + P_{ETS} + P_{12V_{DCDC}}) - P_{Loss,DCDC} \quad \text{Equation 7}$$

The powertrain goes through various distinct modes of operation that are determined by where and how power must flow through the high voltage bus. The vehicle supervisory control software defines when each of these modes occurs and applies the appropriate power and torque control algorithms.

5.2.1 High Voltage Bus Pre-Charge

The only source of permanently energized on-board high voltage energy is the battery pack. The battery therefore must pre-charge the entire high voltage system before any high voltage components can operate. Thus, the DC-DC must be capable of being energized on the low side and subsequently energizing the high side from zero volts. Once a sufficient voltage is established on the high side, the start-up sequence for the fuel cell system can begin.

5.2.2 Fuel Cell Stack Start

The fuel cell system draws high voltage power from the battery pack to operate the balance of plant components (e.g. cathode air compressor) which enables the fuel cell stack to get to open circuit voltage in preparation to close its contactors and provide power to the vehicle. Up to 10 kW of power is required to start the fuel cell stack, which translates to 40 A at 250 V nominal battery voltage.

Additionally, the DC-DC converter must be able to “follow” the stack voltage during the start-up sequence. The fuel cell system requires that the high side bus voltage be within 25 V of the stack-side voltage before contactors are allowed to close. The final voltage of the stack

before contactors attempt to close is variable and depends on temperature, humidity, and other factors. Therefore the DC-DC must have the ability to control the high side voltage to a specific value (the stack-side voltage) as the high side component loads vary. Otherwise, failed starts may occur.

5.2.3 Propulsion – Excess High-Side Power

During normal propulsion there is a power balance which exists on each side of the high voltage system. An imbalance on either side is compensated for by transferring current through the DC-DC power converter to the other side. Every propulsion system component (with the exception of the fuel cell stack) can act as a source or a sink at any time. The system power balance is shown in terms of DC-DC power transfer below. All terms are signed positive when acting as a source and negative when acting as a sink. The DC-DC power term is positive when current is transferred to the low side.

For normal propulsion the control system is concerned with the power balance on the high side of the high voltage system. This is due to a combination of factors, including the slower fuel cell electrical inertia (compared to the battery) and the fact that the battery can act as a source or sink as needed without any battery-level control (it is treated as an automatic energy buffer). If, during normal propulsion the fuel cell power output exceeds the load of the front motor and the high voltage auxiliary loads, the DC-DC must be able to transfer the excess current to the low side (bucking mode), to be absorbed by the battery. This could be due to fuel cell inertia being too slow to respond to a reduction in the front motor load (heavy tip-out), which is shown in Figure 52 of section 7.1.

5.2.4 Propulsion – Excess Low-Side Power

If the fuel cell power output is insufficient to meet the loads of the high voltage auxiliaries and the front motor, current must be transferred to the high side (boost mode) from the battery. This typically occurs during heavy acceleration where the fuel cell cannot increase power output quickly enough to meet the motor demand. It may also occur during charge depleting operation, where the hybrid control strategy prefers a low fuel cell power output.

5.2.5 Regenerative Braking

During regenerative braking, the fuel cell is commanded to zero power output, which maximizes the regenerative braking power that can be transferred through the DC-DC power converter. The fuel cell will internally generate the power required to satisfy all non-motor high side loads (auxiliaries). The DC-DC converter is commanded to maintain fuel cell open-circuit voltage on the high side of the high voltage bus to accomplish this. Thus, maximum current generated by the front motor is transferred to the low side.

5.3 Power Control Algorithm

The power split strategy calculates the total instantaneous vehicle power demand and then determines how much power should come from the battery and how much should be provided by the fuel cell. The output of the power split strategy is the desired instantaneous power output of fuel cell power. It is implied that the battery will provide the balance of the power being consumed by the powertrain.

The power command algorithm uses the desired fuel cell power output, together with information about the state of each powertrain component, to determine how closely the desired power split can be met. The fuel cell power set-point is modified as needed to meet operational or safety requirements. The algorithm then generates the required control signals for the fuel cell and the high voltage DC-DC converter to match, as closely as possible, the desired power split. This algorithm has previously been described in [42].

The power command algorithm translates the desired power split between the fuel cell and the battery into the component control commands required to achieve them. It also overrides the desired power split when it cannot be achieved based on the state of individual components, modifying the desired power output only as far as is needed to maintain safety, while also maximizing the total power output when it cannot be achieved. A flow chart illustrating the calculation process of the power command algorithm is shown in Figure 30.

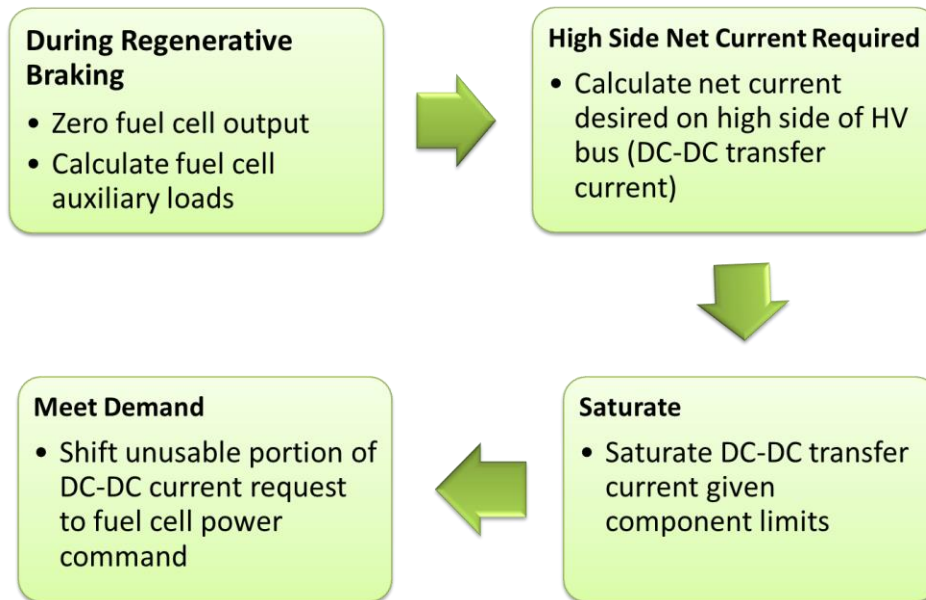


Figure 30: Calculation flow for the power command algorithm

The power command algorithm accepts the desired fuel cell power as a set-point command, and outputs a DC-DC current command and a fuel cell power request. It also receives the instantaneous battery current, DC-DC transfer current, front motor power requirement, fuel cell power availability, the voltage on both sides of the HV bus, and the regenerative braking system status. The top-level view of the algorithm is shown in Appendix E. In the first calculation, the power command algorithm computes the HCS_FCPM_PwrDesired variable, which is simply a reflection of the fuel cell power target for the power command algorithm (Figure 31). It is differentiated from the HCS_FCPM_PwrDesired variable by the desire to maximize the regenerative braking system efficiency. Since the fuel cell and the front motor must both use the DC-DC converter to sink any current produced in regenerative braking mode, it is desirable to idle the fuel cell system during regenerative braking events, thus maximizing the available throughput for regenerative braking current to charge the battery. Thus, if the RegenAtv variable is true, the HCS_FCPM_PwrDesired is zero, else it is the FCPM_PwrDesired.

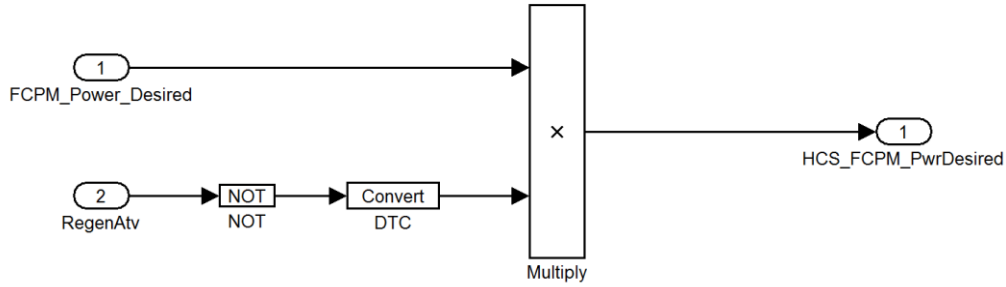


Figure 31: Fuel cell power desired

The next calculation step also works to maximize the use of available regenerative braking energy. Even during idle operation, the fuel cell system requires some power to maintain its balance of plant loads. Instead of allowing the fuel cell to produce this required electrical energy during a regenerative braking event, additional current can be commanded from the front motor. The fuel cell continually reports the required power for internal operation, which is translated into a current requirement by dividing by the voltage of the high side of the HV bus. This current requirement is stored in the variable FCPMcurrRq_Regen for later use whenever the RegenAtv variable is true (Figure 32).




Figure 32: Fuel cell internal current requirement during regenerative braking

The next calculation step computes the DC-DC current command required to achieve the desired power split as closely as possible, within the limits of the powertrain components. First, an energy balance is calculated for the high side of the HV bus, if the desired fuel cell power was to be implemented at the current system traction state (front motor power requirement). Any unbalanced power is output as a DC-DC power transfer requirement, converted to a current requirement with reference to the low side of the HV bus. If regenerative braking is active the FCPMcurrRq_Regen is subtracted from the DC-DC current requirement to supply some regenerative braking power to the fuel cell. The resulting

desired DC-DC current transfer is then dynamically saturated according to the instantaneous current limits reported by the battery (in charge and discharge) and the DC-DC converter. The minimum of these two component current limits is used to saturate the desired DC-DC current.

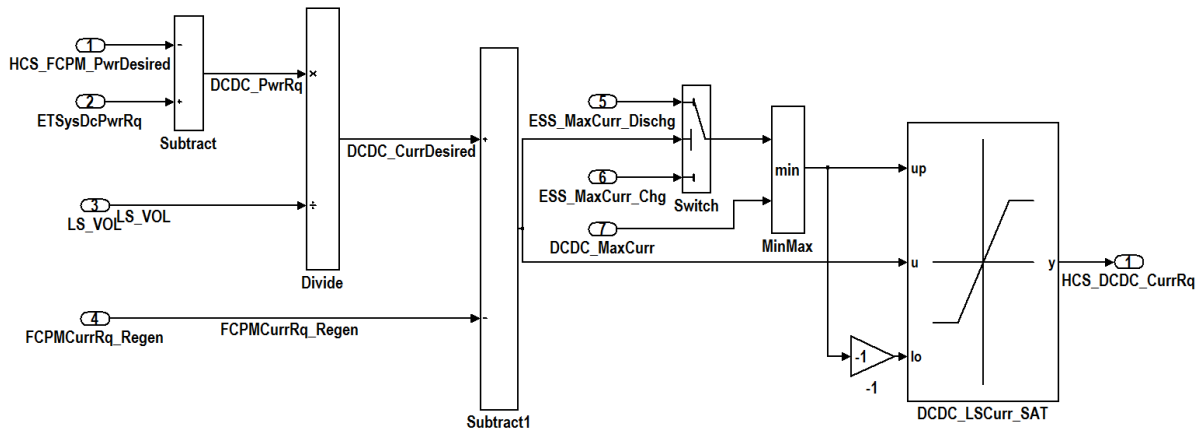


Figure 33: Power control algorithm DC-DC current calculation

The saturated DC-DC current request represents the best-case DC-DC transfer current available under the circumstances. The balance of the vehicle’s power requirements must be met by the fuel cell. The final fuel cell power request is simply the difference between the best-case DC-DC current transfer capability and the total vehicle power requirement (Figure 34). Note that the rear motor power requirement is not part of the power control algorithm’s calculations. It is assumed that the fuel cell set-point calculated by the hybrid control strategy is sufficient to meet the needs of the powertrain. The function of the power control algorithm is to meet this set-point as closely as possible, within the operational limits of the powertrain components. This requires a high side power balance, as shown. The battery pack has sufficient power capability, even in the reduced 3-module form, to supply all of the rear motor power needs (80 kW per, as per Table 2). This power is drawn freely from the battery without control system intervention. Maintaining the appropriate SOC of the battery is one of the primary functions of the hybrid control strategy. As the rear motor draws power from the battery and the SOC falls, the hybrid control strategy requests more power from the fuel cell to replenish it. If it is advantageous for the fuel cell to supply current to the low side for rear

traction, the hybrid control strategy will request this power. Thus, the power control algorithm needs only to balance power on the high side of the bus to meet the fuel cell power set-point provided by the hybrid control strategy.

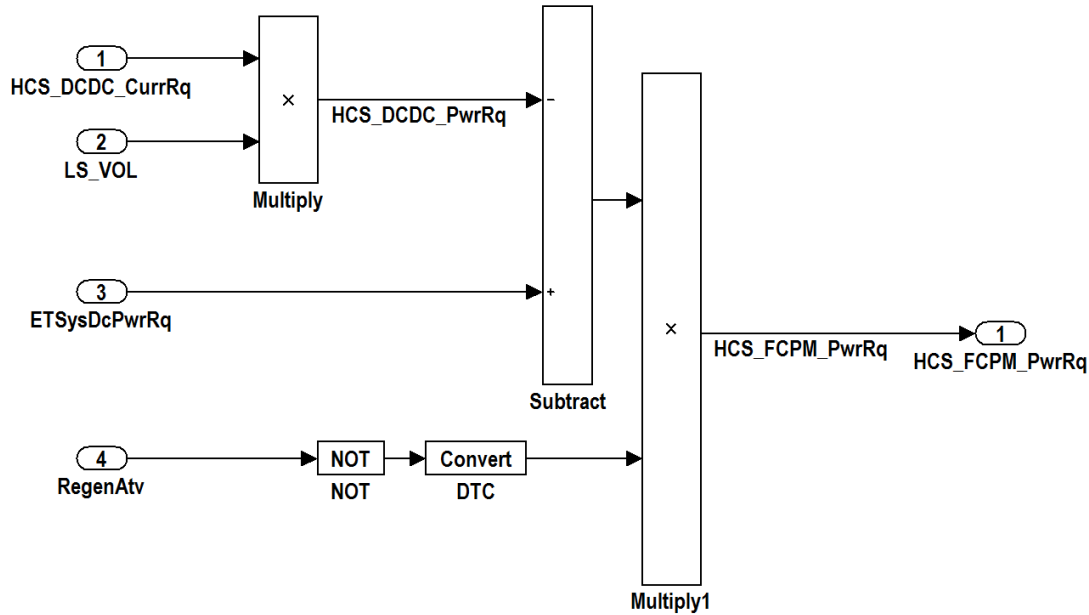


Figure 34: Fuel cell power request calculation

The DC-DC control block for the “driving” mode is shown in Appendix F. Depending on the sign of the DC-DC current transfer request (positive or negative) the DC-DC control code can infer the transfer direction (boost or buck, respectively). When boosting, the fuel cell voltage is forced to its maximum value, resulting in zero power output. When the front motor places a load on the high side the DC-DC converter maintains the maximum voltage set-point by boosting current from the battery, thus forcing the fuel cell to maintain zero power output. However, there is a boost current limit that can be sent to the converter, forcing it to saturate the DC-DC transfer current at the desired value. This current limit supersedes the voltage set-point. Hence, once the transfer current limit is reached the converter no longer enforces the voltage set-point, allowing the fuel cell to meet the remaining front motor load while continuing to provide the requested DC-DC transfer current. A similar strategy is used during bucking, with the fuel cell voltage set to a minimum allowable value. In this way the powertrain controller can indirectly but accurately control the power split (and hence fuel cell

power output) by limiting the DC-DC converter to transfer only a particular fraction of the vehicle power requirement and allowing the fuel cell system to automatically provide the rest. The value of the fuel cell power output that will be achieved is calculated from a high side power balance that includes the DC-DC, front motor, high voltage auxiliaries, and the fuel cell stack, as shown previously.

The benefit of this power control architecture is the de-coupling of the power split strategy, which is concerned with optimizing component and/or system efficiency, from the actual component control algorithm, which is concerned with safe and reliable actuation of the powertrain. Not only is it easier to develop these systems in parallel (they have only a single dependency, the fuel cell power set-point that is passed between them) but it also allows for rapid evaluation of various power split optimization strategies without potentially disturbing the low-level control of the powertrain itself. In more general software terms, this is analogous to writing a component “driver” (the power control algorithm) which is developed with intimate knowledge of the component being controlled, and exposing only a restricted interface to higher-level software layers that may not be aware of all the nuances of component control.

The drawback to this method is that the optimization strategy does not necessarily know what the actual power split will be. It is possible that, knowing any operational and safety constraints, the power split strategy could generate a fuel cell power set-point that, while less optimal than in the unrestricted case, would be more optimal than that which is generated by the power command algorithm. The architecture does allow for this, of course, as all sensor data is available to the power split strategy for use. The power command algorithm would still have the final determination of what can be commanded at any given time.

5.4 Torque Control Algorithm

In an AWD traction system, the torque split strategy computes torque set-points for each motor controller, with the aim to achieve various control goals such as giving traction to wheels with more grip, or optimizing the overall power conversion efficiency of the traction system. Given the time constraints of this project a fixed torque split strategy was

implemented to give basic performance in the EcoCAR competition. This strategy applies torque distribution between the two motors as a fixed percentage of the overall torque demand requested by the driver via the accelerator pedal. The split was set to 70% front and 30% rear, up to the saturation point of the front motor, with the reasoning that the more efficient front motor should deliver most of the traction demands. This gives basic AWD performance needed to compete, but it does not optimize for efficiency or provide intelligent traction control in slippery conditions.

The torque control algorithm is the main focus of this section. The torque control algorithm has several key functions that are needed, even with the most simplified torque split strategy, to ensure the performance, safety, and reliability of the traction system and the overall powertrain. These functions are:

1. Moderating torque delivery to ensure that the high voltage bus stays within component limitations;
2. Modifying the torque requests of the torque split strategy as needed to meet the instantaneous requirements / limitations of each motor system; and,
3. Mitigating the effect of motor faults by shifting all torque to the remaining motor, without affecting the pedal mapping.

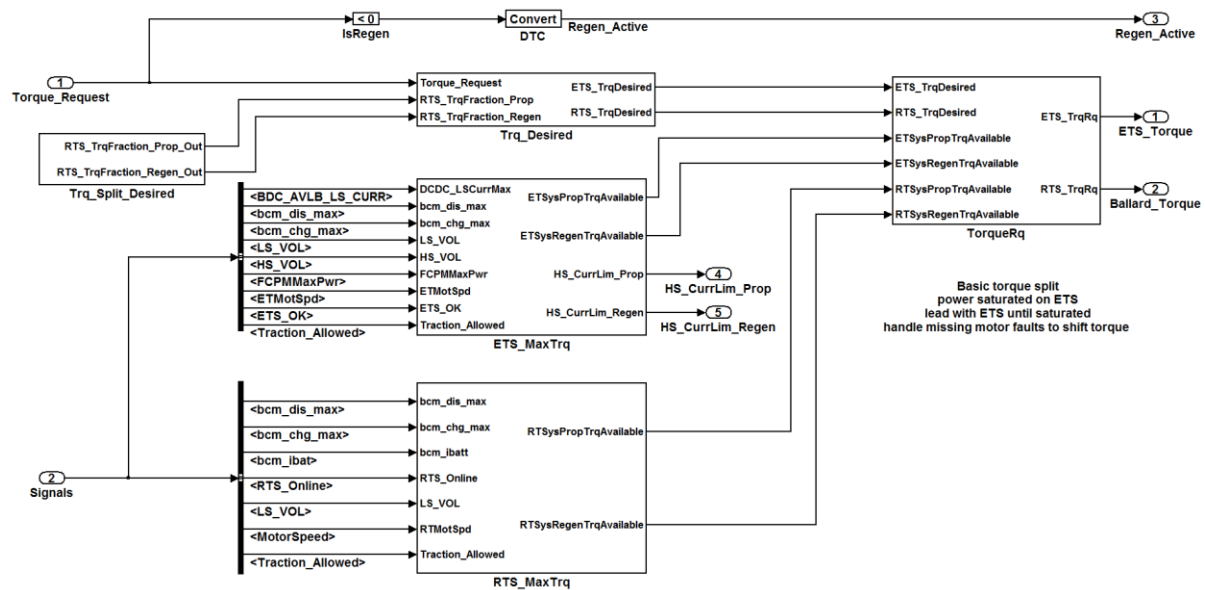


Figure 35: Torque control algorithm

The need for moderating torque delivery based on high voltage bus considerations is a consequence of the maximum electrical power output of the powertrain being lower than the maximum possible combined electrical power demand of the traction motors. Section 2.3 discusses the polarization effect of electro-chemical power sources, section 5.1.2 demonstrates how this affects in the UWAFTEcoCAR powertrain, and section 5.3 shows how this impacts the power control algorithm. Section 3.1 provides a brief overview of powertrain component limitations, including the front motor power inverter DC voltage operating range. In the UWAFTEcoCAR powertrain the maximum combined current demand of the motor system at high speeds is capable of exceeding the current supply limitations of the powertrain. Without torque delivery (i.e. electrical current request) moderation, the motor power inverters will pull more current than the electro-chemical power sources can sustain at the required voltage. Figure 36 shows a plot of DC-DC high and low side voltages and currents taken from on-road vehicle testing when this exact situation occurred. The high side bus voltage falls to the point of the low side voltage at about 3400.78 seconds, and the DC-DC converter attempts to drive 200 A (50 A above the maximum and commanded value) in an attempt to arrest the falling high side voltage. The converter is unable to sustain this condition and faults at about 3401.45 seconds.




