

Optimizing Hybrid Vehicles: Battery Pack Design, Energy Management, and Collaborative Learning

by

Dylan Charles Erb

Submitted to the Department of Mechanical Engineering
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2016

© Dylan Charles Erb, MMXVI. All rights reserved.

The author hereby grants to MIT permission to reproduce and to
distribute publicly paper and electronic copies of this thesis document
in whole or in part in any medium now known or hereafter created.

Author
Department of Mechanical Engineering
May 18, 2016

Certified by.....
Sanjay E. Sarma
Professor of Mechanical Engineering
Thesis Supervisor

Accepted by
Rohan Abeyaratne
Chairman, Department Committee on Graduate Theses

Optimizing Hybrid Vehicles: Battery Pack Design, Energy Management, and Collaborative Learning

by

Dylan Charles Erb

Submitted to the Department of Mechanical Engineering
on May 18, 2016, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

Hybrid vehicles are designed to operate efficiently using data from within the bounds of that single vehicle. With limited access to outside information, hybrid control systems can react to external stimuli but cannot prepare ahead of time, which limits the overall efficiency. This thesis looks beyond the scope of traditional automotive systems to increase efficiency and boost performance of hybrid vehicles while decreasing cost and maximizing drivability.

I cover three main subject areas which are all focused on hybrid vehicle optimization: battery pack design for low cost thermal management, energy management of blended battery packs for weight minimization, and collaborative learning for improved energy management. The outputs of this thesis include insights into the minimization of air-cooling costs in battery packs, a better understanding of the blended battery pack design space, and in-vehicle applications of collaborative learning.

First, I have advanced the understanding of how cell-level properties affect the overall costs of air-cooling in battery packs and have shown how to minimize those costs. Through both simulation and analytical studies, I illustrated that cell size and geometry can be selected to balance the tradeoff between average and delta temperature constraints. At the intersection of these two constraints, a cell size can be found that results in the minimum cost of air-cooling for the parameters of that application. While the simulation results take significant time to compute, the proposed analytical method can be used to narrow down a design space in a fraction of the time through direct calculation.

Second, I have shown how blended battery packs can be sized for a given application and have identified what applications will benefit most from tribridization. Using integer programming for a lower bound, a dynamic programming based method for minimum feasibility, and a peak shaving heuristic as a baseline, I carried out a large study. In this experiment, I quantified how much room for improvement there is for a variety of hybrid applications through the pack-level combination of different energy storage technologies. Based on the results, I extracted rules of thumb for understanding when blended battery packs can improve over uniform packs.

Finally, I have demonstrated how collaborative learning can benefit energy management for hybrid vehicles. I illustrated how the drivability and fuel efficiency of current micro hybrids suffer from the automatic starting and stopping of the engine for relatively short idles. By defining a traffic model and comparing that to real-world datasets, I identified parameters that are most useful in predicting idle time of vehicles in traffic. Using these predictors in combination with machine learning classifiers, I quantified the potential for a Smart/Stop system that harnesses collaborative learning to improve over current automatic engine start/stop performance. The result is a system that can significantly increase drivability while maintaining or even improving upon baseline fuel economy.

Thesis Supervisor: Sanjay E. Sarma

Title: Professor of Mechanical Engineering

Acknowledgments*

*The word “acknowledgements” can’t even hold a candle to the amount of appreciation I have for all the wonderful people who have helped me get where I am today. While I know I can’t possibly capture my gratitude on this page, I hope that these words will at least show that I tried half as well as I should like, but certainly not half as well as half of you deserve.

Of course, the first thank you note goes to Sanjay. You’ve taught me so many lessons in my time here in Building 35, and I have to say that while I anticipated learning a great deal about engineering from you, I now see that I got far more than I bargained for. Sure, the lessons in engineering are great, but your best teachings come from the way you live your life and how you treat the people who pass through it. I’m truly grateful to have been one of them.

Next, I would be remiss if I did not extend my most sincere thanks to Professor Wallace and Professor Frey. Your feedback and support throughout this whole process have been phenomenal. Thank you so much for serving on my committee and for showing such interest in my research.

To my friends and labmates in Building 35, thank you for making every day at work seem like an episode on a sitcom. Thanks especially to Josh for being my automotive enabler, to Isaac for being my resident rabbi, to Sergio for being mi hermano, to Stephen for keeping me honest, to Sumeet for being my oracle, to Rahul for always having a pun ready, to Pranay for putting up with me, to Partha for chasing your dreams, to Yongbin for always laughing, to Alex for tolerating my Spanish, to Laura for all the help, and to Jason for all the bean town.

Thank you to Boston Power and JLR for sponsoring my research. Thanks to Eric Carlson for all the guidance and long conversations about battery packs. Your feedback was awesome.

To my high school physics and electronics teachers, Mr. Curry and Mac, thank you for challenging me to keep building and breaking things. You made high school an aspiring engineer’s paradise.

Thanks to my friends Dieter Brommer, Joey Paquette, Paul Schorfheide, Josh Eickmeier, Adrienne Bowman, Michael Murphy, Roberto Espaillat, Alex Keck and their families for everything.

To my parents and sister, I can't thank you enough for loving me and putting up with me for all these years. You showed me how to be happy from a young age, and I can never thank you enough for that. You taught me how to fix things and cook great food and always appreciate the little things in life. What more can I say? I will never forget where I came from.

Lastly, and most importantly, I would like to thank my best friend and future wife, Nicole. I don't know how you make me laugh so much, but I do know that you somehow manage to brighten every day of my life. Your unconditional support and unwavering adorability constantly amaze me. I can't wait to spend the rest of my life with you.

Contents

1	Introduction	19
1.1	A New Automotive Revolution	20
1.1.1	Automotive Electrification	21
1.1.2	Connected Cars and Vehicle Autonomy	23
1.2	Thesis Contributions	24
1.3	Thesis Outline	26
2	From EVs to Hybrids to Tribrids	27
2.1	Battery Packs	28
2.1.1	Cells	28
2.1.2	Modules	29
2.1.3	Cooling System	29
2.1.4	Battery Management	29
2.1.5	Bus Bars, Fuses, and Contactors	30
2.2	Electric Vehicles	30
2.3	Hybrid Electric Vehicles	31
2.3.1	Full Hybrid Electric Vehicles	32
2.3.2	Plug-in Hybrid Electric Vehicles	32
2.4	Tribrid Electric Vehicles	33
2.4.1	Blended Battery Packs	34
3	Battery Pack Design for Low Cost Thermal Management	37
3.1	Thermal Management in Air-Cooled Battery Packs	38

3.2	Scalable Air-Cooled Battery Packs	39
3.2.1	Types of Lithium-ion Cells	39
3.2.2	Scalable Cylindrical Battery Packs	39
3.2.3	Scalable Rectangular Battery Packs	42
3.2.4	Cost Modeling	45
3.3	Effects of Cell Size in Cylindrical Battery Packs	46
3.3.1	Design Constraints	46
3.3.2	Input Parameters	48
3.3.3	Cylindrical Simulation Results	48
3.3.4	Overall Cylindrical Results	53
3.3.5	Cylindrical Cell Sizing Conclusions	57
3.4	Effects of Geometry in Air-Cooled Battery Packs	58
3.4.1	Design Constraints	58
3.4.2	Input Parameters	59
3.4.3	Geometry Simulation Results	59
3.4.4	Geometry Comparison Conclusions	69
3.5	Analytical Model Based on Tube Bank Theory	70
3.5.1	Inlet Turbulence	71
3.5.2	Analytical Formulation for Cylindrical Packs	71
3.5.3	Comparison of Analytical and Simulated Results for Cylindrical Packs	77
3.6	Conclusion	83
4	Energy Management of Blended Battery Packs for Weight Mini- mization	85
4.1	Energy Storage Types	86
4.2	Blended Battery Packs for Tribridization	87
4.3	Drive Cycles as a Design Requirement	88
4.4	Automotive Use Cases	91
4.5	Uniform Battery Packs	92

4.5.1	Uniform Pack Modeling and Constraints	92
4.5.2	Uniform Battery Pack Sizing	96
4.6	Blended Battery Packs	100
4.6.1	Linear Integer Programming for Lower Bound	102
4.6.2	Dynamic Programming for Minimum Achievable Pack Size . .	111
4.6.3	Peak Shaving Strategy as a Benchmark	123
4.7	Study Results	125
4.7.1	Design of Experiment	126
4.7.2	Hybrid Electric Vehicles	127
4.7.3	Plug-in Hybrid Electric Vehicles	130
4.7.4	Electric Vehicles	132
4.8	Conclusion	134
5	Collaborative Learning for Improved Hybrid Energy Management	137
5.1	Connected Vehicles and Collaborative Learning	137
5.1.1	Connected Vehicle Technology	138
5.1.2	Collaborative Learning	138
5.2	Micro Hybrids and Automatic Engine Start/Stop	139
5.2.1	Micro Hybrid Structure	140
5.2.2	Estimated Fuel Economy Gains	142
5.2.3	Problems with Automatic Engine Start/Stop	143
5.3	JLR Data Set	147
5.3.1	JLR Data Statistics	147
5.3.2	JLR Data Set Limitations	153
5.4	Traffic Simulation Data Set	154
5.4.1	Traffic Modeling with the Intelligent Driver Model	154
5.4.2	Traffic Simulation Data Statistics	159
5.5	Cambridge Data Set	160
5.5.1	Experimental Method	161
5.5.2	Cambridge Data Statistics	162

5.6	Smart/Stop: Making AESS Seamless	165
5.6.1	Short Idle Classification for Smart/Stop	165
5.6.2	Confusion Matrix	165
5.6.3	Learning Feature Selection	166
5.6.4	Classification Methods	171
5.6.5	Cross-Validation	172
5.6.6	Idle Time Prediction	172
5.7	Conclusion	182
6	Conclusions	185
6.1	Insights on Hybrid Vehicle Optimization	185
6.1.1	Battery Pack Design	185
6.1.2	Energy Management and Blended Battery Packs	186
6.1.3	Collaborative Learning for Future Insights	187
6.2	Summary of Contributions	188
6.3	The Death of the Drive Cycle	189
6.4	Future Work	190
A	Result Summary Tables from Chapter 4	193

List of Figures

2-1	Cell Internal Resistance Model	28
2-2	EV Architecture	31
2-3	Parallel HEV Architecture	32
2-4	PHEV Architecture	33
2-5	TEV Architecture	34
2-6	Passively Connected Blended Battery Pack	34
2-7	Blended Battery Pack with Buffer Facing Load	35
2-8	Blended Battery Pack with Baseload Facing Load	36
2-9	Dual Active Parallel Blended Battery Pack	36
3-1	Scalable Module of Cylindrical Cells (Top View)	40
3-2	Scalable Module of Prismatic Cells (Top View)	42
3-3	Scalable Module of Rounded Prismatic Cells (Top View)	43
3-4	Fan Price vs. Static Pressure and CFM (Data and Planar Fit Contours)	46
3-5	Temp vs. CFM for n=2 ($A_{cell} = 2520 \text{ mm}^2, \varepsilon_{mod} = 278 \frac{Wh}{L}$)	49
3-6	Pressure vs. CFM for n=2 ($A_{cell} = 2520 \text{ mm}^2, \varepsilon_{mod} = 278 \frac{Wh}{L}$)	49
3-7	Temp vs. CFM for n=5 ($A_{cell} = 403 \text{ mm}^2, \varepsilon_{mod} = 278 \frac{Wh}{L}$)	51
3-8	Pressure vs. CFM for n=5 ($A_{cell} = 403 \text{ mm}^2, \varepsilon_{mod} = 278 \frac{Wh}{L}$)	51
3-9	Temp vs. CFM for n=8 ($A_{cell} = 158 \text{ mm}^2, \varepsilon_{mod} = 278 \frac{Wh}{L}$)	52
3-10	Pressure vs. CFM for n=8 ($A_{cell} = 158 \text{ mm}^2, \varepsilon_{mod} = 278 \frac{Wh}{L}$)	52
3-11	CFM and Pressure vs. Cell Cross-sectional Area, $\varepsilon_{mod} = 278 \frac{Wh}{L}$	54
3-12	Fan Cost vs. Cell Cross-sectional Area, $\varepsilon_{mod} = 278 \frac{Wh}{L}$	54
3-13	CFM and Pressure vs. Cell Cross-sectional Area, $\varepsilon_{mod} = 306 \frac{Wh}{L}$	55

3-14 Fan Cost vs. Cell Cross-sectional Area, $\varepsilon_{mod} = 306 \frac{Wh}{L}$	55
3-15 CFM and Pressure vs. Cell Cross-sectional Area, $\varepsilon_{mod} = 333 \frac{Wh}{L}$. . .	56
3-16 Fan Cost vs. Cell Cross-sectional Area, $\varepsilon_{mod} = 333 \frac{Wh}{L}$	56
3-17 Delta Temperature [C] Contour Plot for Cylindrical Cells	60
3-18 Average Temperature [C] Contour Plot for Cylindrical Cells	60
3-19 Pressure Drop [Pa] Contour Plot for Cylindrical Cells	61
3-20 CFM and Pressure vs. Cell Cross-sectional Area for Cylindrical Cells	61
3-21 Delta Temperature [C] Contour Plot for Rounded Prismatic Cells . .	63
3-22 Average Temperature [C] Contour Plot for Rounded Prismatic Cells .	63
3-23 Pressure Drop [Pa] Contour Plot for Rounded Prismatic Cells	64
3-24 CFM and Pressure vs. Cell Cross-sectional Area for Rounded Prismatic Cells	64
3-25 Delta Temperature [C] Contour Plot for Prismatic Cells	65
3-26 Average Temperature [C] Contour Plot for Prismatic Cells	66
3-27 Pressure Drop [Pa] Contour Plot for Prismatic Cells	66
3-28 CFM and Pressure vs. Cell Cross-sectional Area for Prismatic Cells .	67
3-29 Fan Cost vs. Cell Area for $\varepsilon_{mod} = 278 [\frac{Wh}{L}]$	68
3-30 Fan Cost vs. Cell Area for $\varepsilon_{mod} = 333 [\frac{Wh}{L}]$	69
3-31 Top View of Flow Paths for Scalable Cylindrical Battery Pack	72
3-32 Mesh Sensitivity for (a) Airflow and (b) Pressure Drop when $\varepsilon_{mod} = 278 [Wh L^{-1}]$	78
3-33 (a) Analytical and (b) Simulated Average Cell Temp [$^{\circ}C$] for Cylindrical Cells $\varepsilon_{mod} = 278 [Wh L^{-1}]$	79
3-34 (a) Analytical and (b) Simulated Delta Temp [$^{\circ}C$] for Cylindrical Cells $\varepsilon_{mod} = 278 [Wh L^{-1}]$	81
3-35 Simulated Temperature for (a) $N = 2$ and (b) $N = 5$ in [$^{\circ}C$] for 20[CFM] and $\varepsilon_{mod} = 278 [Wh L^{-1}]$	81
3-36 (a) Analytical and (b) Simulated Pressure Drop with Temperature Constraint Contour Overlay	82

3-37 (a) Airflow and Pressure Drop Comparison and (b) Estimated Fan Cost Comparison	83
4-1 Gravimetric Power Density vs. Gravimetric Energy Density for Surveyed Cells	89
4-2 Urban Dynamometer Drive Cycle (UDDS)	89
4-3 Heavy Duty Vehicle Urban Dynamometer Drive Cycle (HDVUDDS) .	90
4-4 Uniform Battery Pack Layout and Internal Resistance	93
4-5 Minimum Pack Mass for Uniform Solutions for Light Duty Vehicles on a UDDS Cycle, Numbers Indicate Energy Storage ID as Defined in Table 4.1	97
4-6 Minimum Pack Mass for Uniform Solutions for Heavy Duty Vehicles on a UDDS and HDVUDDS Cycle	99
4-7 Minimum Pack Mass for Integer Programming and Uniform Solutions for the Toyota Prius on a UDDS Cycle	105
4-8 Minimum Pack Mass for Integer Programming and Uniform Solutions for the Honda Insight on a UDDS Cycle	106
4-9 Minimum Pack Mass for Integer Programming and Uniform Solutions for the Hybrid Bus on a HDVUDDS Cycle	107
4-10 Minimum Pack Mass for Integer Programming and Uniform Solutions for the Hybrid Bus on a UDDS Cycle	108
4-11 Minimum Pack Mass for Integer Programming and Uniform Solutions for the Hybrid SUV on a UDDS Cycle	109
4-12 Minimum Pack Mass for Integer Programming and Uniform Solutions for the PHEV on a UDDS Cycle	110
4-13 Minimum Pack Mass for Integer Programming and Uniform Solutions for the Focus EV on a UDDS Cycle	111
4-14 Deterministic Dynamic Programming Based Algorithm for Blended Battery Pack Sizing	114
4-15 Example Portion of Power Cycle for UDDS Cycle	114