Figure 36: Over current condition on DC bus leading to DCDC shutdown

Not only does this under-voltage condition cause loss of functionality, it can be catastrophic for the DC-DC power converter. As discussed in section 5.1.1, the DC-DC power converter is a two-quadrant device, operating only in positive boost and negative buck modes and having a peak current limitation of 150 Amps in either mode. Violating the two-quadrant constraint by allowing the high voltage bus to enter a negative boost or positive buck mode (essentially allowing the high side voltage to fall below the low side voltage) can result in component damage, and the DC-DC power converter typically shuts down to protect itself. This in turn triggers the supervisory controller to orchestrate an emergency powertrain shutdown to mitigate the potential for cascading component damage. Consuming too much current with the front motor can quickly exhaust the current driving capability of the DC-DC converter, and then shortly thereafter the current output capability of the fuel cell. However, with the DC-DC power converter having a smaller current driving capability than the battery, there is still current available for transfer from the low side of the high voltage bus to meet the front motor current requirement. The DC-DC converter has both voltage and current driving set-points (commanded by the power control algorithm as discussed in section 5.3) that it tries to meet. Due to the two-quadrant limitation and potential for converter damage

when this aspect is violated, the DC-DC converter will drive itself above the 150 A limit, pumping more current into the high side of the bus in an attempt to hold up the high side voltage. This over-current condition is only sustained briefly (100 milliseconds in Figure 36) before the converter will eventually fault and cease to operate, removing one of the supervisory controller's methods of regulating bus voltage. If current is still drawn on the front motor the high side voltage will fall as the bus capacitance discharges. Damage to the DCDC converter is possible if the powertrain is not safely shutdown before the fuel cell voltage crosses into the battery's operational range (see vehicle polarization plot in Figure 28).

To avoid the aforementioned conditions, which at the minimum cause loss of vehicle control and at worst carry the potential for powertrain component damage, the torque requests of the torque split strategy must be moderated by the torque control algorithm to respond to the instantaneous operational limitations of the powertrain. This function is separated from the torque control strategy, with a defined inter-function interface linking them that permits multiple strategies to be implemented and tested simultaneously, without the risk of circumventing the safety limitations enforced by the torque control algorithm.

In the event of a motor fault that is not cause for a system-wide shutdown, the torque control algorithm has the critical function of maintaining maximum vehicle control via the remaining motor, effectively providing one of the vehicle's "limp-home modes". The torque shifting algorithm is a sub-set of the torque control algorithm, and is the final filter for torque commands leaving the supervisory controller. The algorithm block diagram is shown in Figure 37.

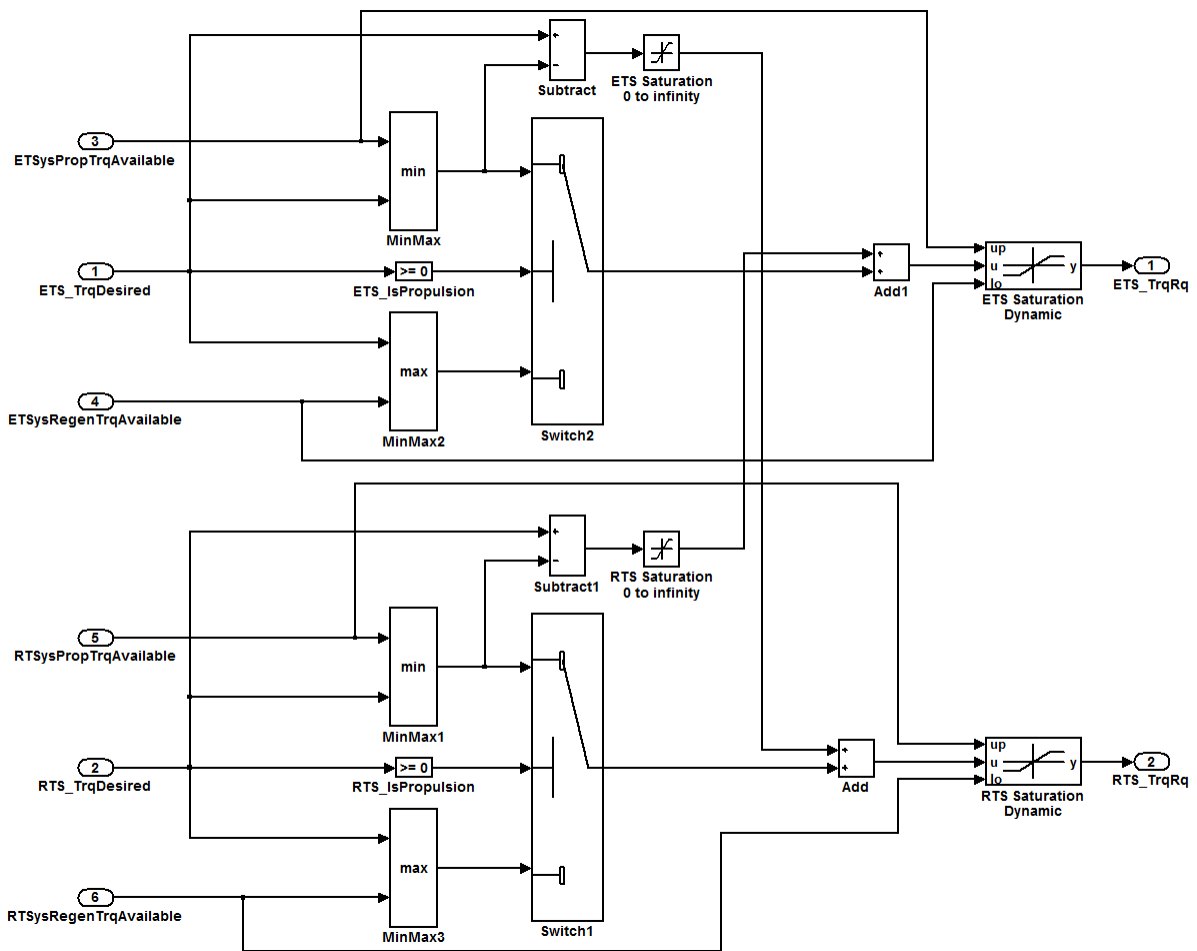


Figure 37: Torque limit and shift algorithm

The torque limit and shift algorithm receives the desired torque demand for each motor (calculated as a percentage of the overall vehicle torque demand, as determined by the torque control strategy), and the instantaneous peak torque capabilities of each motor, both positive (propulsion torque) and negative (regenerative torque). For each motor, the algorithm then compares the instantaneous torque demand with the instantaneous peak torque capability in both propulsion and regenerative operation. Given the sign of the torque demand, only one of these comparisons is valid (e.g. for a negative torque demand, only the comparison with the regenerative torque limit is valid). A switch (single condition if statement) uses the sign of the torque demand as a test to determine which comparison is valid, and the valid comparison is passed through as the smallest of either the torque demand or the applicable torque limit.

This step ensures that the outgoing torque demand does not exceed the applicable torque limit, needed to keep the high voltage bus within limits.

Simultaneously, for each motor the algorithm computes the differences between the maximum allowable propulsion allowed and the allowable propulsion torque that is passed to the switch case. This difference represents the un-met torque demand that the algorithm can attempt to shift to the other motor. When the propulsion torque demand is less than the maximum allowable propulsion torque, these are the same value and the desired shift torque is zero. When the propulsion torque demand exceeds the maximum allowable propulsion torque, the difference is computed and added to the torque request of the other motor, as output from the switch case as being the maximum torque permissible for that motor.

Finally, the combined allowable torque of each motor, plus any shift torque added from the other motor, is re-saturated based on the instantaneous torque limits for propulsion and regenerative operation. This ensures that the added shift torque does not push the torque demand for either motor above the limits imposed by previous calculations. Alternatively this could have also been accomplished by adding the shift torque to the incoming desired torque, before it is compared to the torque limits.

This torque shifting algorithm provides the functionality to dynamically limit and shift torque in a multiple-motor powertrain, maximizing power delivery while also respecting torque limitations that can be externally calculated and fed to the algorithm to satisfy any required system constraint. While this implementation of the torque control algorithm is primarily concerned with high voltage bus voltage constraints the algorithm can be used in conjunction with multiple constraints respecting thermal or other operational limitations, where the most restrictive constraint is fed to the limit and shift algorithm. These additional constraints can be added in EcoCAR2 to expand the functionality of the algorithm and ensure stable operation over a wider range of conditions.

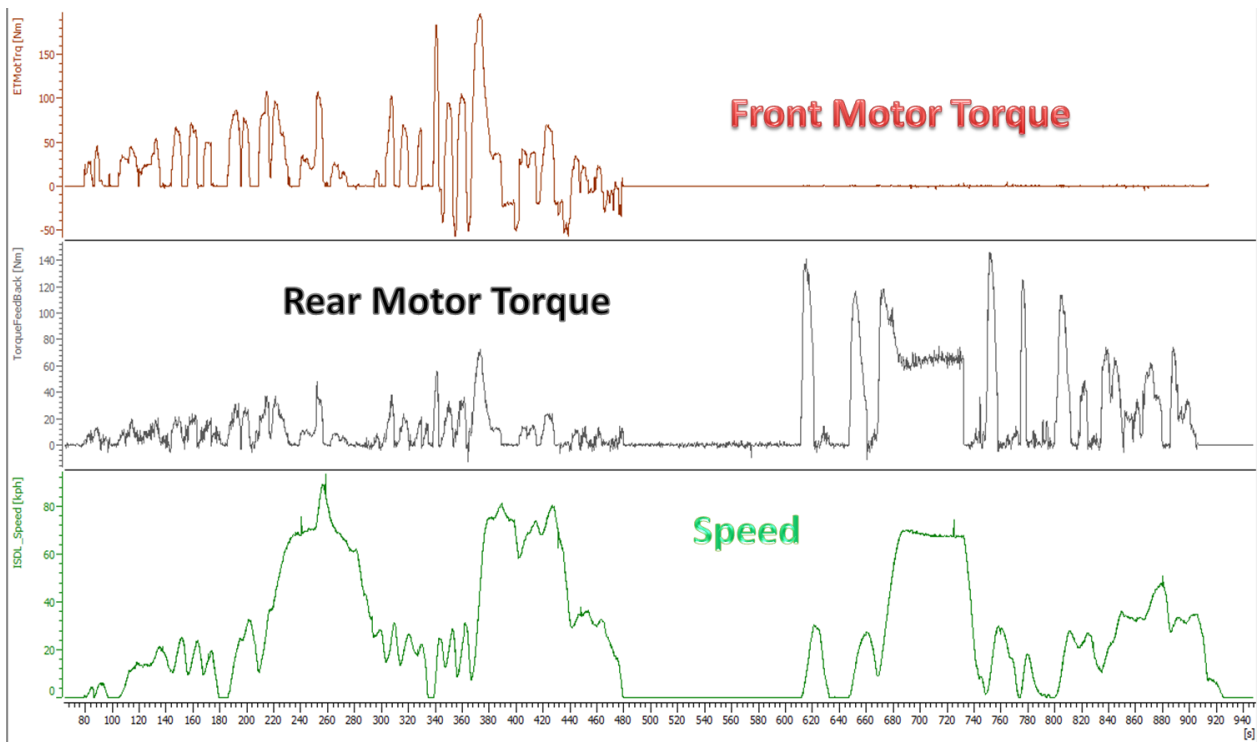


Figure 38: Torque shifting algorithm allows “limp home” on rear motor

The torque control algorithm certainly has room for improvement. It currently does not account for some of the less critical (or at least less expected) operational restrictions, such as the temperature limits of the front and rear motors. These are less critical since the individual component controllers will shut themselves down if limits such as temperature are exceeded, but it is preferable to anticipate the limit and avoid it by reducing current.

Another important improvement is to remove the fixed torque limitation and dynamically adjust current limits according to the state of the high voltage bus. This requires control code to track the voltage on each side of the bus and to predict when the voltage cross-over point will be reached. This can be accomplished via several means, including a PID controller or a state space controller.

Dynamically adjusting the current limits of the torque control algorithm’s limit and shifting capability will improve performance while maintaining function and reliability under all operating conditions. Similarly, the torque split strategy can also improve efficiency by dynamically adjusting the motor torque split. This optimization of the split strategy would

rely on the efficiency maps of each motor, combined with either online calculations or an offline lookup table.

5.5 Control System Design in Project-Based Learning Environment

Control system design in UWAFt's project-based learning environment expands on and sometimes complements the engineering design skills taught in the undergraduate curriculum. In undergraduate design courses primary consideration is given to how a design will perform in ideal or foreseeable conditions, often selected by the student. Designs are evaluated on technical accuracy and completeness within this scope. This is important but does not address the significant consideration that must be given to reliability and robustness in an engineering design. In the UWAFt environment significant consideration is given to reliability and robustness. Sometimes this consideration leads to trade-offs, sacrificing some level of performance or efficiency in order to achieve a certain level of reliability. This was evident in the torque limiting implemented in section 5.4. These considerations are inherent to project-based education and align with real-world engineering design experiences. In real-world design projects the ultimate objective is to generate a working system that will be put into service. This means that reliability and safety must be considered and, where needed, take priority over efficiency and optimization. This is especially true where resources like time, budget, and manpower are limited. The UWAFt environment certainly fits this description, and participating students gain experience making decisions to move a design forward while grappling with these sometimes conflicting priorities. This has personally led to a significantly greater ability to execute projects successfully.

Not only are resources and design priorities a significant consideration in projects, they can be in constant flux. At UWAFt that means a high rate of turnover as terms change, new or revised competition rules or performance standards as the project develops, and unforeseen obstacles that arise simply due to the complex nature of the project and the (relative) inexperience of the participants. The ability to break down a problem into components (i.e. modularize the design), divide those components amongst individuals and/or teams, and re-assemble them to produce the final design is critical. This experience of working in a larger

team, learning task and role differentiation and dividing work tasks into units, is a central theme of project-based learning and it is present in every design project at UWAF, particularly the implementation of the supervisory control code discussed in this Chapter. Designs at UWAF must also often go ahead without all the required information (like what components will be used), or must be changed at the last minute to accommodate changes in other parts of the design or in availability of components. This forces students to formulate the design problem around these barriers. This may mean producing a design process or framework (like that discussed in Chapter 4) that can adapt to changing needs, or it may mean modularizing the design with defined interfaces so that pieces can be tackled by multiple parties and substituted/changed with limited side-effects (as discussed in this Chapter).

The experiences and design work detailed in this Chapter highlight some of the key advantages to working in a project-based learning environment, using examples from the design of the powertrain control system. These include the additional considerations given to reliability and robustness in addition to performance and optimization, the need to adapt to changing circumstances and resources, and the consequent strategies that students can learn and implement to cope with that uncertainty.

Chapter 6

Safety Control System Design

Fault detection and mitigation is a critical element of automotive controls. It is particularly important in the EcoCAR competition where teams are building on-road passenger vehicles. The motivation for implementing robust safety control systems is to protect vehicle operators, those around the vehicle, and the vehicle itself from harm caused by a failure in the vehicle control system.

Appropriate safety analysis, leading to design and operational safety controls, are mandatory elements of real-world engineering design. While their importance and (sometimes) theory is touched upon in general introductory engineering courses, these skills are not often addressed nor practiced in the technical or design project courses of the undergraduate engineering curriculum. In the UWAFI project-based environment designs require safety analyses and validated fault detection and mitigation efforts in order to obtain any measure of success, as the vehicle is not permitted to run without them. Consequently, as they are mandatory elements, student motivation to research, learn, and apply safety-related design and operational controls is high, and safety becomes an embedded element of the core control system knowledge (e.g. signal acquisition and processing) being taught.

This chapter discusses the general approaches to safety analysis and control in the UWAFI EcoCAR vehicle, particularly within the electrical and control systems of the vehicle. Torque security is used as an example of how a safety-critical system is analysed to identify safety risks, how these risks are prioritized, and how priority risks are then mitigated through design and operational controls. User feedback, especially when mitigating actions may cause vehicle behaviour deviating from user intent, is also emphasized.

6.1 Frameworks for Safety Analysis

Failure Modes and Effects Analysis (FMEA) and Fault Tree Analysis (FTA) are two frameworks commonly used in the automotive industry to analyse system safety. These tools are the starting point for system safety design, and are used to identify the safety-critical

elements of the vehicle's real-time control system. FMEA and FTA are, respectively, used as bottom-up and top-down approaches to system analysis. Used together, they are complimentary processes that enable more complete coverage of a system than either process can achieve alone.

FTA is a top-down approach to system safety analysis. System faults are identified through brainstorming sessions, which form the top-level entry for a fault tree [49]. For each fault, immediate causes or events are listed. Those causes or events are further caused by other, lower-level events or combinations thereof. A logical tree of events is thus constructed, leading from the top-level fault to the most basic system events. Thus, FTA is an analysis that begins with known system-level faults and works down through causal errors to reveal root-level component failure mechanisms. FTA serves as a brainstorming tool to generate a list of critical system faults, and also provides insight into how system behaviours lead to these faults. FTA is a good tool to aid in designing controls for safety-related faults for several reasons. FTA shows single and multiple points faults, so it can help designers identify weak points (i.e. single point failures) in the system. It also shows all the intermediary steps between a root (bottom-level) cause and the (top-level) effect, which helps designers identify points of detection and mitigation within a critical sub-system.

FMEA is a bottom-up approach to system safety analysis. Systems are broken down into subsystems and components, and brainstorming sessions identify possible failure modes for each system component [50]. Once a list of failure modes is compiled for each component, the effect(s) of each failure are listed and quantified through a Severity Index (S). Next, the cause(s) of these effect(s) are listed, and quantified with an Occurrence Index (O). Design controls are identified next. In this implementation of FMEA, design controls are divided into two types. Prevention controls involve designing the system to reduce the occurrence of the causes of failure. Detection controls involve implementing methods that can be used to identify when failure causes occur. The ability of design controls to prevent or detect a failure is quantified using a Detection Index (D). The combination of Severity, Occurrence, and Detection indexes is used to rank and prioritize failure modes. FMEA is an analysis that

begins with known root-level component failure mechanisms and works up through the resulting errors to yield system-level faults.

FTA and FMEA provide two independent methods of performing a system analysis. As a deductive method, FTA is a good starting point when effects of failure are known (e.g. from previous experience with the particular application, brainstorming, hazard analysis) but the particular causes in the current system design are unknown [49]. As an inductive method, FMEA is useful when the causes of failure are known (e.g. from experience with components) but the effects on the system are unknown [49]. For both methods it is critical that the right people are involved in the analysis and that an appropriate level of knowledge is available regarding the known causes or effects that serve as a starting point for FMEA and FTA, respectively.

For example, in designing a vehicle control system unintended vehicle movement is a critical effect of a failure in part of the system. During the design process a FTA for unintended vehicle movement would be performed in an effort to understand how certain events within the components of the vehicle control system can lead to this effect. Thus, FTA is an analytical process that is best executed by experts at the system integration or application level. A good FTA session may be preceded a brainstorming session at the system or application level to generate the failure effects that serve as the top-level of the fault trees.

Conversely, FMEA is useful when detailed knowledge of a component and its failure mechanisms are known. For example, in the same vehicle design project, a FMEA performed on a motor controller would identify an open or short on the torque control signal as possible causes of failure. This failure may lead to the effect of unintended vehicle movement. Thus, FMEA is a brainstorming process that lends itself to experts at the component or subsystem level. A good starting point for FMEA is a brainstorming session at the component level, to generate the failure modes that serve as bottom-level events for the analysis.

By approaching system analysis from both ends, top-down and bottom-up, and by involving both system-level and component-level experts in the brainstorming process, the likelihood of identifying and linking all safety-critical causes and events is increased. These methods

are only the beginning of designing a safe control system, however. Once safety-related design elements are identified, they must be addressed with appropriate controls. Those controls may occur in the design phase to reduce the probability of failure causes, and / or they may be implemented within the production system to detect known failure effects and take appropriate action. The focus of this chapter is on the latter through examples related to one of the most safety-critical elements of vehicle controls: torque delivery.