4-16	Discretized Example Portion of Power Cycle for UDDS Cycle	115
4-17	Discretized Example Portion of Power Cycle for UDDS Cycle Zoomed to Red Region from Figure 4-16	117
4-18	DDP Algorithm Substep Evaluation Technique for UDDS Power Cycle Example	117
4-19	DDP Algorithm Step N-1 Evaluation for UDDS Power Cycle Example	120
4-20	DDP Algorithm Step N-2 Evaluation for UDDS Power Cycle Example	121
4-21	DDP Algorithm Step N-3 Evaluation for UDDS Power Cycle Example	122
4-22	DDP Algorithm Best Path for UDDS Power Cycle Example	122
4-23	UDDS Power Cycle Example with Power Baseload Power Thresholds	123
4-24	Peak Shaving Baseload Power Curve for UDDS Power Cycle Example	124
4-25	Peak Shaving Buffer Power Curve for UDDS Power Cycle Example .	124
4-26	Minimum Pack Mass for DDP, Integer Programming, Peak Shaving, and Uniform Solutions for the Toyota Prius on a UDDS Cycle	125
4-27	Minimum Pack Mass for DDP, Integer Programming, Peak Shaving, and Uniform Solutions for the Hybrid Bus on a UDDS Cycle	126
4-28	Minimum Pack Mass for DDP, Integer Programming, Peak Shaving, and Uniform Solutions for the Hybrid SUV on a UDDS Cycle	127
4-29	Gravimetric Power Density vs. Gravimetric Energy Density for Sur- veyed Cells (black) and DDP Solutions (red) for 4 Consecutive Cycles	128
4-30	Least Massive Pack for DDP, Peak Shaving, and Uniform Solution Methods for 4 Consecutive Cycles	129
4-31	Minimum Pack Mass for DDP, Integer Programming, Peak Shaving, and Uniform Solutions for the PHEV on a UDDS Cycle	130
4-32	Gravimetric Power Density vs. Gravimetric Energy Density for Sur- veyed Cells (black) and DDP Solutions (red) for 7 Consecutive Cycles	131
4-33	Least Massive Pack for DDP, Peak Shaving, and Uniform Solution Methods for 7 Consecutive Cycles	132
4-34	Minimum Pack Mass for DDP, Integer Programming, Peak Shaving, and Uniform Solutions for the Focus EV on a UDDS Cycle	133

4-35	Gravimetric Power Density vs. Gravimetric Energy Density for Surveyed Cells (black) and DDP Solutions (red) for 10 Consecutive Cycles	134
4-36	Least Massive Pack for DDP, Peak Shaving, and Uniform Solution Methods for 10 Consecutive Cycles	135
5-1	Connected Vehicle Technologies: (a) V2V, (b) V2I, (c) V2C	138
5-2	Micro Hybrid Topology	140
5-3	Types of Starter Motors: (a) Ring Gear Starters and (b) Belt-Integrated Starter Generator	141
5-4	Idle Time for UDDS Cycle	143
5-5	Performance Comparison of BMW X1 with AESS enabled and Disabled (Gathered Using OBDII Data)	144
5-6	Critical Idle Duration Before AESS Benefits Fuel Economy	145
5-7	Histogram of Trip Durations from the JLR Data Set	148
5-8	Sample Trip from JLR Data Set	149
5-9	Histogram of Idle Durations from the JLR Data	150
5-10	Empirical Cumulative Distribution Function (CDF) of Idle Durations from the JLR Data	150
5-11	CDF of % Idle Time Due to Idles Less than the Threshold	152
5-12	Sample Trip with Driver Disablement from JLR Data Set	153
5-13	1-D Traffic Simulation with Traffic Light	156
5-14	Histogram of Idle Durations from the Simulated Data	160
5-15	Empirical CDF of Idle Durations from Simulated Data	161
5-16	Windshield Mounted LIDAR and Sonar Rig (1 of 2)	162
5-17	Chase Car with LIDAR and Sonar Mounted on Windshield Near Side Mirrors	163
5-18	Map View of Cambridge Idle Data, Greener Markers Correspond to Shorter Idles, Bluer Markers Correspond to Longer Idles	163
5-19	Histogram of Idle Durations from the Cambridge Data	164
5-20	Empirical CDF of Idle Durations from the Cambridge Data	164

5-21	(a) Speed of Vehicle $i-1$ at Start of Idle and (b) Acceleration of Vehicle $i-1$ at Start of Idle vs. Actual Idle Time	168
5-22	(a) Speed of Vehicle $i-2$ at Start of Idle and (b) Acceleration of Vehicle $i-2$ at Start of Idle vs. Actual Idle Time	169
5-23	Speed of Vehicle $i-1$ at Start of Idle vs. Actual Idle Time for Simulated and Cambridge Data	170
5-24	Classification Accuracy for (a) Short and (b) Long Idles Using Position Only with Simulated Data	173
5-25	Classification Accuracy for (a) Short and (b) Long Idles Using Speed of First Car in Front Only with Simulated Data	174
5-26	Classification Accuracy for (a) Short and (b) Long Idles Using Position and Speed of First Car in Front with Simulated Data	175
5-27	Classification Accuracy for (a) Short and (b) Long Idles Using Position and Speeds of First and Second Car in Front with Simulated Data	176
5-28	Classification Accuracy for (a) Short and (b) Long Idles Using Position Only with Cambridge Data	177
5-29	Classification Accuracy for (a) Short and (b) Long Idles Using Speed of First Car in Front Only with Cambridge Data	178
5-30	Classification Accuracy for (a) Short and (b) Long Idles Using Position and Speed of First Car in Front with Cambridge Data	178
5-31	Histogram of Idle Durations for Cambridge Data when AESS is Disabled and when Smart/Stop is Enabled	181

List of Tables

3.1	COMSOL Model Dimensions in Millimeters	48
4.1	Vehicle Energy Storage Information with Associated ID Numbers . .	87
4.2	Energy Storage Parameters for Each ID Number	88
4.3	Use Case Parameters	92
4.4	Minimum Uniform Pack Mass [kg] for Light Duty Vehicles on a UDDS Cycle	98
4.5	Minimum Uniform Pack Types for Light Duty Vehicles on a UDDS Cycle	99
4.6	Minimum Uniform Pack Mass [kg] for the Hybrid Bus on the UDDS and HDVUDDS Cycles	100
4.7	Minimum Uniform Pack Types for the Hybrid Bus on the UDDS and HDVUDDS Cycles	101
5.1	IDM Parameter Values	155
5.2	Traffic Simulation Parameter Values	156
5.3	Example Confusion Matrix	166
5.4	Smart/Stop Results Summary	180
A.1	Overall Results for DDP Solution for 4 Consecutive Cycles	193
A.2	Overall Results for PS Solution for 4 Consecutive Cycles	194
A.3	Overall Results for IP Lower Bound for 4 Consecutive Cycles	194
A.4	Overall Results for Uniform Solution for 4 Consecutive Cycles	194
A.5	Overall Results for DDP Solution for 7 Consecutive Cycles	194
A.6	Overall Results for PS Solution for 7 Consecutive Cycles	195

A.7 Overall Results for IP Lower Bound for 7 Consecutive Cycles	195
A.8 Overall Results for Uniform Solution for 7 Consecutive Cycles	195
A.9 Overall Results for DDP Solution for 10 Consecutive Cycles	195
A.10 Overall Results for PS Solution for 10 Consecutive Cycles	196
A.11 Overall Results for IP Lower Bound for 10 Consecutive Cycles	196
A.12 Overall Results for Uniform Solution for 10 Consecutive Cycles	196

Chapter 1

Introduction

Hybrid vehicles are currently at the forefront of automotive technology. With many extra degrees of freedom with respect to conventional internal combustion vehicles, they are simultaneously pushing the bounds of performance for high-end supercars (e.g. McLaren P1, Ferrari LaFerrari, Porsche 918 Spyder) while stretching a gallon of gas further than otherwise possible for passenger vehicles (e.g. Toyota Prius, Ford Fusion Hybrid, Chevy Volt). As fleet fuel economy requirements continue to rise and emissions regulations are tightened, hybrid and electric vehicles will necessarily be taken from cutting edge to run-of-the-mill.