6.2 Fault Identification in the Torque Control System

Vehicle torque security is accomplished by identifying those critical systems and signals which, if failed, could cause a loss of vehicle control or damage to one or more powertrain components. Those systems and signals are then monitored by software that is capable of detecting failure conditions and taking mitigating action to protect the vehicle and its occupants.

The first step in this process, identifying critical systems and signals, is both the most difficult and the most critical. It is critical because, simply, in order to monitor and protect a safety-critical signal that signal must be known to need protection. It is difficult because identifying those systems and signals that are safety-critical is not easily done using kinds of the formal analytical frameworks that engineers, particularly new engineers, are so reliant on. The FTA and FMEA processes are applied to the torque control system in this section to reveal their respective strengths and weaknesses in performing systems analysis in the UWAFT EcoCAR project.

The torque control system involves all of the components and interfaces that are responsible for translating driver inputs into vehicle motion. To facilitate the top-down FTA for unintended torque actuation, it is useful to generate a system-level schematic of the torque command path, as shown in Figure 39.

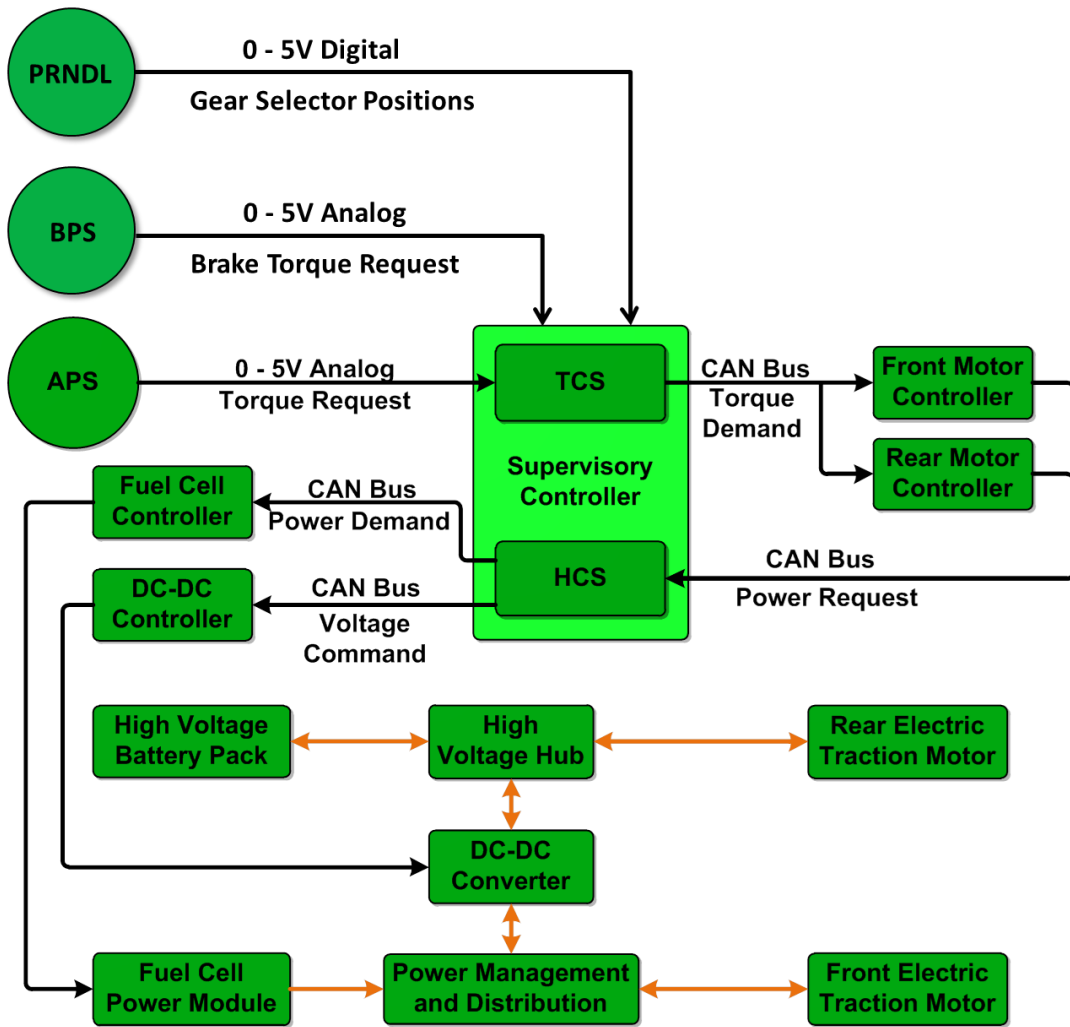


Figure 39: Control path for torque actuation in the EcoCAR vehicle

The control path for torque actuation in the EcoCAR vehicle starts with the driver's input via the accelerator pedal position sensor (APS). This input is sensed by a linear potentiometer, translating the position of the pedal into an analog voltage signal. This voltage is sensed by the supervisor controller, and is converted into a torque request using a pedal map. The pedal map is simply a look-up table that maps pedal positions to torque requests. This map can be tuned for certain objectives. For example, a map that places most of the available torque at low pedal positions results in a vehicle that feels responsive, while distributing torque more evenly throughout the map or towards the end of pedal travel will result in a dull response. The distribution of torque throughout the pedal map is often one of the things that an ECU

modifies when switching between driving modes like “Eco” or “Sport”, where “Eco” mode may limit the torque requests or require a more aggressive pedal position to achieve full torque, with the aim to improve energy efficiency during normal driving.

The desired torque indicated by the pedal map is split by the torque control strategy between the front and rear motors. The torque control algorithm then adjusts these torque requests as required to meet operational and safety limitations, resulting in torque commands for each motor. These torque commands are then sent to each motor controller via CAN bus messages. The motor controllers drive the electrical current required to meet the requested torque level (within the operating limits of the controller), and provide feedback indicating the amount of electrical power required to meet the torque request and the actual torque achieved. The hybrid control strategy uses the motor power requirements to determine the overall vehicle power requirement, splits this requirement between the fuel cell and battery pack, and generates a fuel cell power request. Like the torque control algorithm, the power control algorithm modifies the fuel cell power request to meet operational and safety requirements, and generates a fuel cell power command. Component control algorithms within the supervisory controller then generate CAN bus control signals. The DC-DC receives a voltage set-point and a current limit, and the fuel cell receives a power request. Current is sourced from the battery, fuel cell, and DC-DC converter to achieve the requested torque. The fuel cell controller uses the fuel cell power request to ensure adequate delivery of reactants to the anode and cathode via the hydrogen fuel injectors and air compressor, respectively. The DC-DC converter drives the indicated current in the direction required by the voltage set-point.

6.2.1 Top-Down Safety Analysis: FTA

The fault tree for unintended vehicle movement can be constructed by working backwards from the torque output of the motor controllers, through the various components and interfaces, to the root failure mechanisms. A FTA for unintended movement of the UWAFT EcoCAR is shown in Figure 40. The FTA for the supervisory controller’s is selectively expanded to the root causes associated with the accelerator pedal sensor (APS) in Figure 41.

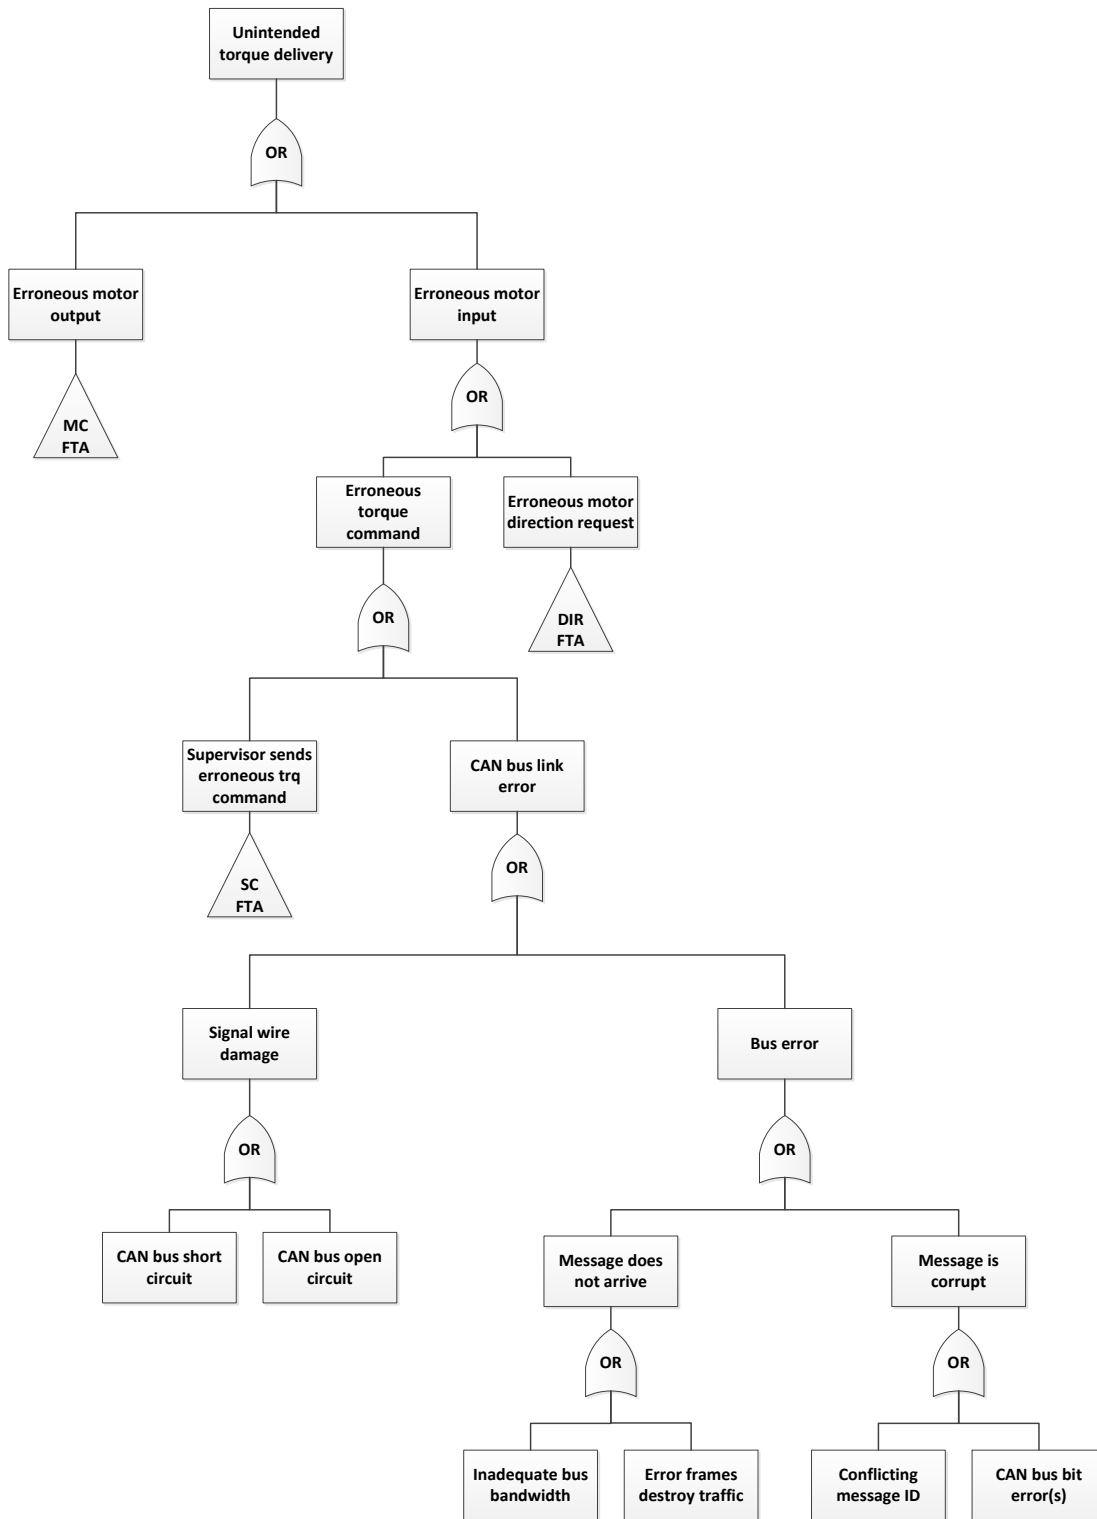


Figure 40: Fault tree for unintended torque delivery



Figure 41: Fault tree for supervisory controller erroneous torque request

The fault tree follows the pattern of examining the direct cause of each event and moving backwards through the system. Starting from the motor controller that produces the torque, faults can arise in the controller itself (via hardware or software) or from the controller inputs. Focusing on the inputs, the motor torque request and motor direction request signals determine the magnitude and direction of the actuated torque. These signals are sent via the CAN bus by the supervisory controller on two separate messages. Focusing on the motor torque request signal, errors can occur due to either an erroneous value sent by the supervisory controller or an error on the physical CAN bus. These faults continue to be traced to the limit of the tree's resolution. For this fault tree, this limit is defined as that point which exposes a cause which can be detected within the vehicle. For example, a conflicting message ID could be caused by any one of the other CAN-connected controllers, but this is irrelevant to the detection and mitigation of the fault within a production system (although furthering the tree to examine the root cause of a conflicting message ID is useful during the design stage).

The framework for executing the FTA considers ECUs and their interfaces. ECU failures are categorized by input, processing (hardware and software), and output faults. Interface failures are examined within a framework appropriate for the type of interface. For example, an analog voltage interface (e.g. APS) can fail due to physical damage (e.g. wire damage), signal degradation (e.g. resistance along the path), or signal noise (e.g. external electromagnetic interference). A linear transducer (APS) can produce an unintended signal due to physical or electrical damage. By applying these structures to each element of the tree it is easier and quicker to obtain complete coverage of possible faults. It is still critical that systems-level experts participate in the FTA to ensure that each component's role in the tree is explored.

6.2.2 Bottom-Up Safety Analysis: FMEA

A bottom up FMEA of each component involved in the torque control system is complex and beyond the scope of this work. An analysis of the supervisory controller serves as an example of how FMEA is applied to elements of the torque control system. For an FMEA analysis the

focus is on the various failure modes of the component. It has been noted that brainstorming with technical experts serves to generate the required failure modes, however as with the FTA it is important to place some structure around the discussion to ensure a complete coverage [51]. For an electronic control unit like the supervisory controller, the FMEA can be structured to examine the controller from the perspective of controller inputs, controller hardware, controller logic, and controller outputs. Hardware-related failure modes (inputs, outputs, and controller hardware) of various types (e.g. digital, analog, CAN) experience failure modes common to their type, and designers can draw on common approaches to identify, detect, and protect against their failure. The software that produces the control logic must be analysed in the particular context of the application, starting with the software outputs and working backwards through possible causes of output failure (essentially its own FTA).

In order to demonstrate the FMEA, an analysis of the accelerator pedal is performed. The resulting table of failure modes is given in Appendix G. Failures are categorized by type (e.g. interface, sensor) and sub-type (e.g. physical, electrical) to assist with failure identification. Many of the failure modes identified in the FTA are covered in the FMEA, however in the FMEA they are documented with details on causes, effects, and detection methods. This additional detail is valuable for improving safety by addressing failure causes or improving detection methods. It is also a better form of design documentation.

The FMEA, having a bottom-up perspective, also reveals additional possible effects that would not be identified in the top-down FTA (where effects must be known in advance). One of the effects of a short to power within the accelerator pedal system is a short circuit of the power supply and consequent brown-out. This brown-out, while not a root failure, can affect other sensors and/or systems supplied by the same hardware, for example the second accelerator pedal sensor in a redundant assembly. One outcome of this analysis would be the recommendation of dual redundancy of the accelerator pedal assembly that extends to include separate power supplies, both from each other and from other vehicle systems. This prevents the loss of a single power supply from causing the complete failure of the safety-critical accelerator pedal assembly.

6.2.3 Top-Down vs. Bottom-Up

This work demonstrates how FTA and FMEA are used together to produce a complete safety analysis of a control system. FTA is a useful engineering tool for tracking steps in a system-wide process and where failures may occur. It shows the interdependence of components and interfaces, and events probabilities propagated along the tree permits a quantitative understanding of the probability of a top-level fault. This system-level information helps in the design and documentation of vehicle sub-systems. FMEA is a useful documentation tool for tracking details on individual failure mechanisms within a system, such as their severity, occurrence rate, detectability, and action items. It helps in the documentation of the overall safety of a system and serves as a repository of safety-related knowledge for components and systems. FMEA focuses on causes and action items for failure modes, where FTA simply demonstrates where failures may occur.

Both the FTA and FMEA are conducted with a common framework to facilitate the brainstorming process and produce reliable coverage of potential failures from each method's point of view. System-wide fault coverage is additionally ensured by performing the analysis from both top-down and bottom-up perspectives. While the methods certainly overlap, the differing perspectives of FMEA and FTA uncover, respectively, failures and faults that can be invisible to the other method. FMEA uses known failures (causes) to uncover unknown faults (effects), while FTA uses known faults (effects) to uncover unknown failures (causes).

6.3 Fault Detection and Mitigation: Implementation and Validation Using HIL

The FMEA and FTA safety analyses of the accelerator pedal yield important failure modes that cannot be completely designed out of the system. Given the severity of a failure in the torque command path, these failures must be detectable during operation, and actions must be taken to mitigate the failure. The FMEA analysis in Appendix G shows that the effects of various accelerator pedal malfunctions can be summarized as zero signal, saturated signal, and unpredictable erroneous signal. These effects are detectable in software by various

means. Again, as with the fault identification of the safety analysis, a structured approach to fault detection helps ensure complete coverage.

The zero and saturated signals are predictable values. The malfunction of a single sensor resulting in these failure effects can be detected by classifying these values as invalid. Typically, some portion of the total electrical output of the sensor, both at the top and bottom of the sensor range, is excluded from normal sensor output. For example, in a 5 V linear transducer like the accelerator pedal, the “zero” mechanical position of the sensor will correspond to a 0.5 V output, and the “full” mechanical position will correspond to a 4.5 V output. An electrical malfunction resulting in a 0 V or 5 V signal is therefore easily detected by the supervisory control software. This is an example of defining valid signal characteristics that can be evaluated using only the sensor signal itself.

Other failure modes can occur that will result in an unpredictable but still erroneous signal, which is more difficult to detect. In this case the sensor value itself is insufficient, however other sensor characteristics can be used. For example, the rate of change of the accelerator pedal signal is limited by the mechanical rate at which the pedal can be moved. Changes in the signal value outside of this limit, either immediately (instantaneous jump) or over time (signal noise) can be detected by the controller. This solution must be carefully calibrated to ensure that false failures are not detected, and is limited in that it cannot detect sensor movement within these tolerances (e.g. a slower sensor drift).

It is also possible to compare the signal to another (redundant) sensor value. If the deviation between the two sensors exceeds a particular threshold the supervisory controller will be able to determine that there is a failure in one of the sensors. This method provides protection for the sensor value across all signal characteristics, however the controller will not be able to determine which sensor value has failed unless one of the previously discussed signal characteristics also exceeds detection thresholds.

The validity of the signal can also be determined by correlating it with the values of other control signals, which can be done in various ways. This amounts to a redundancy check of the sensor value by using related sensory information. For example, if a motor controller

reports a significant current draw when the accelerator pedal is at rest, then there is likely an error in the control system. This error may lie in the accelerator pedal, the motor controller, or the communications system. Using a model of the expected torque delivery as a function of the accelerator pedal position (i.e. the pedal map) provides another layer of signal integrity checking. Again, this type of modelling must be carefully calibrated to avoid flagging false failures.

Some of these detection mechanisms are implemented for the EcoCAR vehicle's accelerator pedal input, as shown in the Simulink sensor block of Figure 42. Dual redundant inputs are used to enhance fault detection. A redundancy check (shown in Figure 43) compares the values of each sensor to detect invalid but in-range signals. Out of range checks (shown in Figure 44) detect shorts to power and ground.