With increasing market penetration of hybrid vehicles, they are expanding into market segments where price and fuel efficiency are paramount over the allure of being “green.” While the costs of such enabling technologies as lithium-ion batteries are continuing to decline, they are starting to approach the underlying material costs [1]. Automotive manufactures must therefore seek out new ways to decrease the costs of their hybrid components such as finding more efficient energy management strategies to reduce the necessary amount of material.

It may appear as though most of the low hanging fruit has been harvested in terms of hybrid vehicle optimization, but I will propose in this thesis that there are still many opportunities for improvement. I explore three main areas of hybrid vehicle optimization that touch on battery pack design, energy management, and collaborative learning.

Although each area of this thesis is distinct, it follows a natural progression that starts with the design of the battery pack – the very heart of hybrid and electric vehicles. Then I will show how the mass of battery packs can be minimized for certain applications by using blended battery packs – a system that we call the tribrid electric vehicle. These minimized packs are only possible through the use of energy management strategies that are enabled by connected vehicle technology. Finally, I will show that collaborative learning can be leveraged in connected vehicles to further optimize energy management through the study of a specific application - automatic engine start/stop systems in micro hybrid vehicles.

In the remainder of this chapter, I will give a brief overview of the new automotive revolution that has driven the majority of the recent advances in automotive technology. This includes automotive electrification, connected vehicles, and autonomy in vehicles. Then, I will highlight the major contributions of this thesis and outline the overarching structure of the chapters to come.

1.1 A New Automotive Revolution

Since the invention of the automobile in the late 19th century, cars have gone through several major revolutions as they have evolved into the vehicles we are familiar with today. They started out as luxury items that only the wealthiest people could afford. Then, Henry Ford brought the Model T to the masses through the introduction of the moving assembly line in the early 20th century.

As vehicles increased in popularity throughout the 20th century, automotive manufacturers continued to attract more customers through the addition of luxuries such as power windows and locks, more powerful headlights, and heated seats that demanded larger and larger electrical systems. Despite the fact that hybrid vehicles were invented very early on (circa 1900), they did not take off due to the relatively low cost of gasoline and lack of regulations on fuel economy and emissions [2].

With the construction of highways and interstates, high speed collisions brought on safety regulations that caused a revolution in automotive safety features such as

seat belts, airbags, and crumple zones. Despite all the changes, many aspects of automobiles have remained largely the same. The vast majority of vehicles still have a single internal combustion engine for propulsion, drive on four wheels, are controlled by a driver using a steering wheel and pedals, and employ control strategies that are purely reactionary.

Vehicles are likely to change more in the next decade than they have changed over the entirety of the last century. These changes are being brought on by a combination of factors. Regulations and public sentiment are driving up fuel efficiency and emissions standards beyond the point where internal combustion alone can keep up. At the same time, the falling cost and increased abundance of computational resources, cellular bandwidth, and high resolution sensors that make vehicles more aware of the environments around them are driving down the costs of active safety systems (e.g. collision-avoidance and lane-keeping) and semi-autonomous features (e.g. adaptive cruise control and parking assist). This new automotive revolution is two-fold: the gradual retirement of the internal combustion engine and the phasing out of the need for a human driver.

1.1.1 Automotive Electrification

Internal combustion engines have a multitude of features that have maintained their popularity in the automotive industry since the inception of the car and throughout the 20th century. The most advantageous of these features are the relatively low cost, high power density, and high energy density of the fossil fuels that power them [3]. In the 21st century, however, the underlying assumptions of the abundance of fossil fuels and lack of regulations on carbon emissions and fuel efficiency are becoming outdated.

When it comes to meeting upcoming fuel efficiency standards while still complying with emissions regulations, internal combustion engines suffer from a major disadvantage – they are irreversible. An engine can convert chemical energy stored in gasoline into mechanical work, but it cannot convert mechanical energy into gasoline. This complicates the recovery of energy that is naturally wasted in a car, of which there are

four main categories: waste heat, conventional braking, engine idling, and dynamic operation through inefficient load regions [4]-[5].

Motors, on the other hand, can convert mechanical work to electrical energy and store it chemically in batteries, so they are inherently bidirectional. Unfortunately, the batteries that power them are still relatively costly and not as energy or power dense as gasoline. All these reasons have led automotive manufacturers to turn to hybridization. An electric motor can easily capture energy through regenerative braking and convert it back into mechanical energy during loading regions where the engine would be least efficient. With a traction motor that can respond instantly to a torque command, the engine can be shut off during idling and started up after the vehicle has begun moving. The strengths of the electric motor address three out of the four main categories of energy wasted in engines (conventional braking, engine idling, and inefficient load regions) while still leveraging the high power and energy density of petroleum. The fourth area, waste heat, can also be recaptured through a variety of approaches including steam cycles, but this will be the last area to be tapped.

Although hybrid vehicles offer a great solution for improving efficiency with respect to conventional vehicles, they will not be the end solution due to further regulations of carbon emissions. Ultimately, it is much easier to power a full electric vehicle off of low-carbon fuels (e.g. solar and wind) than it is to produce low-carbon fuels that can be burned through combustion. If the trend continues, the electric systems in hybrid vehicles will become larger and larger as the costs of batteries decline, giving them longer all-electric range until EVs ultimately take over entirely. In the meantime, however, we must continue to advance the field of hybrid vehicle technology.

As price point is one of the main barriers standing in the way of mass adoption of electrified vehicles, and the high voltage battery makes up a large portion of the underlying costs [6], automotive manufacturers must continue to seek out new ways to optimize the designs of these packs. Since battery cells are the largest contributor to pack cost [7], designers may become fixated on optimizing the cell without fully considering the impacts that cell will have on the cost of the overall pack. To truly minimize the pack costs, the cost of the cell must be considered in conjunction with

the effects of that particular cell choice on the overall pack.

1.1.2 Connected Cars and Vehicle Autonomy

Until recently, vehicles have only been able to acquire the state of their environment from the sensors on board that vehicle. These sensors give the control systems the information they use to reliably control a vehicle, but they offer very little insight about what is coming in the near future. In this way, current vehicles are almost completely reactionary.

In order to support active safety features such as parking distance, lane-keeping, and collision-avoidance, automotive manufacturers are fitting high-end vehicles with sonar and RADAR sensors that extend the perceptive range of vehicles further and further away from the vehicle itself. These give the vehicle information about what is coming in the next few seconds and help the driver with simple reactive tasks but do not reveal anything about what will happen over the next few minutes or hours.

For further safety and traffic flow benefits, research has been carried out on connected vehicle technologies such as vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) communications. While they appear to offer great benefits for applications such as collision-avoidance [8] and vehicle platooning [9], they have not been widely adopted due to lack of regulation and the inextricable link between adoption and the utility of the technology.

Vehicle-to-cloud technology (V2C), on the other hand, is being adopted by automotive manufacturers who wish to carry out over-the-air updates, large-scale telematics, and vehicle usage information. This trend is likely to continue as there are obvious advantages to being able to update firmware in the owner's garage rather than instituting a large recall so that a dealer can flash the vehicle software manually.

With the flow of vehicle data onto the cloud comes the question of how else all of this information might be useful. Aggregation of large datasets enables pattern recognition and machine learning algorithms that could potentially allow vehicles to learn and adapt for the purpose of improving energy efficiency, drivability, user experience, and safety among others. This type of application will require massive

amounts of data and large investments before the bulk of gains can be realized. To justify the costs, the potential benefits thereof need to be understood ahead of time.

Even a connected vehicle with complete knowledge about the world around it can only attempt to predict how a driver will behave, and this can only be done with some degree of certainty. For example, a human driver may become more aggressive based on his or her mood (i.e. accelerating through a yellow light instead of slowing down), and this would be extremely difficult to predict. As semi-autonomous features become standard, the natural next step is fully autonomous vehicles where the future can much more easily be predicted because it will be defined by algorithms. Computer programs will always behave according to the predefined rules they are programmed to follow, and so predicting them is much more straightforward. Until the time when all vehicles are fully autonomous, there will still be uncertainty due to the human element.

1.2 Thesis Contributions

In this thesis, I further the field of hybrid vehicle technology through contributions in three main areas. These contributions are summarized in the following list:

- I have advanced the understanding of how cell-level properties affect the overall costs of air-cooling in battery packs and how to minimize those costs.
 - I define and utilize scalable air-cooled battery packs for the isolation of the effects of different cell features on thermal management characteristics.
 - Through simulation, I show that there is an optimal cell size for a given battery pack that balances the tradeoff between average temperature and delta temperature across a pack.
 - Through simulation, I show that rounded prismatic cells have an advantage over cylindrical and prismatic cells when it comes to air-cooling cost minimization.

- I define and utilize an analytical heat transfer formulation that verifies the conclusions drawn from simulation and provides a simpler means of determining the cell size for minimal cost cooling at early stages in hybrid vehicle design.
- I have advanced the understanding of how blended battery packs can be sized for a given application and what applications will benefit most from tribridization.
 - I define and utilize an integer programming method for determining the lower bound of blended battery pack mass for a given application.
 - I define and utilize a dynamic programming-based method for simultaneously calculating the minimum mass blended battery pack and control strategy thereof.
 - I carry out an experiment using the two optimization methods in addition to a peak shaving heuristic to compare the relative masses of the best blended solution versus the best uniform solution
 - Based on the results of the experiment, I extract rules of thumb for determining if an application can benefit from tribridization.
- I have advanced the understanding of how collaborative learning can benefit energy management for hybrid vehicles.
 - I define and utilize a traffic simulation to understand what parameters are most useful in predicting vehicle idle time.
 - I design and utilize a data acquisition system for gathering vehicle idle data in the real-world.
 - I analyze idle data from a large data set, and quantify the effectiveness and identify problems with stop/start systems in practice.
 - I quantify the performance of machine learning techniques on the simulated and real-world data for the purpose of a Smart/Stop system that increases drivability while maintaining fuel economy.

1.3 Thesis Outline

The remainder of this thesis is structured as follows. Chapter 2 provides an overview of different conventional hybrid and electric vehicle topologies, introduces the tribrid electric vehicle, and discusses the different blended battery pack topologies that enable it. Chapter 3 covers battery pack design for the minimization of the cost of air-cooled thermal management. I show first with simulation and then with an analytical formulation that these costs can be minimized by selecting the appropriate cell size for the chosen application. In Chapter 4, I go into depth on how blended battery packs can be sized without determining a predefined power split strategy, and I use this method to carry out a large experiment that shows the potential for weight reduction with respect to uniform battery packs for a variety of applications. At the end of Chapter 4, I point out that the energy management strategy that minimizes the size of a blended battery pack is not always practically realizable without future knowledge of the drive cycle, which highlights the need for predictive algorithms in hybrid vehicles. In Chapter 5, I quantify the potential benefits of using these types of predictive algorithms to improve the automatic engine start/stop systems in micro hybrid vehicles. I show that the drivability can be significantly improved without sacrificing fuel economy gains. Finally, I conclude this thesis in Chapter 6 with a discussion of the overall insights of this thesis and the future direction of hybrid vehicle optimization.

Chapter 2

From EVs to Hybrids to Tribbrids

Before we dig into the specifics of optimizing particular systems within hybrid vehicles, it is important to understand the entire spectrum of vehicle electrification from conventional vehicles all the way to full electric vehicles. Even modern day internal combustion vehicles can be considered electrified to a lesser degree.

Since an engine cannot output appreciable torque until it is running at hundreds of RPM, it also cannot start itself. The latest direct injection engines can sometimes be made to start themselves but cannot self-start reliably across a wide range of temperatures [10]. Because of this, automotive manufacturers have turned to electric starter motors to get engines going since the late 1800's. The starter motor supplements the engine torque enough at low speeds for the engine to continue to turn over on its own. Although this system is taken for granted now, it is credited with enabling mass adoption of internal combustion vehicles and represents the most basic form of vehicle electrification.

Prior to examining the different types of electrified vehicles in further depth, we will discuss the fundamental element of vehicle electrification that is responsible for storing energy – the battery pack.




Figure 2-1: Cell Internal Resistance Model

2.1 Battery Packs

For automotive applications, a single battery cell does not have a high enough power or energy capacity. This necessitates putting multiple cells into arrays that multiply the voltage and current capability by the number of cells in series and parallel. To limit the current to reasonable levels (in order to minimize conductive losses), tractive automotive electrical systems usually range from between 48-400 VDC depending on how much power is required [11]. As most automotive cells are around 3 to 4 VDC nominal, this translates to tens or hundreds of cells in series. How much current (A) and capacity (kWh) are required at this voltage will determine how many parallel cells are required. These series and parallel arrays of cells and all the components that support them are called battery packs. In the following sections, we will introduce the main components of battery packs. Sizing of uniform battery packs will be covered in more detail in Chapter 4.

2.1.1 Cells

Battery cells store electrical energy in the form of chemical potential and can be modeled simply as an ideal voltage source in series with an internal resistance, as shown in 2-1. Cells come in many different form factors, sizes, and chemistries. Each one has its own advantages. Because of their relatively high energy density, power

density, affordability, and improvements in safety, lithium-ion cells have become the industry standard for all forms of electrified vehicles [12]. Ultra-capacitors are also becoming more popular for applications where extreme power density and cycle life is required but energy density is not important.

2.1.2 Modules

Since cells are relatively small, they are usually arranged into modules before being packaged into the final battery pack enclosure. These modules serve a variety of purposes. They mechanically secure the cells, provide pathways for coolant flow, house fusible links, hold bus bars, compress cells when needed, simplify shipping and assembly, and define the general mechanical architecture of the battery pack.

2.1.3 Cooling System

Lithium-ion cells must be maintained within a certain temperature range in order to remain efficient and prevent any damage to the cells. As the cell model in Figure 2-1 suggests, the internal resistance of battery cells makes them inherently inefficient, and so any current flowing through the cell produces heat (i^2R_i losses) which must be removed from the battery pack. The amount of heat generated depends on the internal resistance of the cell and the current demanded by the application. For some applications, natural conduction and convection of heat out of the battery pack is enough to keep the cells within their recommended temperature range. For others, active cooling systems such as forced air-cooling or liquid cooling must be utilized.

2.1.4 Battery Management

In addition to a usable temperature range, lithium-ion cells must be maintained within acceptable voltage, current, and power limits [13]. Due to differences in capacity and internal resistance, the cells within a pack also tend to become unbalanced. Since the usable energy in a battery pack is limited by the cell with the lowest state-of-charge (SOC), they must be balanced periodically to ensure that all the available

energy in the battery pack can be used [14]. The battery management system (BMS) ensures that all these cell-level constraints are adhered to at all times. As SOC cannot be measured directly, the BMS also computes an estimated level, usually based on coulomb counting and the occasional calibration with open circuit voltage [15]. Additionally, the BMS usually tracks the state-of-health (SOH) of the battery pack in order to estimate how much life the pack has left. Battery management is an enormous task, and one of the main challenges in safely utilizing lithium-ion cells.