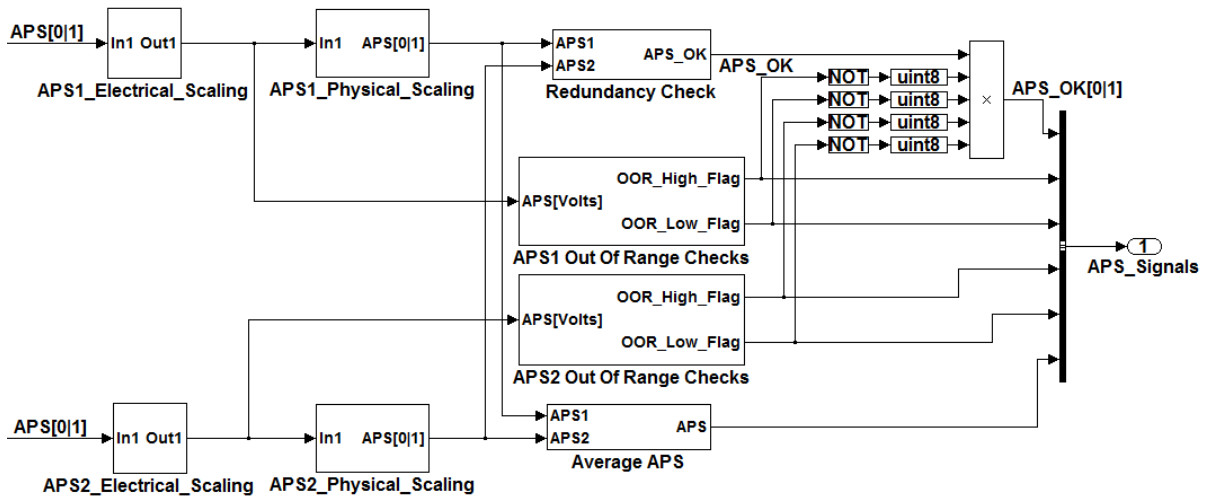


Figure 42: Accelerator pedal sensor block for EcoCAR vehicle




Figure 43: Accelerator pedal out of range checks




Figure 44: Accelerator pedal redundancy check

When a fault is detected, appropriate mitigating actions are needed that are commensurate with the severity of the fault. In the case of protecting the accelerator pedal signal, the consequences of unintended torque actuation are high, necessitating a severe reaction such as limiting or shutting down the traction system. The out of range and redundancy checks combine to form the APS_OK signal, which is used as an input to the Traction_Allowed gate (Figure 45). This gate checks the value of several signals to ensure that it is safe to enable the torque control system. If the APS_OK signal is false, Traction_Allowed also becomes false and the torque output is zeroed via the torque control algorithm described in section 5.4.

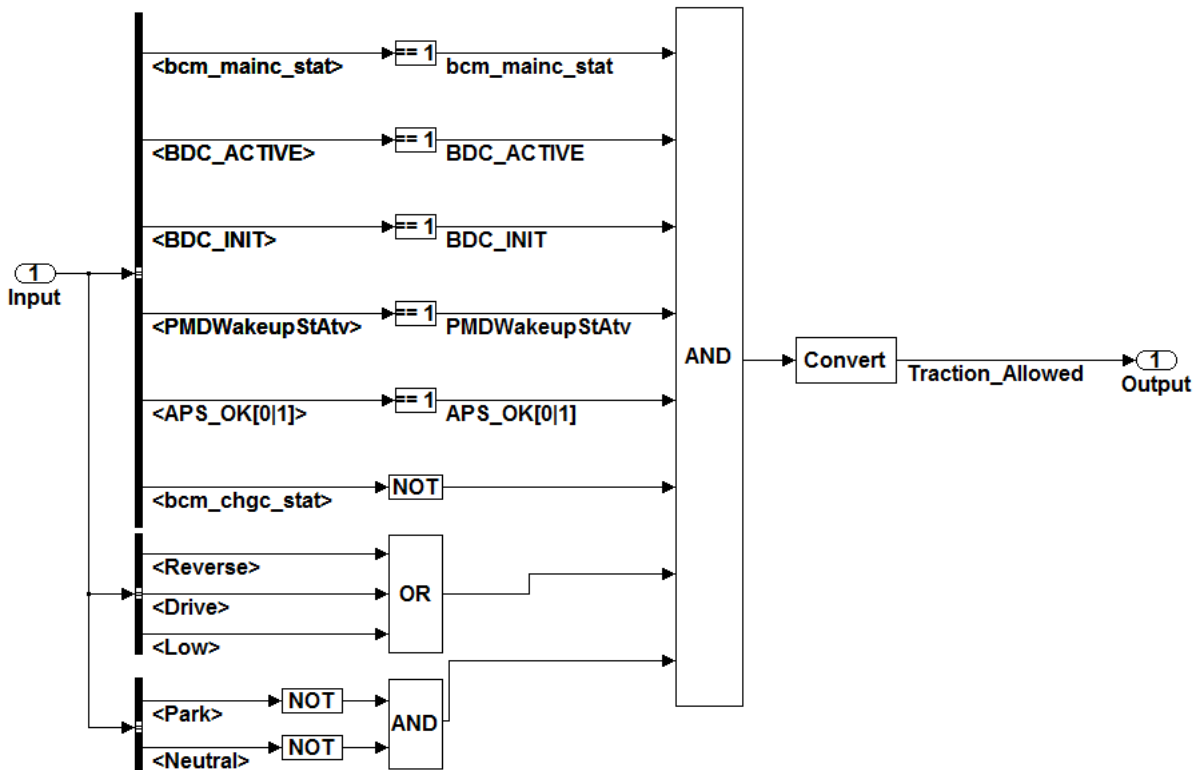


Figure 45: Traction_Allowed gate using APS_OK signal

6.4 Validation Using Hardware-in-the-Loop Testing

Hardware-in-the-loop testing enables the design team to validate not only the control algorithms but also their implementation on the controller hardware, before the vehicle is available for testing. This enables more parallel development and decreases the time required to produce a verified control system. Hardware-in-the-loop testing involves connecting the production controller, running production control code, to a simulator. The simulator runs a model of the plant (in this case the UWAFTEcoCAR vehicle) in real-time and replicates the I/O that the controller would see in the actual application. The hardware-in-the-loop (HIL) setup used to validate the control code for the EcoCAR vehicle is shown in Figure 46, and the PC interface is shown in Figure 47.




Figure 46: Simulator hardware connected to the supervisory controller and PC interface



Figure 47: PC interface for the EcoCAR HIL setup

The general setup for the HIL simulator, including a plant model, some GUI elements, and some I/O, was already available in year 2 of the competition. It has been extended to include additional GUI elements and additional I/O handling (for example dual pedal sensors, PRNDL signals, more CAN messages) to replicate the parts of the vehicle that are involved in the torque command path. This allows the simulator to fully interact with the torque control pathways of the supervisory controller, and enables torque-related faults to be tested.

The HIL simulator is useful for validating fault detection and mitigation code on production hardware before the vehicle is available for testing. To extend the case study of protecting the torque control system’s accelerator pedal signal, several faults are simulated in the HIL setup. The controller detects the faults and provides feedback within the simulator and on the vehicle’s user feedback LCD display, which is wired into the simulator. The HIL simulator is flexible, allowing for parts of the setup to be composed of physical hardware (e.g. supervisory controller, user feedback LCD, CAN bus) and other parts to be simulated (e.g.

accelerator and brake pedals, vehicle powertrain components). The PC interface to the simulator allows real-time interaction with the plant model. In this way the accelerator and brake pedal signals are input via the PC interface. They are then converted to an analog voltage by the HIL simulator and sent to the supervisory controller. The supervisory controller thus receives the same input as if it is connected to the actual vehicle. A driver model also runs as part of the plant simulation, which allows the simulator to follow a pre-determined drive cycle.

Over the course of a simulated drive cycle various failures are invoked both in hardware and via the PC interface. These include failures of the accelerator pedal control signal (short to ground, short to power, mismatch between redundant sensor values), the CAN bus (break of the CAN network wires, spurious torque request) and the motor controller (loss of torque). The result of shorting one of the accelerator pedal analog signals to supply power is shown in Figure 48. Accelerator pedal signal 2 is shorted to 5V, which is detected by the out of range checks. Torque is disabled and the user is given feedback on the LCD display.




Figure 48: Simulation of a short to power for accelerator pedal signal 2

The simulator is also used to verify the “brake pedal wins” logic that disables all torque actuation (ignores the accelerator pedal input) whenever the brake pedal is applied. One key

element learned from experience helps to ensure that this safety feature does not interfere with the driver's intentions. The threshold for the brake pedal wins logic is set above the lower threshold for brake torque actuation. That is, the driver can depress the brake pedal slightly without causing the brake pedal wins logic to enable. This accomplishes two goals. First, it avoids the potential startle that will occur if the driver accidentally disturbs the brake pedal while driving. Second, it allows for tolerance in the default position of the brake pedal. Over time the brake pedal's default position may shift due to electrical or physical drift. On another vehicle it has been observed that setting the brake pedal wins threshold too low caused the accelerator pedal to become disabled without any driver input, as the pedal return mechanism had deteriorated over time. This is a significant safety issue, especially as the driver was on an incline at the time. Additionally, without feedback the driver had no ability to correct the issue. In this case simply reporting the error, either with a simple indicator lamp or via text in the driver feedback display, would have allowed the driver to adjust the brake pedal and regain control. Figure 49 shows the results of a brake pedal input that exceeds the 15% application threshold set for the brake pedal wins logic. ETS and RTS torque are cut to zero even though the driver is applying 24% to the accelerator pedal. Additionally, the LCD informs the driver that the accelerator input is being ignored due to the brake pedal application.




Figure 49: Results of the brake pedal wins logic executing on the simulator

6.5 When the Analysis Fails: Mishandling the ETS Torque Enable Signal

Even with software- and hardware-in-the-loop testing, some errors can slip through undetected and become noticeable only when the right circumstances occur in the vehicle. As previously mentioned, fault identification methods like FMEA and FTA are only as good as the system- and component-level knowledge that drives them. An example occurred during a road test of the EcoCAR with only a week remaining until the final competition of year 3. The vehicle experienced a total loss of control during a full acceleration test. The vehicle could not be restarted and had to be towed back to the garage. A preliminary review of the CAN log data from the event showed that the ETS experienced a fault at the time of the failure. Further attempts to start the ETS were unsuccessful, and it was determined that the entire inverter had likely failed and needed to be replaced. It would be impossible to ship a new motor and inverter in time for replacement before the competition, so the team's only choice was to attempt a replacement at the competition, before being able to enter events. However, GM was unable to provide a new ETS if the cause of the failure was unknown, as

it would likely fail again. The author undertook an extensive review of the vehicle systems and the event CAN log data in order to pinpoint the cause of the failure.

The first indication of trouble was that the fuse linking the DC-DC converter to the fuel cell system had been opened. This pointed to an overcurrent condition through the DC-DC converter. The CAN log data showed, among other faults, an active flag for the fault code ETDiagHVOvrVltFlt, which is an over-voltage condition in the ETS. This led to a plot of the voltage measurement at the ETS in the time leading up to the failure event. This plot is reproduced in Figure 50 below.




Figure 50: Powertrain voltage plots leading to ETS failure event

During hard acceleration events the fuel cell delivers maximum power and the voltage sages due to polarization. As shown in Figure 28, at high enough fuel cell output power levels the voltages of the two buses can near the cross-over point. This was known to be a critical failure mechanism, and so detection for this condition was added to the diagnostic code

block. A mitigation action was required once the high to low side bus voltage ratio was detected to have reached 5% of unity, or 1.05. One of the mitigating actions for this condition was to disable the ETS “torque enable” signal. This signal is a physical control wire leading from the inverter to the supervisor. It is driven low by the supervisor when torque is allowed, and floats high to disable torque. This mitigating action was placed in the code as a back-stop measure to prevent potential damage the DC-DC converter. The primary means of protecting the DC-DC was to never reach this condition in the first place through calibration of the power control algorithm. Unfortunately, there were two major oversights when this code was added. First, the fault condition was not latched. That is, after disabling torque the high side voltage would rise, the fault condition would clear, the ETS torque would be re-enabled, the high side voltage would again fall (if the power demand was still sufficiently high), and the cycle would repeat. This caused a “bucking” sensation in the vehicle prior to failure. The second issue was with the control signal itself. The ETS torque enable signal is not simply a controller-to-controller request; it physically cuts off the switching capability of the power inverter and effectively opens the DC high voltage circuit within the inverter. This causes a build-up of voltage within the inverter. The effect is more clearly shown in Figure 51, where the time scale is reduced. When the green line (the ratio of high to low side voltage, R) reaches the critical threshold (purple line) of 1.05, the ETS torque enable signal (red line) goes to zero. The ETS inverter voltage simultaneously spikes due to the large inductance in the circuit. Figure 51 shows an acceleration event just prior to the actual failure. In this event the voltage spikes were not enough to cause component damage (although the driver recalls making a note to investigate the vehicle’s bucking motion after the test).

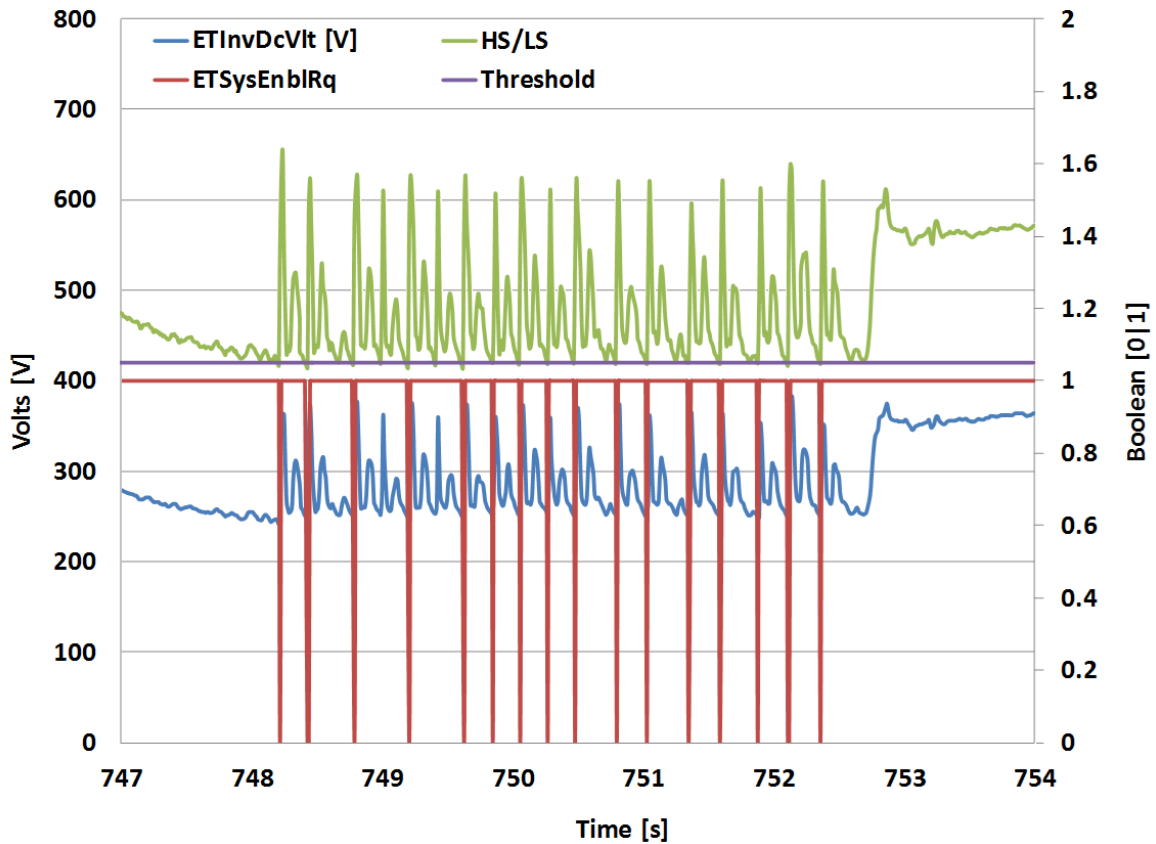


Figure 51: Effect of floating the ETS torque enable signal on bus high side voltage

When the failure event occurs, the ETS torque enable line floats, and the voltage recorded by high-side components exceeds 600 V; the actual voltage reached inside the inverter is unknown as the event likely progressed faster than the CAN logging equipment could track. The internal damage led to a short-circuit condition inside the inverter. The ETS low voltage connector was inspected after the event and had melted, with the ground pin having burned and collapsed, indicating significant current had used this path to escape the inverter. This caused the DC-DC fuse and fuses inside the PMD to open, saving the vehicle from further damage. The supervisory control code detected the unexpected loss of several powertrain components and executed a “quick stop” of the powertrain. This was indicated to the driver (the author) via the LCD feedback display, and the vehicle was pulled to the side of the road. There were no outward signs of damage, but the supervisory controller could not re-start the

vehicle as the DC-DC circuit, essential to all modes of operation (EV and Hybrid), had been broken, meaning that the PMD could not be started.

With the failure mechanism identified, the team was able to obtain a new ETS and replace it at the final competition, with enough time remaining to complete all competition events. The vehicle did not experience another fault after this event. The main cause of this event was a simple but fundamental misunderstanding of the function of the “ETS torque enable” signal. The team was aware of the potential consequences of opening a high current contact on an inductive device like the inverter, but was unaware that this was the function of the “torque enable” signal. This issue was resolved with better calibration of the power control algorithm to prevent the cross-over condition from being reached, by leaving the ETS torque enable signal “enabled” during all but the most critical events (treating it as an emergency stop), and by latching the disabled “torque enable” signal so that it would not fluctuate as it did in Figure 50.

This event highlights two things relevant to safety in the vehicle control system. First, it reinforces the need to carefully separate the component control functions of the supervisory software from the more casually modified code (such as control strategies, vehicle state machines, etc.) that can change rapidly during development. The component control blocks are the last line of defense in ensuring safe operation of the powertrain and they must be controlled, documented, and tested by knowledgeable individuals to ensure reliability and safety. Second, an FMEA of the ETS could have revealed this failure mode if the effects of switching the torque enable signal were known. The author, while possessing the most knowledge of the ETS on the team at the time, did not have the required understanding to anticipate this failure. This highlights the critical need to have the right people with the right knowledge involved in FTA or FMEA analyses. In this case the switching was caused by an error in the supervisory control code, which could have been corrected. More generally it is possible that the signal wire for the torque enable line could become damaged or experience electrical noise that yields similar behaviour. Given the serious effects of this failure mode, additional safeties to tolerate this type of signal failure could be built into the ETS while still providing the supervisory controller with an emergency disable control signal for the ETS.

6.6 Design for Safety in a Project-Based Learning Environment

The UWAFt project-based learning environment incorporates safety as an integral element of all design tasks, including the vehicle control system. This is a critical part of engineering practice that is not otherwise found in undergraduate education, at least in the author's experience. Control system design at UWAFt includes FMEA and FTA analyses to map control system processes, identify failure modes and failure effects, qualify their criticality, and design detection and mitigation strategies to increase safety. These safety analyses and their outcomes (e.g. signal protection schemes) are embedded into the technical control system concepts that students learn at UWAFt (e.g. signal acquisition and processing) as mandatory elements as opposed to optional add-ons. Safety thus becomes an integral part of the subject matter and has a greater chance of being subsequently applied to future control system projects that students might encounter. The knowledge of how to perform safety analyses and design robust, fault-tolerant control systems has been critical to the author's ability to successfully execute projects after leaving UWAFt.

System validation is another core skill that follows naturally from learning in a project-based environment. In order to achieve a successful outcome the project must be functional, and that means it must be tested, in whole and in parts. The safety requirements of UWAFt vehicle projects demand that safety-critical subsystems and behaviours be tested prior to putting the vehicle into service. Again, validation is embedded into the project, closing the loop on the design process by verifying that the design-intent is being achieved. Validation is taught in the undergraduate engineering curriculum with the aim to verify the performance of a design, but not to verify its safety.

The success of UWAFt's project-based learning environment in teaching safety design skills is two-fold. First and more generally, student motivation to acquire new skills is typically high when those skills are being applied to real, implementable designs. Second and more specifically, by embedding safety analysis and validation processes into the control system "curriculum" safety becomes an integral part of that technical knowledge and gives students a significant advantage in future designs.

Chapter 7

On-Road Vehicle Testing

All of the design, simulation, testing, and construction culminated in on-road evaluations of the EcoCAR's capabilities. These tests are designed to evaluate the vehicle's dynamic performance (acceleration, braking, and cornering), energy efficiency, emissions (both generated by the vehicle and incurred by the processing of the fuel(s) upstream), and consumer acceptability (noise, vibration, harshness, consumer features). The acceleration, energy efficiency, and emissions performance are a direct consequence of the powertrain design and execution, and are reviewed here. The specifications evaluated by the on-road testing are outlined in Table 9, with values for the original vehicle, the targets for the re-engineered vehicle, and the results of on-road evaluation. The assumed energy density, greenhouse gas, and petroleum energy use assumptions of the various competition fuels are shown in Table 10. Details from this discussion have been previously published in [52].

Table 9: Vehicle performance metrics evaluated during on-road testing

METRIC		REQUIREMENT		
EcoCAR	Production VUE	Competition	Team VTS	Year 3 Result
Accel 0-60 (s)	10.6 s	≤ 14 s	10 s	9.30 s
Accel 50-70 (s)	7.2 s	≤ 10 s	7 s	8.43 s
Towing Capacity (kg, lb)	680 kg (1,500 lb)	≥680 kg @ 3.5% 20 min @ 72 km·h ⁻¹ (45 mi·h ⁻¹)	680 kg (1,500 lb)	PASS
Cargo Capacity (m ³ , ft ³)	0.83 m ³	Height: 457 mm (18 in) Depth: 686 mm (27 in) Width: 762 mm (30 in)	Height: 457 mm (18 in) Depth: 686 mm (27 in) Width: 762 mm (30 in)	PASS
Passenger Capacity	5	≥ 4	2	2
Braking 60 – 0 (m, ft)	38 m – 43 m (123 – 140 ft)	< 51.8 m (170 ft)	45 m (147.6 ft)	49.5 m (162.3 ft)
Mass (kg, lb)	1,758 kg (3,875 lb)	≤ 2,268 kg (5,000 lb)	2,318 kg (5,111 lb)	2,258 kg (4,979 lb)
Starting Time (s)	≤ 2 s	≤ 15 s	≤ 15 s	< 15 s
Ground Clearance (mm, in)	198 mm (7.8 in)	≥ 178 mm (7.0 in)	165.1 mm (6.5 in)	158.8 mm (6.25 in)

Range (km, mi)	> 580 km (360 mi)	≥ 320 km (200 mi)	350 km (220 mi)	273 km (170 mi)
Fuel Consumption, CAFE Unadjusted, Combined, Team: U.F. Weighted l/100 km	8.3 L per 100km (28.3 mpg _{ge})	7.4 L per 100km (32 mpg _{ge})	3.2 L per 100km (73.5 mpg _{ge})	5.5 L per 100km (42.8 mpg _{ge})
Charge Depleting Fuel Consumption (l/km)	N/A	N/A	0 l/100km (0 mpg _{ge})	N/A
Charge Sustaining Fuel Consumption (l/km)	N/A	N/A	4.9 l/100km (47.7 mpg _{ge})	N/A
Charge Depleting Range (km)	N/A	N/A	50 km (31.3 mi)	21.7 km (13.5 mi)
Petroleum Use (kWh/km)	0.85 kWh·km ⁻¹	0.77 kWh·km ⁻¹	0.013 kWh·km ⁻¹	0.013 kWh·km ⁻¹
Emissions	Tier II Bin 5	Tier II Bin 5	N/A	Tier II Bin 1
WTW GHG Emissions (g/km)	250 g·km ⁻¹	224 g·km ⁻¹	192 g·km ⁻¹	243 g·km ⁻¹

Table 10: Competition fuel data

FUEL	ENERGY CONTENT (kWh·kg ⁻¹)	GHG (gCO _{2eq} ·kWh ⁻¹)	PEU (kWh _{petroleum} ·kWh _{fuel} ⁻¹)
Hydrogen	33.3	397.49	0.0147
Electricity	-	699.18	0.0785
E10	11.61	63.33	0.0932
E85	8.17	1.57	0.0832
B20	11.59	1.99	0.0642

7.1 Acceleration

Vehicle acceleration from a standstill to 96 km·h⁻¹ was measured to be 9.30 seconds, while acceleration from 80 to 112 km·h⁻¹ was accomplished in 8.43 seconds. The 80-112 km·h⁻¹ run took longer than the team expected. These values could be reduced if the torque output limitations discussed in section 5.4 are optimized to dynamic adjust according to bus voltage conditions, as opposed to a fixed reduction.

The acceleration event shows the benefit of re-arranging the motor configuration to permit full power delivery despite the restrictions of the Brusa DC-DC converter. Figure 52 shows plots of vehicle speed, battery and fuel cell power, and front and rear motor torque during one of the 0 to 60 mph acceleration runs. The fuel cell outputs up to 95 kW of power and the front motor achieves up to 350 Nm of torque. The battery is able to provide up to 60 kW of power to support traction, which would not be possible if the DC-DC was located between the battery and the rear motor. The DC-DC plays another critical role in the acceleration event by automatically compensating for fuel cell inertia. When the driver first applies the accelerator pedal a full torque request is sent to both electric motors. For full acceleration the motor speed (and electrical power required) increases faster than the fuel cell can increase its power output. The DC-DC converter maintains the correct high side voltage by transferring current from the battery to the front motor during periods where the fuel cell output is insufficient, thus improving motor response. Similarly, when the driver reaches the desired speed and releases the accelerator the DC-DC converter automatically transfers current from the fuel cell system to the battery, compensating again for fuel cell system electrical inertia and preventing over-voltage of the high side. This effect is evident at 667 seconds in Figure 52, where the torque request to the motors is reduced (but not negative, hence no regenerative braking) and a significant power transfer of roughly 30 kW is required to unload the high side of the bus.

However care must still be taken to ensure that the instantaneous front motor electrical power requirement does not exceed the instantaneous maximum power available from both the fuel cell and the DC-DC converter. If this occurs the DC-DC converter, in an attempt to maintain high side voltage, will continue to increase the current transferred from the battery beyond its maximum capability, causing shutdown and the potential for damage. Thus the converter provides the ability to compensate for a lack of fuel cell power output to improve performance, but there is a limit on this ability that must be respected.

While acceleration was close to the team's target values, Figure 52 shows room for improvement. The peak battery power output was only 60 kW out of 100 kW available. The rear motor also peaked at 120 Nm out of a possible 160 Nm allowed by the control strategy.

This is due to a torque-limiting strategy that was implemented to address the issue of over-current conditions in the DC-DC converter, which had been experienced several times during testing. Without sufficient time to develop and validate more complex controls the team implemented a general reduction in overall vehicle torque requests, experimentally adjusting the gain until the DC-DC converter current limits were not exceeded during full acceleration.

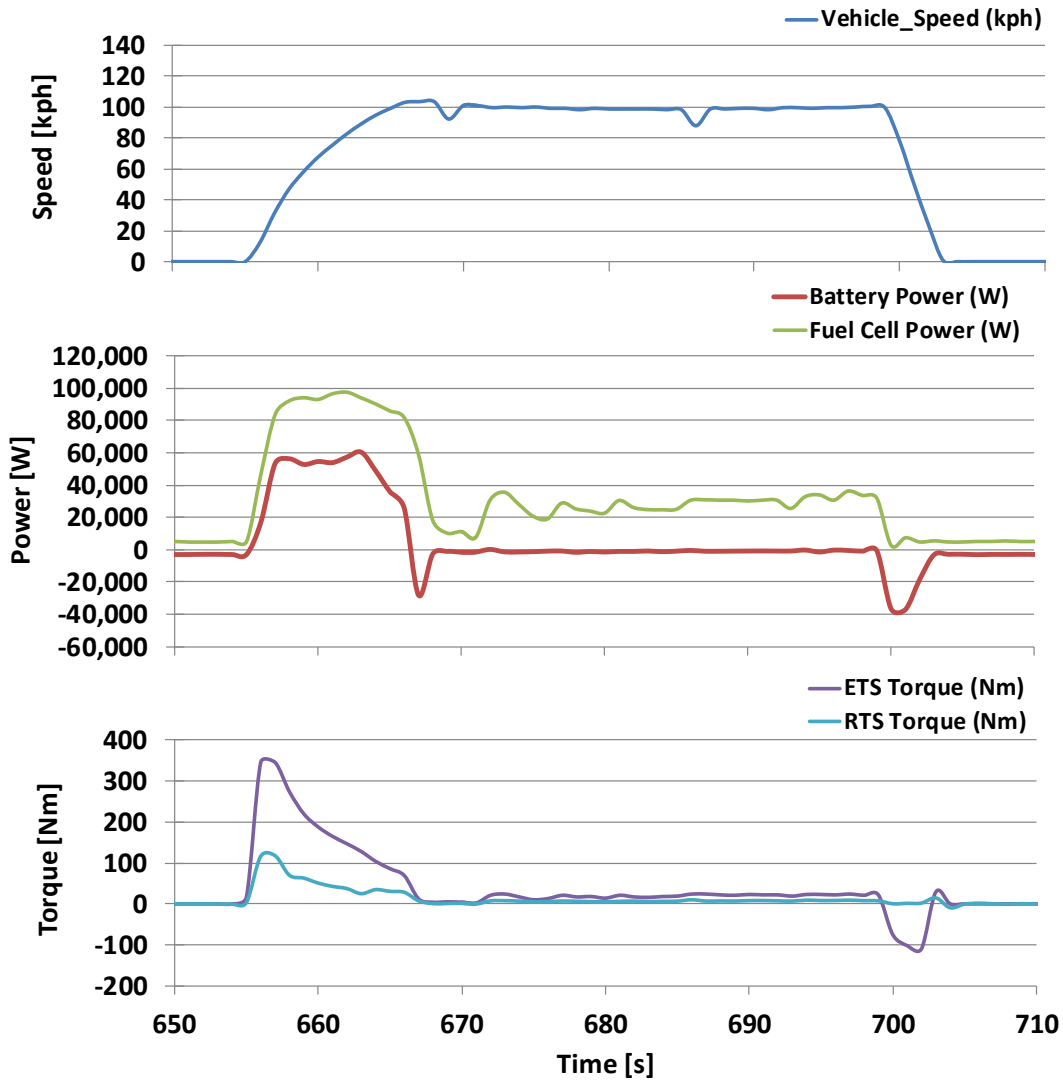


Figure 52: Speed, power, and torque for 0-60 mph acceleration and braking test

7.2 Braking

Braking performance was poorer than the team expected. Several mechanical reasons have been proposed for the high braking distance, mostly relating to tire and brake condition. The DC-DC converter also plays a role, however, due to regenerative braking. The DC-DC converter limits front motor regenerative braking to 38 kW, which is clearly evident in Figure 52 (observe the red battery power line during the regenerative braking event at 700 seconds). This limits the maximum regenerative torque to a little over 100 Nm. The original, higher power converter would have allowed for the full 50 kW regenerative capability of the battery to be realized, reducing the braking distance. Simultaneously using the rear motor for regenerative braking would also make up for 12 kW shortfall, however this was not implemented in time for the competition.

7.3 Energy Efficiency

The energy consumption of the vehicle is evaluated in three separate trips, termed Schedule A, Schedule B, and Schedule C. These schedules have distances of 33 km, 66 km, and 166 km, respectively. Each driving schedule is comprised of the same base drive cycles, representing city and highway driving. The schedules are therefore only differentiated by the number of times those cycles are repeated. The driving schedules are executed by trained drivers on a closed and planned course, with the aim to maximize the repeatability of the tests.

Route statistics for all three schedules are shown in Table 11. Energy consumption figures for each schedule, for both electricity and hydrogen, are shown in Table 12. The UWAFT EcoCAR drive cycle from Schedule A is shown in Figure 53.

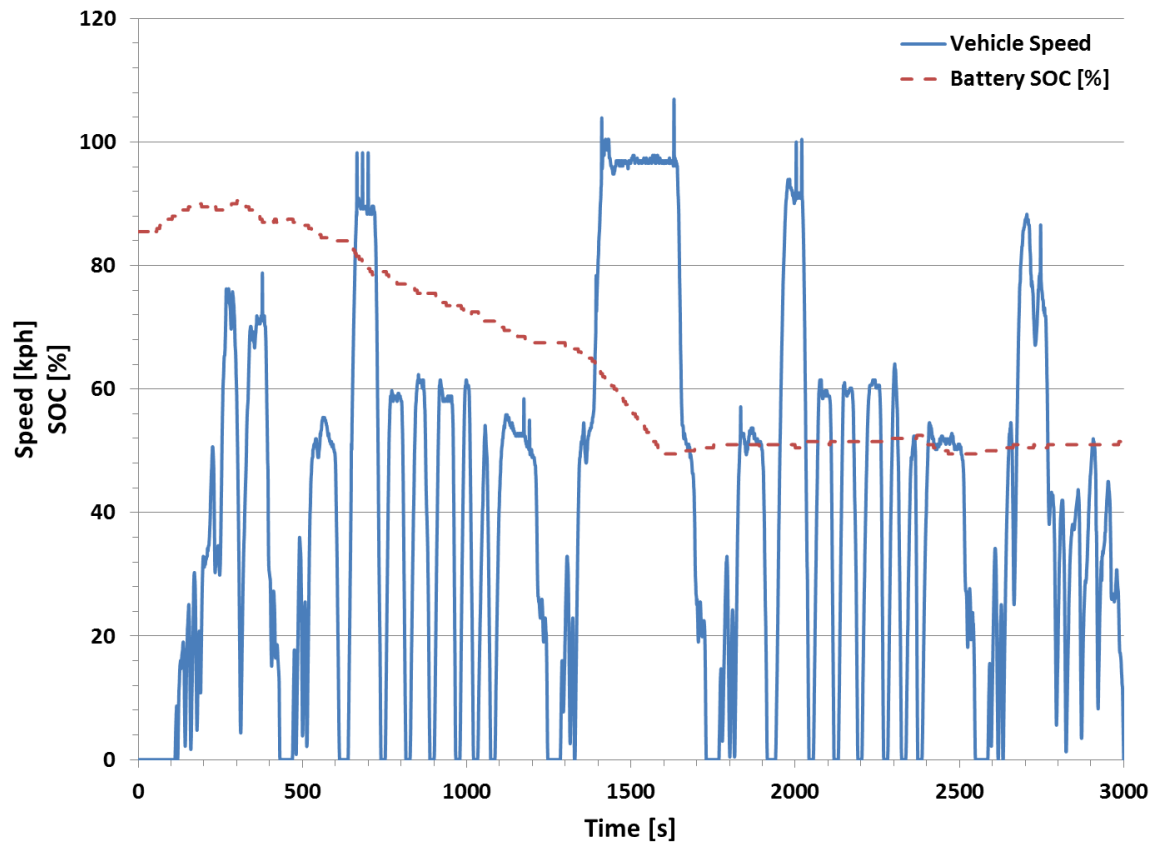


Figure 53: E&EC drive cycle, schedule A

Table 11: Emissions and Energy Consumption Event - Drive Schedule Statistics

DRIVE SCHEDULE	DISTANCE (km)	AVG. SPEED (km·h ⁻¹)	MAX. SPEED (km·h ⁻¹)	AVG. ACCELERATION (m·s ⁻²)
A	32.7	36.3	106.9	0.183
B	66.0	42.8	106.1	0.182
C	166.0	44.7	107.4	0.191
TOTAL	264.8	41.3	107.4	0.185

Table 12: Emissions and Energy Consumption Event – Vehicle Energy Use

DRIVE	H2 USED (kWh)	AC ELECTRICITY USED (kWh)	ENERGY USE (kWh·km ⁻¹)	ENERGY USE (L _{eq} per 100 km)	FUEL ENERGY USE (L _{eq} per 100 km)
A	13.6	5.63	0.59	6.61	4.67
B	30.6	7.12	0.57	6.41	5.20
C	85.1	5.89	0.55	6.16	5.77
TOTAL	3.88	18.6	0.56	6.28	5.49

The vehicle achieved a utility factor weighted fuel consumption of 5.5 L per 100 km, which is 72% greater than the VTS target of 3.2 L per 100 km. There are several reasons why this target was not achieved. The utility factor weighted fuel economy is calculated according to Equation 8.

$$FC_{UFW} = K \cdot \left[\left(\frac{FC_{CD(l)}}{km_{CD}} \right) (UF) + \left(\frac{FC_{CS(l)}}{km_{CS}} \right) (1 - UF) \right] \quad \text{Equation 8}$$

In this formula there are three contributions to the fuel consumption, the charge-depleting (CD) fuel consumption (L per 100 km), the charge-sustaining (CS) fuel consumption (L per 100 km), and the utility factor. The constant “K” simply allows conversion between fuel consumption units. The CD and CS fuel consumption figures are the per kilometer hydrogen fuel use (typically in kilograms) consumed during the charge-depleting and charge-sustaining portions of the drive, respectively. They represent the energy efficiency of the vehicle in each of these modes from a hydrogen consumption perspective; however electrical energy from the grid (stored in the battery) is not counted in the fuel consumption calculation. Hence as the reliance on grid electricity increases (increasing the utility factor), the fuel consumption figure becomes increasingly separated from an indication of vehicle energy efficiency.

The utility factor represents the fraction of the drive which is completed while discharging the battery pack, hence displacing fuel (here, hydrogen) with grid electricity. One contributing factor to the higher fuel consumption is the loss of battery capacity when the pack was reduced from 330 V to 250 V. This in turn reduced the charge-depleting range and hence the utility factor. Since the vehicle relies mostly on battery energy during charge-depleting operation, and since battery electricity is not counted as “fuel”, a large utility factor

amplifies the weighting of the small fuel consumption in charge-depleting mode, and so it has a significant effect on the overall fuel economy. Utility factor is also highly non-linear, so a small loss in CD range has a significant effect on the utility factor at low ranges.

The significance of the utility factor is evident in the VTS. The targeted 3.2 L per 100 km weighted fuel consumption and the 4.9 L per 100 km charge-sustaining fuel consumption are differentiated by the grid electricity which offsets hydrogen fuel use in the first part of the drive cycle. Thus the utility factor weighted fuel economy is not a clear measure of vehicle efficiency, unlike fuel economy in conventional powertrains.

The charge-depleting fuel consumption was assumed to be zero during simulation and development of the VTS. The vehicle control strategy at that time included a pure battery electric mode. After powertrain component changes and control system testing the team decided to run the fuel cell during both charge-depleting and charge-sustaining operation, causing the fuel consumption in CD mode to be non-zero. This has the simultaneous effect of increasing km_{CD} and reducing FC_{CD} . The net effect requires more analysis and comparison to the original powertrain simulations, however this may have contributed to the higher fuel consumption.

One of the largest effects on the utility factor was that the state of charge of the battery was sustained at 50%, where original simulations called for 20-30%, and even lower was possible since battery longevity is not a significant concern in the competition. The maximum state of charge was also reduced from 100% to between 90% and 95% during competition. These set-point changes were made to the control strategy to guarantee reliability in the absence of sufficient testing time. For example the team observed that the battery controller would sometimes skip from a low state of charge directly to zero, which would trigger an unexpected vehicle shutdown. The charge-sustaining limit was raised to prevent this from occurring during competition.

Another significant error in the simulated efficiency was neglecting the efficiency of the battery charger. The efficiency of charging and discharging the cells of the battery pack (due

to polarization losses) was accounted for, but the auxiliary loads (e.g. pumps, fans, controllers) used during charging were not accounted for in the simulation.

As discussed in Chapter 5, the torque split and hybrid power split control strategies can also be improved. For example, by disabling the rear motor until it is absolutely needed (and simultaneously solving the drivability issues resulting from enabling/disabling the rear motor during driving) the vehicle may be able to save power and reduce operation in inefficient low-torque regions. The hybrid control strategy could also be modified to reduce the amount of time the fuel cell operates below the 10 kW power level. Below this level fuel cell system efficiency drops quickly below 50%.

7.4 Range

Range is the product of energy efficiency, measured in kWh/km, and total useable energy, measured in kWh. The vehicle range is determined by dividing the vehicle energy efficiency (evaluated as described in section 7.3) by the vehicle’s energy storage capacity. This extrapolation only applies to the vehicle’s main fuel source (gasoline, diesel, hydrogen). Teams only receive credit for the amount of electric range that the vehicle actually demonstrates, as this is part of the vehicle control strategy being evaluated.

The UWAFTEcoCAR achieved 85% of the competition range requirement and 78% of the team’s objective. Some of this shortfall is attributable to a several reductions in the vehicle’s usable energy in the final year of the competition, while the rest can be attributed to efficiency gaps. These reductions are summarized in Table 13, in kWh and as a fraction of the 150 kWh of usable energy that the team has originally expected to have on-board.

Table 13: Vehicle usable energy reductions in year 3

REDUCTION	(kWh)	(%)
H2 Tank Loss	9.99	6.7%
Battery Size Loss	1.94	1.3%
Battery CS Loss	3.39	2.3%
Battery Charge Loss	0.48	0.3%

The battery pack was the source of several losses in usable energy. The loss of battery capacity in moving to a 250 V pack (see section 5.1) had the obvious effect of reducing the total energy by 2.42 kWh, and usable energy by 1.94 kWh or 1.3%. Charge-sustaining at a higher state of charge than originally planned also reduced the total usable energy, costing an additional 3.39 kWh or 2.3%. Additionally, due to a programming error in the supervisory control code, the vehicle could not be started if the battery voltage was too high. This meant that the team could only charge to about 95%, incurring a loss of another (approximately) 0.5 kWh or 0.3%.