2.1.5 Bus Bars, Fuses, and Contactors

There are many other components in a battery pack in addition to cells, cooling systems, and battery management. The list of important items includes bus bars, fuses, and contactors. Bus bars and wires electrically connect the cells to each other and must be sized to carry the appropriate level of current. At the very least, a pack level fuse will be included to prevent damage from a short circuit. When parallel connections are utilized, fusible links are included to prevent propagation of failure from an internal cell short circuit. The main pack contactor disconnects the high voltage contacts of the battery pack from all external loads. This gives the battery management system the ability to remove the battery pack from the circuit when the key is turned off or in the event of a fault condition.

2.2 Electric Vehicles

The most electrified, and perhaps the simplest, topology is the full electric vehicle or the EV. There are many different ways to build an EV, but we will focus on the simple example of a rear wheel drive configuration with a fixed gear ratio. Figure 2-2 illustrates a simplified block diagram of how this might be laid out in a vehicle. The only source of energy is the high voltage battery pack, which is usually on the order of 300-400 VDC and will be located in a protected region of the vehicle.

When the vehicle is running, the high voltage battery pack supplies power to the motor controller which inverts the DC to three phase AC and delivers it to the




Figure 2-2: EV Architecture

traction motor according to the torque demanded by the driver. That traction motor delivers the mechanical torque through fixed gear ratio to a differential that transfers the torque to the half shafts which turn the wheels. The battery can be recharged from the electrical grid with an onboard or off-board charger.

The main advantage of an EV is also the main disadvantage – it doesn't have an internal combustion engine. This allows it to operate very efficiently and eliminate all of the systems that go along with an engine (gas tank, exhaust, large coolant system, etc), but also means that all of the energy must be stored in lithium-ion batteries, which are on the order of 10 times less energy dense than gasoline. Because of this, perhaps the largest challenge with EVs is limited range.

2.3 Hybrid Electric Vehicles

The category of hybrid electric vehicles covers a great deal of different topologies. Hybrids attempt to harness the efficiency gains of higher levels of electrification while still benefitting from the energy density of gasoline to provide practically unlimited range. They range from micro hybrids, to mild hybrids, to full hybrids, to plug-in hybrids. Micro hybrids and mild hybrids will be covered in greater detail in Chapter 5, so we will discuss full and plug-in hybrids in this section.




Figure 2-3: Parallel HEV Architecture

2.3.1 Full Hybrid Electric Vehicles

Full hybrid electric vehicles (HEVs) are produced in three main types: series, parallel, and series-parallel [16]. A series hybrid is essentially just an EV with an engine that is utilized like a generator to charge the batteries while driving. The only connection between the engine and the wheels is electrical, and the maximum power output to the wheels is equal to the power output of the traction motor.

A parallel hybrid (pictured in Figure 2-3) couples the engine and traction motor mechanically through a transmission so that both can simultaneously send power (in parallel) to the wheels. As with an EV, these traction motors usually use three-phase electric machines that are controlled using variable frequency inverters.

A series-parallel HEV (e.g. Toyota Prius) allows some combination of series and parallel modes usually with two different electric machines coupled to the engine through a planetary gear set or some other flexible arrangement. This is the most flexible of the full hybrid vehicles, but also the most complicated.

2.3.2 Plug-in Hybrid Electric Vehicles

A plug-hybrid electric vehicle (PHEV) is nearly identical to an HEV with two major differences: the vehicle can be charged from the grid like an EV, and it usually has a larger battery pack to accept more charge [17]. An example PHEV topology is




Figure 2-4: PHEV Architecture

shown in Figure 2-4. Just like an HEV, there are series, parallel, and series-parallel configurations.

In order to utilize the extra energy available in the battery pack, PHEVs operate in two different general control modes: charge depletion and charge sustaining modes. These modes follow directly from the main goal of a PHEV, which is to act like a pure EV (charge depleting mode) until there is insufficient charge remaining in the battery to maintain this all-electric mode. At this point, the vehicle switches from pseudo-EV to a pseudo-HEV and maintains the SOC of the battery at some target level (charge sustaining mode).

2.4 Tribrid Electric Vehicles

The next leap is to hybridize the vehicle once again (which we call tribridization), but this time at the battery pack level. This tribrid electric vehicle (TEV) has two different energy storage packs (collectively called a blended battery pack) that operate together to deliver the same output as a normal (uniform) battery pack [18]. The purpose of tribridization is to harness a baseload with relatively high energy density and a buffer with relatively high power density to ultimately reduce the overall mass, size, or cost of the high voltage system. Other benefits can include increased life of the battery pack and higher system efficiency depending on the cycle. One example




Figure 2-5: TEV Architecture




Figure 2-6: Passively Connected Blended Battery Pack

of a TEV topology is shown in Figure 2-5. There are even more potential TEV architectures than there are for HEVs and PHEVs combined. This is due to the many different topologies of blended battery packs. Although we will not go into great detail for each, we will outline a several different blended topologies in the section below.

2.4.1 Blended Battery Packs

The simplest type of blended battery pack is accomplished through a passive parallel connection, so it requires no power electronics [19]-[20], which means that the power split between the baseload and buffer is determined only by the impedance mismatch




Figure 2-7: Blended Battery Pack with Buffer Facing Load

between the two [21]-[23]. A block diagram showing this type of connection is depicted in Figure 2-6. Linking the DC voltage of the two sources means that very little of the SOC range of the buffer can be utilized. This topology can also suffer from relatively high leakage currents when the buffer is an ultracapacitor [24].

The next option for blended battery packs is to decouple the voltages of the buffer and baseload through the use of a single additional power electronic converter as shown in Figure 2-7. When the buffer is an ultracapacitor, this topology is known as ultracapacitor facing load, or UCFL. Placing the baseload behind the DC/DC converter means that all the energy delivered to or from the baseload is subject to the efficiency loss from the converter (typically between 94-95% efficient) [24]. On the other hand the energy traveling to and from the buffer does not suffer any additional losses. Since the buffer tends to have a less stable SOC, this configuration also results in a less stable DC bus voltage [25] - [26]. If the bulk of the energy is being sourced from the baseload, this configuration may not be the best solution in terms of efficiency.

The third blended battery pack topology is shown in Figure 2-8. Just like the previous configuration, this topology uses one additional power electronic converter but places it between the buffer and the load. When a battery is utilized for the baseload, this is known as battery facing load (BFL). Placing the baseload across the load stabilizes the DC bus while allowing the voltage (and SOC) of the buffer to vary




Figure 2-8: Blended Battery Pack with Baseload Facing Load




Figure 2-9: Dual Active Parallel Blended Battery Pack

freely, which translates to better utilization capability [27]-[28]. The disadvantage here is that the DC/DC converter must be sized to handle the full power of the buffer, which may be relatively high, and any energy that goes into and then out of the buffer suffers twice from the efficiency loss thereof [29].

Putting the BFL and UCFL topologies together yields the final and most complex configuration. This type of blended battery pack is called the dual active parallel (DAP) topology and is shown in Figure 2-9. Despite offering the greatest flexibility, as the DC bus voltage can be held at an arbitrary level, this configuration suffers from the highest converter losses [30] - [31].

Chapter 3

Battery Pack Design for Low Cost Thermal Management

In this chapter, I demonstrate how cell-level properties such as size and form factor have large impacts on the minimum possible cost of pack-level thermal management for air-cooled battery packs. In order to do this, I first introduce the idea of a scalable air-cooled battery packs where cell size and geometry can be manipulated while keeping other important attributes constant (like module energy density and outer dimensions). This allows for a range of cell sizes and shapes to be compared fairly in terms of cost of thermal management.

I present several simulation studies using scalable battery packs that highlight the tradeoff between the average and delta temperature constraints. For each of three different form factors, I show that there is a cell size that minimizes the cost of air-cooling. Next, I define an analytical (as opposed to numerical) heat transfer model based on tube bank theory that allows for the quick calculation of pack-level thermal management properties for scalable battery packs. Finally, I use this analytical model to confirm and validate the simulation results for cylindrical cells. This analytical model provides a much faster and scalable method than simulation for early battery pack design studies.