Another source of the range shortfall turned out to be an over-estimation of the usable hydrogen available to the vehicle in the original simulations. The hydrogen storage system has a minimum allowable pressure that was not accounted for in simulation. That is, the final 0.3 kg of stored hydrogen cannot be used by the vehicle, resulting in a loss of 9.99 kWh or 6.7%. This combination of factors led to a 10.5% reduction in the vehicle's usable energy, accounting for some of the 22% range shortfall. The remaining 11.5% range loss is due to differences in the expected / simulated efficiency of the components, the vehicle chassis, and the control strategy, some of which are addressed in section 7.3.

7.5 Petroleum Energy Use and Greenhouse Gas Emissions

The petroleum energy use (PEU) and greenhouse gas emissions (GHG) results, compared to the simulated values, reflect the aforementioned differences in usable energy capacity and energy conversion efficiency for hydrogen and electricity.

The petroleum energy use target was achieved; however greenhouse gas emissions were underestimated. Since GHG and PEU factors are a constant for each fuel throughout the competition the only variable is energy efficiency in kWh·km⁻¹ for each fuel, weighted by the utility factor. Thus all of the factors discussed in section 7.3 directly impact PEU and GHG.

Chapter 8

Conclusions and Recommendations

8.1 Conclusions

This thesis describes the development of the UWAFTEcoCAR vehicle which is a hydrogen fuel cell plug-in hybrid electric vehicle designed and built by students at the University of Waterloo. The vehicle uses a unique split high voltage bus architecture that places a power source and motor on each side of the DC-DC converter, allowing maximum utilization of each power source while using a small DC-DC unit. A total of 150 kWh of energy is carried on-board, split between 4.2 kg of hydrogen stored at 70 MPa and 9.68 kWh of lithium-ion batteries. The fuel cell and battery combine to deliver up to 195 kW of power to two electric motors that can together provide up to 540 Nm of (pre-gearbox) torque. A 150 A DC-DC converter links both sides of the powertrain. The vehicle successfully completed all dynamic evaluations at the EcoCAR competition, placing 3rd out of 16 North American universities.

The thesis also describes UWAFTE, which is a project-based learning environment where students gain core engineering technical knowledge and higher-level application-specific knowledge by executing design projects as a part of a vehicle development team. These design projects are highly inter-dependent and encourage teamwork within and across disciplines. They also expose students to essential elements of engineering design such as dealing with uncertainty, consideration for reliability and robustness, and safety analysis. These elements are not typically explored in the undergraduate engineering curriculum but are required of practicing engineers.

The team's successful competition performance is attributed to the choice of a highly efficient hydrogen powertrain, hard work by team members to integrate that powertrain, and good design processes to balance performance, reliability, and safety in the final implementation. The thermal, powertrain, and safety control system design projects demonstrate some of the processes that were followed and how they benefit both student learning and team outcomes.

The thermal control system design project discussed in Chapter 4 highlights several sources of uncertainty in real-world engineering design. While students are already familiar with the core engineering concepts needed to complete the design (fluid mechanics, heat transfer), these inherent uncertainties present new challenges that often halt the progress of new team members' design projects. The design process discussed in Chapter 4 is used to help undergraduate students identify sources of uncertainty in design projects and to apply strategies to handle them. A lack of information with regards to coolant and air flow restrictions within the low temperature cooling system is handled by making (documented) assumptions about the pressure drop within components and performing the analysis over a range of air flow rates to demonstrate design feasibility under a variety of conditions. This process is not only a tool to help individual students, it also helps the team organize the design so that future students can refine it as conditions change or more information becomes available. Under full load conditions, and assuming 40°C ambient air with a flow rate of at least 2000 CFM, a Bosch PCA12V pump and Dana Thermal HX-260 radiator are selected to yield a coolant flow rate of 15 LPM and 17.1 kW of heat dissipation. Further study of the air and coolant flow restrictions will increase the certainty of the design performance.

Engineering design also involves considerations that go far beyond function or performance. In safety-critical applications steps must be taken to minimize failures and ensure the reliability of a design. In case of failures or unexpected conditions a robust design should also continue to provide some level of functionality. This is important in EcoCAR, where the design must function through hundreds of kilometers of fault-free real-world driving. The design decisions described in Chapter 5 demonstrate how the final powertrain architecture control algorithms are impacted by these considerations.

In the developed vehicle the rear motor is re-located to the low side of the high voltage bus to maintain current availability with a smaller DC-DC. Due to the loss of two DC-DC quadrants of operation the battery voltage is reduced from 330 V to 250 V. This maximizes reliability by simplifying the algorithms that control voltage and current on both sides of the powertrain high voltage bus.

The power control algorithm takes optimal control set points from the power split strategy and processes them to arrive at a commanded power split that attempts to meet the efficiency set point provided by the power split while also satisfying driver performance demands and component operational requirements (e.g. SOC targets and DC-DC current limit). Similarly, the torque control algorithm receives torque set points from the torque split strategy and processes them to ensure the reliability and robustness of the powertrain. Torque set points for each motor are limited to prevent voltage crossover between the high and low sides of the bus and to respect operational limitations reported by the motor controllers. Torque limited on one motor is shifted to the other motor if capacity exists, which in the extreme case of a motor fault provides a level of robustness to the design. These control projects give students the opportunity to add reliability and robustness considerations to their decision making process in an environment where success depends heavily on them.

Vehicle design projects must incorporate safety analyses to ensure that passengers are protected from unexpected failures. FMEA and FTA analyses are applied to the torque control pathways, using top-down and bottom-up approaches to identify and link together component failures and vehicle faults. This is demonstrated by examining the torque control subsystem and the role of the accelerator pedal sensor. Analysis identifies specific failures that can lead to a vehicle level unintended acceleration fault, including invalid in-range and out-of-range sensor values, lost CAN bus communications, and concurrent application of the brake and accelerator pedal. These are detected by the supervisory controller and mitigating actions are taken to protect and warn the driver. These behaviours are validated on a hardware-in-the-loop simulator platform. Effective implementation of these safety analyses is aided by a structured approach that categorizes types of failures according to the type of sensor or controller being examined. It is also critically important that the brainstorming sessions that generate starting points for FTA and FMEA analyses involve a team of system-level and component-level experts; the failure of the ETS torque enable signal is an example of the consequences of incomplete failure mode coverage. Students who learn control design within UWAFST learn to integrate safety considerations into every step of the design process,

from signal acquisition to algorithm design to system networking. This makes safety an inherent part of the core technical knowledge that students learn and apply in future projects.

The competition on-road vehicle evaluation provides an opportunity to examine how design decisions impacted the vehicle's performance. Vehicle acceleration is 9.3 seconds from 0 to 102 km·h⁻¹ and 8.4 seconds for 80 to 112 km·h⁻¹. Torque reductions implemented to protect the new DC-DC converter from voltage crossover led to slower than expected vehicle acceleration at high power outputs (e.g. 80 to 112 km·h⁻¹). The vehicle achieved a range of 273 km, 78% of the expected 350 km. Some of this loss is attributed to reductions in usable energy that were partly caused by decisions made to prioritize the potential reliability of the vehicle over performance, and partly by overestimating usable fuel capacity. In total, 134.2 kWh or 89% of the expected 150 kWh was actually available during competition.

Project-based learning within UWAFST is shown to have significant potential benefits. Students gain skills to complement and extend the technical knowledge learned in the classroom, skills which are valuable in real-world engineering design projects. However, some challenges exist. One issue faced by UWAFST students is that an unstructured learning approach requires more motivation to apply self-directed learning. Design processes like the one discussed in Chapter 4 allows the team to provide help to students who need guidance as they navigate new design problems. Another commonly raised issue is that an unstructured and self-directed learning format may fail to teach the desired concepts. This can be mitigated by supporting project-based learning with tutorials and mentoring. The EcoCAR experience provides this support for a range of subject matter, including safety analyses and control system design, to provide students with guidance and assistance where needed.

Student teams like UWAFST can offer a rich learning environment for undergraduate engineering students. The design projects and resources that they provide can significantly advance student knowledge, experience, and skills in a way that complements the technical knowledge gained in the classroom. Finding ways to provide these experiences to more undergraduate students, either outside or within existing core courses, has the potential to greatly enhance the value of graduates.

8.2 Recommendations

Recommendations from each Chapter pertain to both technical execution of the vehicle design and project-based learning. These are summarized below.

1. Further study of coolant and air flow rates for thermal management systems.

One source of uncertainty in the thermal design project is the resistance to coolant and air flow within the system. Some assumptions and sensitivity analyses help to understand how these uncertain variables impact the final design when such information is lacking. It would be ideal to generate a test rig that can perform pressure drop experiments when components become available to refine the system model and improve the quality of future designs.

2. Use design projects as an opportunity to enhance student approaches to uncertainty

The design case study in Chapter 4 highlights opportunities for students to improve their ability to approach problems that lack the clarity and completeness of a textbook analysis. These conditions resemble those faced by practicing engineers and can help students get past a “one right answer” mentality, generating an analysis that considers a range of solutions based on assumptions and “what-if” or “worst-case” scenarios that can be refined and validated. It may be beneficial to examine how these skills can be explored in the curriculum.

3. Explore efficiency impact of separating strategy from control

The separation of strategy and control discussed in Chapter 5 has the benefit of simplifying parallel development of these two control units and also making rapid prototyping of new “optimizing” control strategies safer. However, by not considering any (or as many) of the instantaneous operational restrictions of powertrain components, the control strategies may generate the most efficient set points under ideal conditions but not that which is best for current conditions. Adding these considerations to the control strategies should be considered.

4. Implement dynamic torque moderation in response to high voltage bus conditions

Due to polarization electric vehicles must take care to limit voltage drop under load, typically at low battery SOC or temperature where polarization effects are greatest. The EcoCAR powertrain, with a split high voltage bus and limited two-quadrant DC-DC, is prone to this issue under all conditions. The simple torque reductions discussed in Chapter 5 should be expanded to dynamically set maximum torque set points as a function of the current and

predicted future bus voltage. This approach requires more validation effort to ensure reliability but will unlock additional vehicle performance.

5. Offer tutorials and help sessions to enhance the value of project-based learning

Self-directed student learning can require significantly more motivation and time than lecture-based instruction, especially when the material is unfamiliar. New UWAFB students would benefit from the kind of tutorials that are offered at AVTC competition workshops, where an overview of the theory behind various core vehicle design activities helps students to get familiar with the concepts they need to learn in order to complete their project. This does not have to be a complete lecture of all relevant material, simply a primer to get started.

6. Consider the impact of operational requirements on usable energy and efficiency

In predicting the efficiency and range of the vehicle one key mistake was neglecting to account for the operational requirements of the real vehicle. Things like not being able to use all the fuel in the hydrogen tanks, having to run auxiliary loads during charging, and limiting battery SOC contributed to reduced efficiency and range in competition. These factors can add together to have a significant impact on the final performance of the real vehicle.

7. More integration of student teams and project-based learning in the curriculum

Engineering student teams like UWAFB offer learning experiences and resources that are unmatched in the undergraduate engineering curriculum. The problems faced by teams add additional dimensions of uncertainty, reliability, robustness, and safety to the engineering design process. Using case studies like those in Chapters 4-6 as a basis for project-based learning opportunities in core undergraduate engineering courses can enhance students' ability to apply core technical material in real-world engineering design problems.

References

- [1] W. McDowall and M. Eames, "Forecasts, scenarios, visions, backcasts and roadmaps to the hydrogen economy: A review of the hydrogen futures literature," *Energy Policy*, vol. 34, no. 11, pp. 1236–1250, Jul. 2006.
- [2] J. C. Perrenet, "The Suitability of Problem-based Learning for Engineering Education: Theory and practice," *Teach. High. Educ.*, vol. 5, no. 3, pp. 345–358, Jul. 2000.
- [3] N. Hosseinzadeh, "Application of Project-Based Learning (PBL) to the Teaching of Electrical Power Systems Engineering," *IEEE Trans. Educ.*, vol. 55, no. 4, pp. 495–501, Nov. 2012.
- [4] Julie E. Mills and David F. Treagust, "Engineering Education - Is Problem-Based or Project-Based Learning the Answer?," *Australas. J. Eng. Educ.*, Dec. 2003.
- [5] US Energy Information Administration, "International Energy Statistics," *US Energy Information Administration*, 2012. [Online]. Available: <http://www.eia.gov/cfapps/ipdbproject/IEDIndex3.cfm>.
- [6] Natural Resources Canada, "Energy Use Data Handbook: 1990 to 2008," Natural Resources Canada, Office of Energy Efficiency, Ottawa, Ontario, 2011.
- [7] Canadian Engineering Accreditation Board, "Canadian Engineering Accreditation Board Accreditation Criteria and Procedures 2008." Sep-2008.
- [8] Pennwell Corporation, "Worldwide Look at Reserves and Production," *Oil Gas J.*, vol. 109, no. 49, p. 28, 2011.
- [9] BP plc, "BP Statistical Review of World Energy 2012," BP plc, London, Jun. 2012.
- [10] N. A. Owen, "The status of conventional world oil reserves—Hype or cause for concern?," *Energy Policy*, vol. 38, no. 8, pp. 4743–4749, Aug. 2010.
- [11] US Energy Information Administration, "Canada," *US Energy Information Administration*, 17-Sep-2012. [Online]. Available: <http://www.eia.gov/countries/cab.cfm?fips=CA>.
- [12] The World Bank, "World Development Indicators: Traffic and congestion," *The World Bank*, 2013. [Online]. Available: <http://wdi.worldbank.org/table/3.13>.
- [13] D. D. Parrish, "Air quality progress in North American megacities: A review," *Atmos. Environ.*, vol. 45, no. 39, pp. 7015–7025, Dec. 2011.
- [14] J. M. Baldasano, E. Valera, and P. Jiménez, "Air quality data from large cities," *Sci. Total Environ.*, vol. 307, no. 1–3, pp. 141–165, May 2003.
- [15] D. Dodman, "Blaming cities for climate change? An analysis of urban greenhouse gas emissions inventories," *Environ. Urban.*, vol. 21, no. 1, pp. 185–201, Apr. 2009.
- [16] US Environmental Protection Agency, "Inventory of U.S. Greenhouse Gas Emissions and Sinks: 1990-2011," U.S. Environmental Protection Agency, Washington, DC, EPA 430-R-13-001, Apr. 2013.
- [17] "Understanding General Motors Hybrid Electric Vehicles." General Motors, 07-Jul-2006.
- [18] "EcoCAR Non Year Specific Rules." Argonne National Laboratories, 25-Sep-2010.
- [19] National Institute of Standards and Technology, "Isothermal Properties for Hydrogen," *NIST Chemistry WebBook*. [Online]. Available: http://webbook.nist.gov/cgi/fluid.cgi?TUnit=C&PUnit=MPa&DUnit=kg%2Fm3&HUnit=kJ%2Fkg&WUnit=m%2Fs&VisUnit=Pa*s&STUnit=N%2Fm&Type=IsoTherm&RefState=DEF&Action=Page&ID=C1333740.
- [20] G. Frenette, "Economic & commercial viability of hydrogen fuel cell vehicles from an automotive manufacturer perspective," *Int. J. Hydrog. Energy*, vol. 34, no. 9, pp. 3578–3588, May 2009.

- [21] Chunzhi He, Sanket Desai, Garth Brown, and Srinivas Bollepalli, "PEM Fuel Cell Catalysts: Cost, Performance, and Durability," *Electrochem. Soc. Interface*, vol. Fall 2005, pp. 41–44.
- [22] B. C. H. Steele and A. Heinzl, "Materials for fuel-cell technologies," *Nature*, vol. 414, no. 6861, pp. 345–352, Nov. 2001.
- [23] G. Frenette and D. Forthoffer, "Economic & commercial viability of hydrogen fuel cell vehicles from an automotive manufacturer perspective," *Int. J. Hydrog. Energy*, vol. 34, no. 9, pp. 3578–3588, May 2009.
- [24] "Honda FCX Clarity - Vehicle Specifications - Official Web Site." [Online]. Available: <http://automobiles.honda.com/fcx-clarity/specifications.aspx>. [Accessed: 22-Nov-2013].
- [25] US Environmental Protection Agency, "2013 Nissan Leaf," *www.fueleconomy.gov*. [Online]. Available: <http://www.fueleconomy.gov/feg/Find.do?action=sbs&id=33558>.
- [26] R. Ono, "Minimum ignition energy of hydrogen-air mixture: Effects of humidity and spark duration," *J. Electrostat.*, vol. 65, no. 2, pp. 87–93.
- [27] Fuel Cell & Hydrogen Energy Association, "Hydrogen Safety." [Online]. Available: http://www.fchea.org/core/import/PDFs/factsheets/Hydrogen%20Safety_NEW.pdf.
- [28] R. Zimmer, "Let's go green with hydrogen! The general public's perspective," *Int. J. Hydrog. Energy*, vol. 37, no. 22, pp. 17502–17508, Nov. 2012.
- [29] US Environmental Protection Agency, "Dynamometer Drive Schedules," *United States Environmental Protection Agency*. [Online]. Available: <http://www.epa.gov/nvfel/testing/dynamometer.htm>. [Accessed: 07-Jul-2013].
- [30] Government of Canada, "Vehicle Warm-Up," *Natural Resources Canada*, 21-Apr-2009. [Online]. Available: <http://oee.nrcan.gc.ca/transportation/idling/14743>. [Accessed: 22-Nov-2013].
- [31] Society of Automotive Engineers, "Recommended Practice for Measuring the Exhaust Emissions and Fuel Economy of Hybrid-Electric Vehicles, Including Plug-In Hybrid Vehicles." Society of Automotive Engineers, Jun-2010.
- [32] Environmental Protection Agency, "Many Factors Affect Fuel Economy," *www.fueleconomy.gov*. [Online]. Available: <http://www.fueleconomy.gov/feg/factors.shtml>. [Accessed: 18-Aug-2013].
- [33] US Environmental Protection Agency, "Compare side-by-side," *www.fueleconomy.gov*. [Online]. Available: <http://www.fueleconomy.gov/feg/Find.do?action=sbs&id=33374&id=33372>.
- [34] I. Nonaka, "SECI, Ba and Leadership: a Unified Model of Dynamic Knowledge Creation," *Long Range Plann.*, vol. 33, no. 1, pp. 5–34.
- [35] J. Savery, "Overview of Problem-based Learning: Definitions and Distinctions," *Interdiscip. J. Probl.-Based Learn.*, vol. 1, no. 1, May 2006.
- [36] W. Hung, J. Harpole Bailey, and D. H. Jonassen, "Exploring the Tensions of Problem-Based Learning: Insights From Research," *New Dir. Teach. Learn.*, vol. 2003, no. 95, pp. 13–23, 2003.
- [37] Harold B. White, "Dan tries problem-based learning: A case study," *Improve Acad.*, vol. 15, pp. 75–91, 1996.
- [38] J. A. Colliver, "Effectiveness of Problem-based Learning Curricula: Research and Theory," *Acad. Med. March 2000*, vol. 75, no. 3, pp. 259–266, 2000.
- [39] D. A. Kolb, *Experiential learning: experience as a source of learning and development*. Englewood Cliffs, N.J.: Prentice-Hall, 1984.
- [40] D. A. Kolb, R. E. Boyatzis, and C. Mainemelis, "Experiential learning theory: Previous research and new directions," in *Perspectives on Thinking, Learning, and Cognitive Styles*, Routledge, 2001.
- [41] S. Kundu and M. W. Fowler, "Use of Engineering Design Competitions for Undergraduate and Capstone Projects," *Chem. Eng. Educ.*, vol. 43, no. 2, pp. 131–136, Mar. 2009.