3.1 Thermal Management in Air-Cooled Battery Packs

When designing automotive battery packs, cells are selected based on a plethora of factors including, among others, energy density, power density, cycle life, and extreme temperature performance [40] - [44]. Capacity is chosen based on pack level requirements including peak C rate, vehicle range for electric operation (and the resulting required capacity), battery-life concerns, etc [45] - [48]. In packs that are sufficiently small (e.g. hybrids), the cell capacity may simply be chosen as the pack capacity, which results in a single string of series cells, which is known as a 1P pack (for one parallel). The number of series cells is indicated in a similar manner with an S, so a pack with 72 cells in series and 1 in parallel would be a 72S1P pack.

Cell selection must occur early on in the design process when the design of the enclosure that will contain and cool the cells is far from finalized, if started at all. Because of this, it is difficult to predict the pack-level implications of one cell over another. For all these reasons, the actual physical size of the cells is more likely to be a constraint than a variable in the battery pack design process.

When packs become sufficiently large (e.g. EVs), a new degree of freedom emerges. The question becomes whether it is better to choose larger form factor cells (fewer parallel cells) or to connect many smaller cells in parallel. At a glance, it may seem obvious that given the option, the number of cells should be minimized for simplicity's sake. The true answer, unfortunately, is not so simple. In reality, choosing a 1P over a 2P or 3P architecture can have far-reaching effects on pack level costs, including material costs, such as cell interconnects; fusing; crimp connections; thermal management hardware (fans, pumps, blowers); and manufacturing costs such as tab welding, assembly costs, and so on [1], [49].

While the overall pack cost function is composed of numerous variables, many of which are functions of cell size, this chapter seeks to isolate the effects of cell size, cell geometry, inlet turbulence, and cell alignment on pack level thermal management. Furthermore, the study is limited to air-cooled packs because of the stark differences between liquid and air-cooled battery packs, and so each deserves special attention.

3.2 Scalable Air-Cooled Battery Packs

3.2.1 Types of Lithium-ion Cells

Lithium-ion cells come in all different geometries and sizes. Many factors contribute to the decision of what battery cell geometry will best serve an application. For applications with small numbers of cells in a tightly constrained space, the most suitable geometry is one that efficiently fills the available space envelope. For example, cylindrical cells fit well into a cylindrical flashlight casing, and a flat polymer cell is perfect for a low profile cell phone. In large packs, on the other hand, packing efficiency becomes less of a differentiating factor [50] - [52], and other factors must be considered. Due to economies of scale, some geometries/sizes are available more cheaply than others, especially for relatively low volume players (Tesla Motors early on) [40]. Additionally, some cell geometries (such as pouch cells) require external compression to prevent excessive deformation over life [53]. In the case of air-cooled battery packs, the relative performance of different cell shapes is not well understood. To compound the problem, each cell geometry is available in many different capacities, and selecting the best size is difficult [54].

3.2.2 Scalable Cylindrical Battery Packs

Before any analysis can take place, a scalable pack topology must be defined that allows for a fair comparison between different cell sizes for a given battery module. Here we will examine cylindrical cells in square packs (Figure 3-1), although the process is extensible to other geometries.

This scalable geometry attempts to hold important design constraints constant while primarily varying cell size and module energy density. Cell energy densities, module dimensions, total cell volume, cell height, and total heat generation are held constant, while the number of cells is varied (and along with it, the diameter of the cells). This assumes that the volume of cell casing materials is small with respect to active materials.




Figure 3-1: Scalable Module of Cylindrical Cells (Top View)

As input parameters, the model accepts cell energy density (ε_{cell}), total module energy (E_{mod}), cell height (H), total heat output (\dot{Q}_{tot}), and outer gap (G_{out}). The variables of concern are number of cells per side (n) and desired module energy density (ε_{mod}).

With these inputs defined, the important derived parameters are as follows. The volume of the module (V) can be found based on the total module energy and module energy density.

$$V = \frac{E_{mod}}{\varepsilon_{mod}} \quad (3.1)$$

Since the module is square, the side length (L) is

$$L = \sqrt{\frac{V}{H}} = \sqrt{\frac{E_{mod}}{H\epsilon_{mod}}} \quad (3.2)$$

Similarly, the total cell cross-sectional area (A_{tot}) is

$$A_{tot} = \frac{E_{mod}}{H\epsilon_{cell}} \quad (3.3)$$

and the area of a single cell (A_{cell}) becomes

$$A_{cell} = \frac{A_{tot}}{n^2} = \frac{E_{mod}}{n^2 H \epsilon_{cell}} \quad (3.4)$$

which yields the cell diameter (D)

$$D = \sqrt{\frac{4A_{cell}}{\pi}} = 2\sqrt{\frac{E_{mod}}{\pi n^2 H \epsilon_{cell}}} \quad (3.5)$$

The spaces between centers (S_T) and (S_L) are then

$$S_T = S_L = \frac{L - D}{n - 1} \quad (3.6)$$

and the block side (B_{side}) and block height (B_{height}) are

$$B_{side} = L + 2G_{out} \quad (3.7)$$

$$B_{height} = H + 2G_{out} \quad (3.8)$$

and finally, the heat output of a single cell (\dot{Q}_{cell}) is

$$\dot{Q}_{cell} = \frac{\dot{Q}_{tot}}{n^2} \quad (3.9)$$

With these parameters known, the pack geometry is defined and can be scaled to any n -by- n array of cells.




Figure 3-2: Scalable Module of Prismatic Cells (Top View)

3.2.3 Scalable Rectangular Battery Packs

To better understand how these different geometries affect the overall, pack-level thermal management properties, a scalable battery module from 3.2.2 is extended to other popular geometries. In order to best compare the modules containing different shaped cells, cross-sectional area for each cell in a given n -by- n arrangement is held constant. Gap spacing is recalculated to maintain the same module-level energy density. While an n -by- n arrangement of cylindrical cells results in a square pack, prismatic and rounded prismatic cells yield a rectangular pack, and this further complicates the calculations.

As input parameters, the model accepts cell energy density (ε_{cell}), total module




Figure 3-3: Scalable Module of Rounded Prismatic Cells (Top View)

energy (E_{mod}), cell height (H), total heat output (\dot{Q}_{tot}), and outer gap (G_{out}). The variables of concern are number of cells per side (n) and desired module energy density (ε_{mod}).

The rounded prismatic cell geometry shown in Figure 3-3 is essentially a cylindrical cell that is elongated along the longitudinal axis. In order to elongate the cell while maintaining the overall volume and height of the cell, the cell must shrink along the transverse axis. If we assume the cell has been elongated by some fraction (or multiplier α) of the new diameter, then the rounded prismatic area is

$$A_{RP} = \alpha D_{RP}^2 + \frac{\pi D_{RP}^2}{4} \quad (3.10)$$

Since we require that $A_{RP} = A_{CY}$ [54]

$$A_{cell} = \frac{E_{mod}}{n^2 H \varepsilon_{cell}} = D_{RP}^2 \left(\alpha + \frac{\pi}{4} \right); \quad (3.11)$$

$$so \ D_{RP} = \sqrt{\frac{E_{mod}}{n^2 H \varepsilon_{cell} \left(\alpha + \frac{\pi}{4} \right)}}. \quad (3.12)$$

Knowing that the volume target is

$$V = \frac{E_{mod}}{\varepsilon_{mod}} = H(D_{RP} + (n-1)S_T)(D_{RP} + (n-1)S_L) \quad (3.13)$$

and the longitudinal spacing is

$$S_L = S_T + \alpha D_{RP} \quad (3.14)$$

we can expand the expression of volume and solve for the space between the cells using a simple quadratic formula while taking the positive root.

$$\begin{aligned} S_T &= \frac{-b + \sqrt{b^2 - 4ac}}{2a}; \text{ where } \mathbf{a} = H(n-1)^2; \\ \mathbf{b} &= H(D_{RP}(2+\alpha)(n-1) + \alpha D_{RP}(n-1)^2) \\ \text{and } \mathbf{c} &= H(D_{RP}^2(1+\alpha n)) - \frac{E_{mod}}{\varepsilon_{mod}} \end{aligned} \quad (3.15)$$

In order to extend this general pack geometry to prismatic cells (Figure 3-2), the width can be assumed to be equal to the diameter of the rounded prismatic cell

$$w = D_{RP} \quad (3.16)$$

By maintaining that the cross-sectional areas are all equal

$$A_{CY} = A_{RP} = A_{PR} = \frac{E_{mod}}{(nH\varepsilon_{cell})} \quad (3.17)$$

we can find the length of the prismatic cell by division

$$l = \frac{A_{PR}}{w} = \sqrt{\frac{E_{mod}(\alpha + \frac{\pi}{4})}{n^2 H \epsilon_{cell}}} \quad (3.18)$$

In both cases, the spacing between cells is the same. With these parameters known, the pack geometry is fully defined and can be scaled to any n -by- n array of cells.