- [42] Michael Giannikouris, Dr. Roydon Fraser, and Dr. Michael Fowler, "Development of a Fuel Cell Plug-in Hybrid Electric Vehicle at the University of Waterloo," in *Proceedings, EV2011 Conference and Tradeshow*, 2011, pp. 1–13.
- [43] A. K. Koch, "Hybrid Controls Development and Optimization of a Fuel Cell Hybrid Powertrain," University of Waterloo, Waterloo, Ontario, Canada, 2012.
- [44] "Power Management and Distribution Unit (PMD) Component Technical Specification [GM Confidential]." General Motors, 22-Nov-2006.
- [45] "Power Inverter Module (PIM) Component Technical Specification [GM Confidential]." General Motors, 27-Jan-2009.
- [46] "FCS Loads and interface info for Ecocar Team [GM Confidential]." General Motors.
- [47] Microelectronics Heat Transfer Laboratory, "Fluid properties calculator," 1997. [Online]. Available: <http://www.mhtl.uwaterloo.ca/old/onlinetools/airprop/airprop.html>.
- [48] C. J. Atman, R. S. Adams, M. E. Cardella, J. Turns, S. Mosborg, and J. Saleem, "Engineering Design Processes: A Comparison of Students and Expert Practitioners," *J. Eng. Educ.*, vol. 96, no. 4, pp. 359–379, 2007.
- [49] Padma Sundaram, "Fault Tree Analysis," presented at the EcoCAR Tech Talk, University of Waterloo, 22-Apr-2009.
- [50] Tom Achatz, "What is FMEA?," presented at the EcoCAR Tech Talk, University of Waterloo, 29-Apr-2009.
- [51] J. J. Stadler, "Software failure modes and effects analysis," *2013 Proc. Annu. Reliab. Maintainab. Symp. RAMS*, pp. 1–5, Jan. 2013.
- [52] Michael Giannikouris, Mark Cremasco, Carl Chan, and Josh Lo, "DC-DC Power Converter for the University of Waterloo Fuel Cell Plug-in Hybrid Electric Vehicle." University of Waterloo, Jul-2011.
- [53] US Department of Energy, "Comparison of Fuel Cell Technologies," *Energy Efficiency and Renewable Energy*, Feb-2011. [Online]. Available: http://www1.eere.energy.gov/hydrogenandfuelcells/fuelcells/pdfs/fc_comparison_chart.pdf.
- [54] Frank M. White, *Fluid Mechanics*, Sixth. McGraw-Hill, 2006.
- [55] "Engineering Data," *Crane Engineering*. [Online]. Available: <http://www.craneengineering.net/products/valves/documents/craneEngineeringValveEngineeringData.pdf>.

Appendix A

Tabulated Comparison of Fuel Cell Technologies

Table 14: Comparison of fuel cell technologies [53]

Fuel Cell Type	Common Electrolyte	Operating Temperature	Typical Stack Size	Efficiency	Applications	Advantages	Disadvantages
Polymer Electrolyte Membrane (PEM)	Perfluoro sulfonic acid	50-100°C 122-212° typically 80°C	< 1kW-100kW	60% transportation 35% stationary	<ul style="list-style-type: none"> • Backup power • Portable power • Distributed generation • Transportation • Specialty vehicles 	<ul style="list-style-type: none"> • Solid electrolyte reduces corrosion & electrolyte management problems • Low temperature • Quick start-up 	<ul style="list-style-type: none"> • Expensive catalysts • Sensitive to fuel impurities • Low temperature waste heat
Alkaline (AFC)	Aqueous solution of potassium hydroxide soaked in a matrix	90-100°C 194-212°F	10-100 kW	60%	<ul style="list-style-type: none"> • Military • Space 	<ul style="list-style-type: none"> • Cathode reaction faster in alkaline electrolyte, leads to high performance • Low cost components 	<ul style="list-style-type: none"> • Sensitive to CO₂ in fuel and air • Electrolyte management
Phosphoric Acid (PAFC)	Phosphoric acid soaked in a matrix	150-200°C 302-392°F	400 kW 100 kW module	40%	<ul style="list-style-type: none"> • Distributed generation 	<ul style="list-style-type: none"> • Higher temperature enables CHP • Increased tolerance to fuel impurities 	<ul style="list-style-type: none"> • Pt catalyst • Long start up time • Low current and power
Molten Carbonate (MCFC)	Solution of lithium, sodium, and/or potassium carbonates, soaked in a matrix	600-700°C 1112-1292°F	300 kW-3 MW 300 kW module	45-50%	<ul style="list-style-type: none"> • Electric utility • Distributed generation 	<ul style="list-style-type: none"> • High efficiency • Fuel flexibility • Can use a variety of catalysts • Suitable for CHP 	<ul style="list-style-type: none"> • High temperature corrosion and breakdown of cell components • Long start up time • Low power density
Solid Oxide (SOFC)	Ytria stabilized zirconia	700-1000°C 1202-1832°F	1 kW-2 MW	60%	<ul style="list-style-type: none"> • Auxiliary power • Electric utility • Distributed generation 	<ul style="list-style-type: none"> • High efficiency • Fuel flexibility • Can use a variety of catalysts • Solid electrolyte • Suitable for CHP & CHHP • Hybrid/GT cycle 	<ul style="list-style-type: none"> • High temperature corrosion and breakdown of cell components • High temperature operation requires long start up time and limits

Appendix B

Underbody View of Final Vehicle Showing Powertrain Integration

GM Electric Traction System
Power: 120 kW Peak, 80 kW Con't
Voltage: 240 – 450 VDC



Ballard Power Systems Motor
Power: 67 kW Peak, 32 kW Con't
Voltage: 100 – 350 VDC

A123 Energy Storage System
Power + Energy: 100kW, 9.7 kWh
Voltage: 240 – 255 VDC



Brusa BCD412 Bidirectional DC/DC
Power: 38 [kW] peak
Voltage: 100 – 425 VDC



Brusa NLG513-Sx Charger
Power: 3.3 kW
Voltage: 260 – 520 VDC



Figure 54: Vehicle underbody view showing powertrain integration

Appendix C

Front Cooling Loop Thermal-Fluid Calculations

Table 15: 50/50 ethylene glycol/water properties at 70°C [47]

PROPERTY	VALUE
Density ($\text{kg}\cdot\text{m}^{-3}$)	1029
Dynamic Viscosity ($\text{kg}\cdot\text{m}^{-1}\cdot\text{s}^{-1}$)	0.00107
Kinematic Viscosity ($\text{m}^2\cdot\text{s}^{-1}$)	1.0388E-06
Specific Heat ($\text{J}\cdot\text{kg}^{-1}\cdot\text{K}^{-1}$)	3398

Table 16: Pipe properties, friction factor from White [54]

PROPERTY	VALUE
Pipe Diameter (m)	0.019
Pipe Cross-Sectional Area (m)	0.000284
Pipe Length (m)	5
Pipe Roughness (mm)	0.01
Relative Roughness	0.003684
Number of 90° Bends	10
90° Bend K-Factor	30*f

Table 17: Low temperature cooling loop thermal and fluid specifications

COMPONENT	MAX. COOLANT INLET TEMPERATURE (°C)	MIN.COOLANT FLOW RATE ($\text{L}\cdot\text{min}^{-1}$)	PRESSURE DROP AT MIN. FLOW RATE (Pa)	PEAK THERMAL POWER (W)	PRESSURE DROP @ 21 LPM (Pa)	PRESSURE DROP COEFFICIENT ($\text{Pa}\cdot\text{LPM}^{-2}$)
PIM	65	13	10,000	3,942	55,000	125
DU	69	13	15,000	11,158		
PMD	68	5	15,000	2,000		

Table 18: Pressure drop equations

PARAMETER	EQUATION	
Straight pipe pressure drop	$p_{straight} = \frac{1}{2} \rho f \left(\frac{L}{D} \right) V^2$	Equation 9
Straight pipe friction factor (Re > 4000)	$f_{straight} \cong \left\{ -1.8 \log \left[\left(\frac{6.9}{Re} \right) + \left(\frac{\varepsilon/d}{3.7} \right)^{1.11} \right] \right\}^{-2}$ [54]	Equation 10
Straight pipe friction factor (Re < 2300)	$f_{straight} = \frac{64}{Re}$	Equation 11
90° pipe pressure drop	$p_{90} = \frac{1}{2} \rho K_{90} V^2$	Equation 12
90° pipe loss coefficient	$K_{90} = 30 f_{straight}$ for r/D = 10 [55]	Equation 13
Component / radiator pressure drop	$p_{comp/radiator} = \frac{1}{2} \rho K_{comp/radiator} Q^2$	Equation 14
Loop pressure drop	$p_{loop} = p_{straight} + p_{90} + p_{comp} + p_{radiator}$	Equation 15
Reynolds number	$Re = \frac{\rho V}{\nu}$	Equation 16
Flow velocity	$V = \frac{Q}{A}$	Equation 17
Pipe cross-sectional area	$A = \pi \frac{D^2}{4}$	Equation 18

Table 19: Low temperature cooling loop pressure drop calculations

Flow (Q) (L·min ⁻¹)	Flow (Q) (m ³ ·s ⁻¹)	Flow Velocity (V) (m·s ⁻¹)	Re	f	K ₉₀ (Pa·s ² ·m ⁻¹)	Pipe drop (P _{straight}) (Pa)	Bend drop (P ₉₀) (Pa)	Component drop (P _{comp}) (Pa)	Radiator drop (P _{radiator}) (Pa)	Loop drop (P _{loop}) (Pa)
1	1.667E-05	0.059	1075	0.060	1.786	28	3	125	0	156
5	8.333E-05	0.294	5376	0.037	1.121	437	50	3,118	945	4,550
6	1.000E-04	0.353	6451	0.035	1.064	598	68	4,490	1,246	6,402
7	1.167E-04	0.411	7526	0.034	1.020	779	89	6,111	1,572	8,551
8	1.333E-04	0.470	8601	0.033	0.984	982	112	7,982	1,922	10,998
9	1.500E-04	0.529	9676	0.032	0.954	1,205	137	10,102	2,297	13,741
10	1.667E-04	0.588	10752	0.031	0.929	1,448	165	12,472	2,696	16,781
15	2.500E-04	0.882	16127	0.028	0.841	2,952	337	28,061	5,058	36,408
20	3.333E-04	1.176	21503	0.026	0.789	4,920	561	49,887	8,032	63,399
25	4.167E-04	1.470	26879	0.025	0.752	7,333	836	77,948	11,617	97,734
30	5.000E-04	1.763	32255	0.024	0.725	10,183	1,161	112,245	15,814	139,402
35	5.833E-04	2.057	37631	0.023	0.705	13,460	1,534	152,778	20,622	188,394
40	6.667E-04	2.351	43006	0.023	0.688	17,160	1,956	199,546	26,041	244,703
45	7.500E-04	2.645	48382	0.022	0.674	21,278	2,426	252,551	32,072	308,326
50	8.333E-04	2.939	53758	0.022	0.662	25,810	2,942	311,791	38,715	379,259
15.1	2.513E-04	0.886	16214	0.028	0.420	2,980	170	28,364	5,058	36,572

Table 20: Heat transfer equations

PARAMETER	EQUATION	
Radiator coolant heat transfer rate	$\dot{q}_R(W) = U_R \left(\frac{W}{^\circ C} \right) (T_{c,i,R} - T_{a,i,R})(^\circ C)$	Equation 19
Radiator coolant inlet temperature	$T_{c,i,R}(^\circ C) = \frac{\dot{q}_R(W)}{U_R \left(\frac{W}{^\circ C} \right)} + T_{a,i,R}(^\circ C)$	Equation 20
Component coolant heat transfer rate	$\dot{q}_C(W) = \rho \left(\frac{kg}{m^3} \right) \dot{Q}_c \left(\frac{m^3}{s} \right) c_{p,c} \left(\frac{J}{kg \cdot ^\circ C} \right) (T_{c,o,C} - T_{c,i,C})(^\circ C)$	Equation 21
Component coolant inlet temperature	$T_{c,o,C} = \frac{\dot{q}_C(W)}{\rho \left(\frac{kg}{m^3} \right) \dot{Q}_c \left(\frac{m^3}{s} \right) c_{p,c} \left(\frac{J}{kg \cdot ^\circ C} \right)} + T_{c,i,C}$	Equation 22

Table 21: Low temperature cooling loop temperature calculations for loop order FPIM, DU, PMD, radiator air flow rate 4857 [CFM]

Flow (Q) (L·min ⁻¹)	Flow (Q) (m ³ ·s ⁻¹)	R (W·C ⁻¹)	Tc,o,PMD (°C) Tc,i,R (°C)	Tc,o,R (°C) Tc,i,FPIM (°C)	Tc,o,FPIM (°C) Tc,i,DU (°C)	Tc,o,DU (°C) Tc,i,PMD (°C)	Tmax, FPIM (°C)	Tmax, DU (°C)	Tmax, PMD (°C)	q (W)
1	1.667E-05	39.9	468.4	175.4	242.9	434.1	65	69	68	17,100
2	3.333E-05	78.2	258.6	112.1	145.9	241.5	65	69	68	17,100
3	5.000E-05	115.0	188.7	91.1	113.6	177.3	65	69	68	17,100
4	6.667E-05	150.2	153.8	80.6	97.5	145.3	65	69	68	17,100
5	8.333E-05	184.0	132.9	74.3	87.8	126.1	65	69	68	17,100
6	1.000E-04	216.5	119.0	70.2	81.4	113.3	65	69	68	17,100
7	1.167E-04	247.6	109.1	67.2	76.9	104.2	65	69	68	17,100
8	1.333E-04	277.4	101.6	65.0	73.5	97.4	65	69	68	17,100
9	1.500E-04	306.0	95.9	63.3	70.8	92.1	65	69	68	17,100
10	1.667E-04	333.4	91.3	62.0	68.7	87.9	65	69	68	17,100
15	2.500E-04	454.5	77.6	58.1	62.6	75.3	65	69	68	17,100
20	3.333E-04	552.7	70.9	56.3	59.7	69.2	65	69	68	17,100
25	4.167E-04	632.4	67.0	55.3	58.0	65.7	65	69	68	17,100
30	5.000E-04	697.4	64.5	54.8	57.0	63.4	65	69	68	17,100
8.0	1.330E-04	276.9	101.8	65.0	73.5	97.5	65	69	68	17,100
9.9	1.644E-04	329.8	91.9	62.1	69.0	88.4	65	69	68	17,100
21.5	3.577E-04	577.8	69.6	55.9	59.1	68.0	65	69	68	17,100

Table 22: Low temperature cooling loop temperature calculations for loop order PMD, FPIM, DU, radiator air flow rate 694 [CFM]

Flow (Q) (L·min ⁻¹)	Flow (Q) (m ³ ·s ⁻¹)	R (W·C ⁻¹)	Tc,o,PMD (°C) Tc,i,R (°C)	Tc,o,R (°C) Tc,i,FPIM (°C)	Tc,o,FPIM (°C) Tc,i,DU (°C)	Tc,o,DU (°C) Tc,i,PMD (°C)	Tmax, FPIM (°C)	Tmax, DU (°C)	Tmax, PMD (°C)	q (W)
1	1.667E-05	32.0	574.3	280.9	315.2	382.9	65	69	68	17,100
2	3.333E-05	61.7	317.3	170.6	187.7	221.6	65	69	68	17,100
3	5.000E-05	89.1	231.9	134.1	145.5	168.1	65	69	68	17,100
4	6.667E-05	114.5	189.4	116.0	124.6	141.5	65	69	68	17,100
5	8.333E-05	137.9	164.0	105.3	112.2	125.7	65	69	68	17,100
6	1.000E-04	159.4	147.3	98.4	104.1	115.4	65	69	68	17,100
7	1.167E-04	179.2	135.4	93.5	98.4	108.1	65	69	68	17,100
8	1.333E-04	197.4	126.6	90.0	94.2	102.7	65	69	68	17,100
9	1.500E-04	214.0	119.9	87.3	91.1	98.6	65	69	68	17,100
10	1.667E-04	229.2	114.6	85.3	88.7	95.5	65	69	68	17,100
15	2.500E-04	286.7	99.7	80.1	82.4	86.9	65	69	68	17,100
20	3.333E-04	320.3	93.4	78.7	80.4	83.8	65	69	68	17,100
25	4.167E-04	338.1	90.6	78.8	80.2	82.9	65	69	68	17,100
30	5.000E-04	346.0	89.4	79.6	80.8	83.0	65	69	68	17,100
50.0	8.333E-04	346.5	89.4	83.5	84.2	85.5	65	69	68	17,100
50.0	8.333E-04	346.5	89.4	83.5	84.2	85.5	65	69	68	17,100
50.0	8.333E-04	346.5	89.4	83.5	84.2	85.5	65	69	68	17,100

Table 23: Low temperature cooling loop temperature calculations for loop order PMD, FPIM, DU, radiator air flow rate 2082 [CFM]

Flow (Q) (L·min ⁻¹)	Flow (Q) (m ³ ·s ⁻¹)	R (W·C ⁻¹)	T _{c,o,PMD} (°C) T _{c,i,R} (°C)	T _{c,o,R} (°C) T _{c,i,FPIM} (°C)	T _{c,o,FPIM} (°C) T _{c,i,DU} (°C)	T _{c,o,DU} (°C) T _{c,i,PMD} (°C)	T _{max,FPIM} (°C)	T _{max,DU} (°C)	T _{max,PMD} (°C)	q (W)
1	1.667E-05	39.6	471.6	178.1	212.5	280.1	65	69	68	17,100
2	3.333E-05	77.1	261.7	115.0	132.2	166.0	65	69	68	17,100
3	5.000E-05	112.6	191.8	94.0	105.5	128.0	65	69	68	17,100
4	6.667E-05	146.2	157.0	83.6	92.2	109.1	65	69	68	17,100
5	8.333E-05	177.9	136.1	77.4	84.3	97.8	65	69	68	17,100
6	1.000E-04	207.9	122.3	73.3	79.1	90.3	65	69	68	17,100
7	1.167E-04	236.2	112.4	70.5	75.4	85.0	65	69	68	17,100
8	1.333E-04	263.0	105.0	68.3	72.6	81.1	65	69	68	17,100
9	1.500E-04	288.2	99.3	66.7	70.5	78.1	65	69	68	17,100
10	1.667E-04	312.1	94.8	65.5	68.9	75.6	65	69	68	17,100
15	2.500E-04	412.4	81.5	61.9	64.2	68.7	65	69	68	17,100
20	3.333E-04	486.9	75.1	60.4	62.2	65.5	65	69	68	17,100
25	4.167E-04	542.3	71.5	59.8	61.2	63.9	65	69	68	17,100
30	5.000E-04	583.6	69.3	59.5	60.7	62.9	65	69	68	17,100
8.2	1.365E-04	267.9	103.8	68.0	72.2	80.4	65	69	68	17,100
13.8	2.292E-04	390.0	83.8	62.5	65.0	69.9	65	69	68	17,100
14.7	2.445E-04	406.6	82.1	62.0	64.4	69.0	65	69	68	17,100

Table 24: Low temperature cooling loop temperature calculations for loop order PMD, FPIM, DU, radiator air flow rate 3469 [CFM]

Flow (Q) (L·min ⁻¹)	Flow (Q) (m ³ ·s ⁻¹)	R (W·C ⁻¹)	Tc,o,PMD (°C) Tc,i,R (°C)	Tc,o,R (°C) Tc,i,FPIM (°C)	Tc,o,FPIM (°C) Tc,i,DU (°C)	Tc,o,DU (°C) Tc,i,PMD (°C)	Tmax, FPIM (°C)	Tmax, DU (°C)	Tmax, PMD (°C)	q (W)
1	1.667E-05	39.9	468.1	174.6	208.9	276.6	65	69	68	17,100
2	3.333E-05	78.1	259.0	112.2	129.4	163.2	65	69	68	17,100
3	5.000E-05	114.5	189.3	91.5	102.9	125.5	65	69	68	17,100
4	6.667E-05	149.3	154.5	81.2	89.8	106.7	65	69	68	17,100
5	8.333E-05	182.5	133.7	75.0	81.9	95.4	65	69	68	17,100
6	1.000E-04	214.2	119.8	70.9	76.7	87.9	65	69	68	17,100
7	1.167E-04	244.4	110.0	68.0	73.0	82.6	65	69	68	17,100
8	1.333E-04	273.2	102.6	65.9	70.2	78.6	65	69	68	17,100
9	1.500E-04	300.8	96.9	64.3	68.1	75.6	65	69	68	17,100
10	1.667E-04	327.0	92.3	63.0	66.4	73.1	65	69	68	17,100
15	2.500E-04	441.1	78.8	59.2	61.5	66.0	65	69	68	17,100
20	3.333E-04	531.2	72.2	57.5	59.2	62.6	65	69	68	17,100
25	4.167E-04	602.4	68.4	56.6	58.0	60.7	65	69	68	17,100
30	5.000E-04	659.1	65.9	56.2	57.3	59.6	65	69	68	17,100
7.0	1.170E-04	245.0	109.8	68.0	72.9	82.5	65	69	68	17,100
11.0	1.836E-04	352.4	88.5	61.9	65.0	71.1	65	69	68	17,100
12.4	2.062E-04	384.5	84.5	60.8	63.5	69.0	65	69	68	17,100

Table 25: Low temperature cooling loop temperature calculations for loop order PMD, FPIM, DU, radiator air flow rate 4857 [CFM]

Flow (Q) (L·min ⁻¹)	Flow (Q) (m ³ ·s ⁻¹)	R (W·C ⁻¹)	Tc,o,PMD (°C) Tc,i,R (°C)	Tc,o,R (°C) Tc,i,FPIM (°C)	Tc,o,FPIM (°C) Tc,i,DU (°C)	Tc,o,DU (°C) Tc,i,PMD (°C)	Tmax, FPIM (°C)	Tmax, DU (°C)	Tmax, PMD (°C)	q (W)
1	1.667E-05	39.9	468.4	175.4	209.7	277.2	65	69	68	17,100
2	3.333E-05	78.2	258.6	112.1	129.3	163.0	65	69	68	17,100
3	5.000E-05	115.0	188.7	91.1	102.5	125.0	65	69	68	17,100
4	6.667E-05	150.2	153.8	80.6	89.2	106.0	65	69	68	17,100
5	8.333E-05	184.0	132.9	74.3	81.2	94.7	65	69	68	17,100
6	1.000E-04	216.5	119.0	70.2	75.9	87.1	65	69	68	17,100
7	1.167E-04	247.6	109.1	67.2	72.1	81.8	65	69	68	17,100
8	1.333E-04	277.4	101.6	65.0	69.3	77.7	65	69	68	17,100
9	1.500E-04	306.0	95.9	63.3	67.1	74.6	65	69	68	17,100
10	1.667E-04	333.4	91.3	62.0	65.4	72.2	65	69	68	17,100
15	2.500E-04	454.5	77.6	58.1	60.4	64.9	65	69	68	17,100
20	3.333E-04	552.7	70.9	56.3	58.0	61.4	65	69	68	17,100
25	4.167E-04	632.4	67.0	55.3	56.7	59.4	65	69	68	17,100
30	5.000E-04	697.4	64.5	54.8	55.9	58.1	65	69	68	17,100
6.7	1.117E-04	238.5	111.7	67.9	73.1	83.1	65	69	68	17,100
10.3	1.713E-04	340.8	90.2	61.6	65.0	71.5	65	69	68	17,100
11.7	1.946E-04	376.9	85.4	60.2	63.2	69.0	65	69	68	17,100

Table 26: Low temperature cooling loop temperature calculations for loop order PMD, FPIM, DU, radiator air flow rate 6245 [CFM]

Flow (Q) (L·min ⁻¹)	Flow (Q) (m ³ ·s ⁻¹)	R (W·C ⁻¹)	T _{c,o,PMD} (°C) T _{c,i,R} (°C)	T _{c,o,R} (°C) T _{c,i,FPIM} (°C)	T _{c,o,FPIM} (°C) T _{c,i,DU} (°C)	T _{c,o,DU} (°C) T _{c,i,PMD} (°C)	T _{max,FPIM} (°C)	T _{max,DU} (°C)	T _{max,PMD} (°C)	q (W)
1	1.667E-05	40.0	467.8	174.3	208.7	276.3	65	69	68	17,100
2	3.333E-05	78.4	258.0	111.3	128.4	162.3	65	69	68	17,100
3	5.000E-05	115.5	188.1	90.3	101.7	124.3	65	69	68	17,100
4	6.667E-05	151.1	153.2	79.8	88.4	105.3	65	69	68	17,100
5	8.333E-05	185.3	132.3	73.6	80.4	94.0	65	69	68	17,100
6	1.000E-04	218.3	118.3	69.4	75.1	86.4	65	69	68	17,100
7	1.167E-04	250.0	108.4	66.5	71.4	81.0	65	69	68	17,100
8	1.333E-04	280.5	101.0	64.3	68.6	77.0	65	69	68	17,100
9	1.500E-04	309.9	95.2	62.6	66.4	73.9	65	69	68	17,100
10	1.667E-04	338.1	90.6	61.2	64.7	71.4	65	69	68	17,100
15	2.500E-04	463.9	76.9	57.3	59.6	64.1	65	69	68	17,100
20	3.333E-04	567.6	70.1	55.5	57.2	60.6	65	69	68	17,100
25	4.167E-04	653.2	66.2	54.4	55.8	58.5	65	69	68	17,100
30	5.000E-04	724.1	63.6	53.8	55.0	57.2	65	69	68	17,100
6.4	1.074E-04	232.5	113.5	68.0	73.3	83.8	65	69	68	17,100
9.8	1.631E-04	332.2	91.5	61.5	65.0	71.9	65	69	68	17,100
11.2	1.871E-04	371.3	86.1	59.9	63.0	69.0	65	69	68	17,100

Appendix D

EcoCAR Powertrain High Voltage Schematic

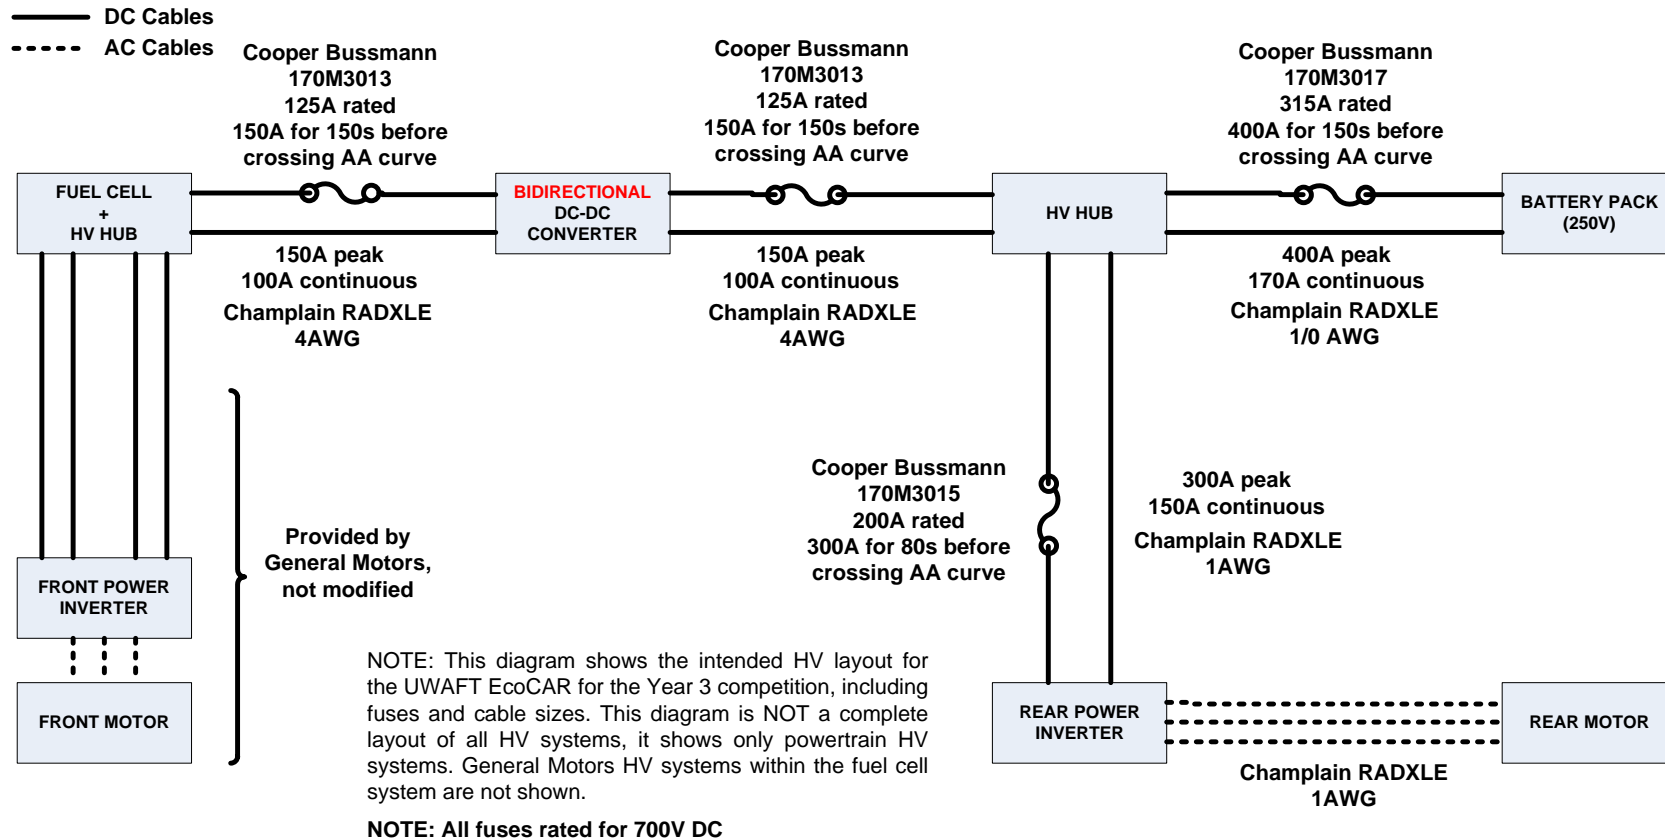


Figure 55: EcoCAR powertrain high voltage schematic

Appendix E

Power Control Algorithm Top-Level Overview

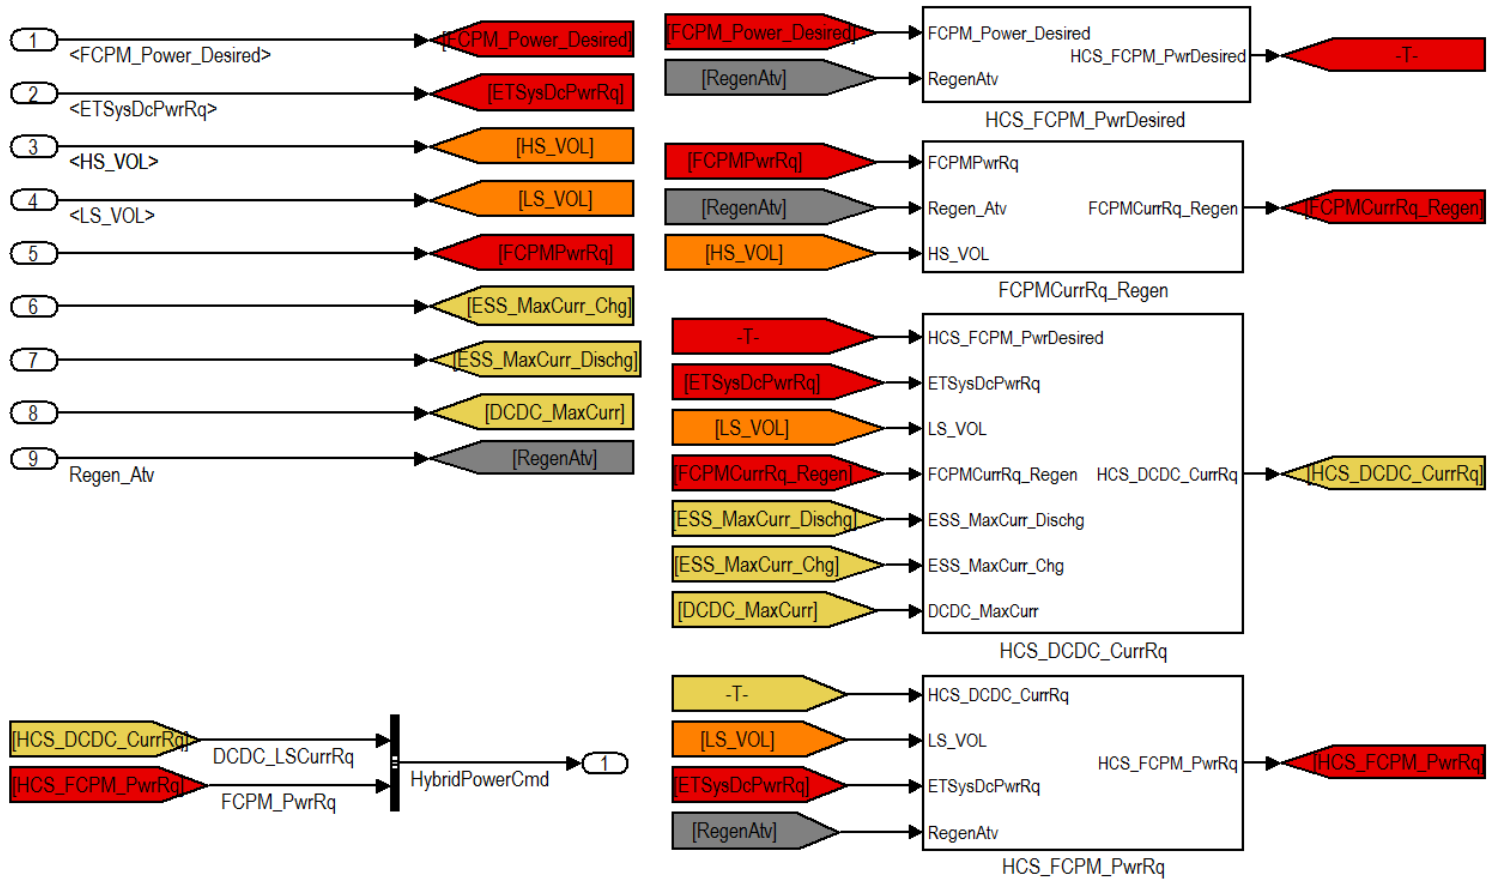


Figure 56: Power control algorithm top-level view

Appendix F

DC-DC Component Control Block for Driving Mode

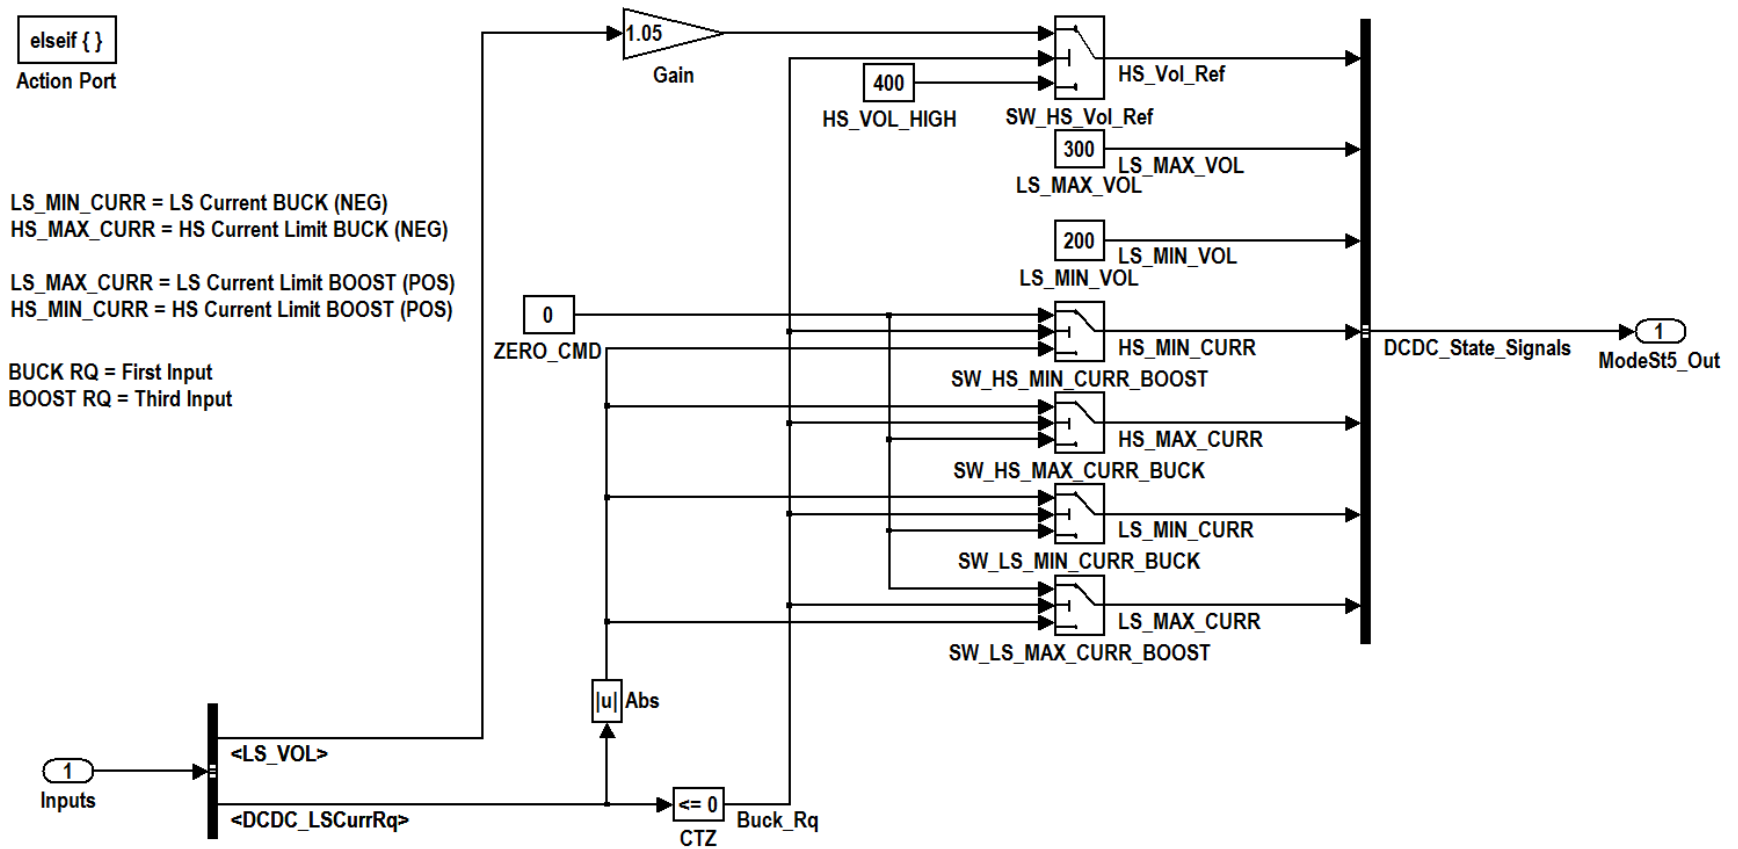


Figure 57: DC-DC component control block for "driving" mode

Appendix G

FMEA for Accelerator Pedal Sensor

Table 27: FMEA of the accelerator pedal sensor

COMPONENT	FAILURE TYPE	FAILURE MODE	EFFECT OF FAILURE	CAUSE OF FAILURE	DESIGN CONTROLS PREVENTION	DESIGN CONTROLS DETECTION	RPN
Accelerator Pedal Sensor	Interface - Physical - Open	Power open circuit	Zero signal	Detached connector Poor connector pin crimp Broken wire	Proper crimping process / tooling Avoid routing wire(s) near hot or sharp objects Conduit or loom to protect wire(s)	Define lower limit of valid range to be above zero signal	
	Interface - Physical - Open	Ground open circuit	Zero signal	Detached connector Poor connector pin crimp Broken wire	Proper crimping process / tooling Avoid routing wire(s) near hot or sharp objects Conduit or loom to protect wire(s)	Define lower limit of valid range to be above zero signal	
	Interface - Physical - Open	Signal open circuit	Zero signal	Detached connector Poor connector pin crimp Broken wire	Proper crimping process / tooling Avoid routing wire(s) near hot or sharp objects Conduit or loom to protect wire(s)	Define lower limit of valid range to be above zero signal	
	Interface - Physical - Short	Signal short to power	Saturated signal Source brown-out (p)	Broken wires Incorrect wire hookup	Avoid routing wire(s) near hot or sharp objects Conduit or loom to protect wire(s)	Define upper limit of valid range to be below zero signal Supply redundant sensor from independent power source	
	Interface - Physical - Short	Signal short to ground	Zero signal	Broken wires Sensor internal damage Incorrect wire hookup	Avoid routing wire(s) near hot or sharp objects Conduit or loom to protect wire(s)	Define lower limit of valid range to be above zero signal	
	Interface - Physical - Short	Power short to ground	Zero signal Source brown-out	Broken wires Sensor internal damage Incorrect wire hookup	Avoid routing wire(s) near hot or sharp objects Conduit or loom to protect wire(s)	Define lower limit of valid range to be above zero signal Supply redundant sensor from independent power source	
	Interface - Electrical - Attenuation	Signal attenuation	Erroneous signal, unknown value	Signal path resistance too high	Proper crimping process / tooling Select appropriate conductor gauge	Compare signal to redundant (backup) sensor	
	Interface - Electrical - Interference	Signal interference	Erroneous signal, unknown value	Electro-magnetic interference Ground point voltage drift Power supply transient	Avoid routing near high current conductors Avoid ground loops Power supply with sufficient load capacity Use appropriate power line filtering elements	Compare signal to redundant (backup) sensor Signal gradient range checks	
Sensor - Physical - Offset	Sensor mechanism stuck at non-zero position		Object behind pedal Mechanical friction, sticking Mounting bracket shifted	Design for conditions			
Sensor - Electrical - Offset	Sensor output drifts from zero position		Sensor internal damage	Design for conditions	Compare signal to redundant (backup) sensor		