3.2.4 Cost Modeling

To assign some dollar cost of a fan with a particular airflow and pressure capability, we compiled prices from an online retailer (Digikey.com) and plotted them versus airflow and static pressure as shown in Figure 3-4. A planar function shows a relatively good fit and achieves an R-squared value of 0.68 in this case. Using this plane defined by the normal vector (-0.019, -0.169, 0.985) through the origin, we can take airflow and pressure as inputs and calculate an estimated dollar cost. Although the actual costs will change over time, this function can always be updated. This best fit plane yields an objective function defined by

$$C = \frac{0.019\dot{V} + 0.169\Delta P}{0.985}, \quad (3.19)$$

where C is [\$], \dot{V} is [CFM] and ΔP is [Pa].

In Figure 3-4, the contour lines represent the planar function, and the dots show the individual fan data points used to generate that planar function. This represents a more accurate cost estimation of blowers at low volumes than high volumes but should provide a reasonable metric for relative comparisons in this study. Without this combined metric, it would be difficult to differentiate between subtle discrepancies in required airflow and resulting pressure drop [55].




Figure 3-4: Fan Price vs. Static Pressure and CFM (Data and Planar Fit Contours)

3.3 Effects of Cell Size in Cylindrical Battery Packs

3.3.1 Design Constraints

To find the necessary airflow and resulting pressure drop for each pack geometry, we developed a COMSOL Multiphysics 4.4 model. The model uses the geometric equations derived above to form a scalable battery pack. Just as above, the model takes cell energy density, module dimensions, total cell volume, cell height, and total heat generation as inputs, while the number of cells per side (n) and module energy density (ε_{mod}) can be varied. Ambient air is forced through the front face of the pack with a specified inlet turbulent intensity (defined as the root-mean-square of the turbulent velocity fluctuations divided by the mean velocity) and exits through the back face (cylinders in cross flow). The side walls are taken to be thermal insulators with wall boundary conditions.

The three most important features from a thermal management standpoint are the average temperature (\bar{T}), delta temperature (ΔT), and pressure drop (ΔP). Average

temperature is one of the largest factors in rate of pack degradation, and excessive delta temperatures lead to ununiform aging of the pack. Since the pack capacity is limited by the weakest cell, large temperature gradients will lead to premature end-of-life. In order to keep the cells cool, we flow air across the pack at a sufficient volumetric flow rate, for which we must necessarily develop a pressure drop. We compute the average temperature by summing the temperature of each cell and dividing by the total number of cells

$$\bar{T} = \frac{1}{n^2} \sum_{i=1}^{n^2} \bar{T}_{cell_i}. \quad (3.20)$$

We find the delta temperature of the pack by subtracting the minimum cell temperature from the maximum cell temperature

$$\Delta T = \max(T_{cell}) - \min(T_{cell}). \quad (3.21)$$

The pressure drop of the pack is taken to be the average pressure at the front face of the pack minus the average pressure at the back face of the pack

$$\Delta P = \bar{P}_{front} - \bar{P}_{back}. \quad (3.22)$$

To illustrate the process, we selected a maximum allowable average temperature of $40^\circ C$ and ΔT of $5^\circ C$ since these values are commonly used in industry. Inlet air temperature was $27^\circ C$, and the inlet turbulent intensity was set to 0.1. For each flow rate in a range of inlet volumetric airflows from 20 to 300 cubic feet per minute (CFM), the resulting temperatures of the cells and the pressure drop from the front to the back were evaluated with the COMSOL model. The maximum of the two resulting flow rates (one for \bar{T} and one for ΔT) is taken as the requirement (*min(CFM) such that $\bar{T} \leq 40^\circ C$ and $\Delta T \leq 5^\circ C$*), and the resulting pressure drop is interpolated from the pressure plot. With the required airflow and pressure drop defined, the cost of the fan can be estimated using the fan cost function.

Table 3.1: COMSOL Model Dimensions in Millimeters

n	D	S_T, $\varepsilon_{mod} = 278$	S_T, $\varepsilon_{mod} = 306$	S_T, $\varepsilon_{mod} = 333$
2	56.6	74.4	68.2	63.1
3	37.7	46.6	43.6	41.0
4	28.3	34.2	32.2	30.5
5	22.6	27.1	25.5	24.3
6	18.9	22.4	21.2	20.2
7	16.2	19.1	18.1	17.2
8	14.1	16.7	15.8	15.1

3.3.2 Input Parameters

In this analysis, the model input parameters are cell energy density $\varepsilon_{cell} = 474 \text{ Wh/L}$, total module energy $E_{mod} = 309 \text{ Wh}$, cell height $H = 64.8 \text{ mm}$, total heat output $\dot{Q}_{tot} = 42 \text{ W}$, and outer gap $G_{out} = 9.8 \text{ mm}$. The ranges of variables are number of cells per side $n = [2 : 8]$ and desired module energy density $\varepsilon_{mod} = [278, 306, 333] \text{ Wh/L}$. These numbers were chosen to be within the practical limitations of current technology. The cell diameters and spaces between centers for each of the simulated configurations are shown in Table 3.1.

3.3.3 Cylindrical Simulation Results

Some representative samples of the COMSOL simulation outputs are shown in Figures 3-5 through 3-10. Each of these figures shows the lesser of the three simulated module energy densities ($\varepsilon_{mod} = 278 \text{ Wh/L}$). On each of the temperature plots, the critical flow rates at which the average and delta temperature constraints are met are highlighted with dotted vertical lines (black for ΔT and cyan for \bar{T}). The higher of the two critical flow rates becomes the driving constraint and is represented on the pressure plot as a dotted vertical line.

Figure 3-5 portrays the cell temperatures verses flow rate for the 2x2 array ($A_{cell} = 2520 \text{ mm}^2$ which corresponds to a relatively large diameter of about 56 mm and 4 cells total). As expected, the maximum, minimum, and average temperatures of the




Figure 3-5: Temp vs. CFM for $n=2$ ($A_{cell} = 2520 \text{ mm}^2, \epsilon_{mod} = 278 \frac{Wh}{L}$)




Figure 3-6: Pressure vs. CFM for $n=2$ ($A_{cell} = 2520 \text{ mm}^2, \epsilon_{mod} = 278 \frac{Wh}{L}$)

cells decrease as the flow rate is increased, but cooling is diminished as the surface temperature approaches the ambient air temperature of $27^\circ C$. It is also important to

note that the spread between the maximum and minimum temperatures is decreasing as the flow rate increases.

In this example, the flow rate required to maintain an average temperature of less than $40^{\circ}C$ is higher than the flow rate required to maintain a ΔT of less than $5^{\circ}C$. Therefore, the minimum flow rate that satisfies the constraints is driven by the average temperature requirement. As shown in Figure 3-6, the pressure drop across the pack increases with flow rate. Using the higher constraint, the required pressure is shown to be about $200 Pa$ in Figure 3-6.

Increasing the number of cells to a 5x5 array ($A_{cell} = 403 mm^2$ which is roughly $22 mm$ in diameter and 25 cells total) but holding all other parameters constant yields some interesting changes to the temperature verses flow rate curves as illustrated in Figure 3-7. As compared to the 2x2 array, the temperatures drop substantially for any given flow rate. The result is that the required CFM for both constraints decreases significantly. At this point, the airflow required to satisfy the ΔT requirement is only slightly less than the airflow for the average temperature requirement, but the average temperature requirement still defines the required flow rate.

Contrary to the change in the temperature curve, the pressure drop for a given flow rate increases to about double that of the 2x2 array (from a pressure of less than $900 Pa$ to almost $1800 Pa$) as illustrated in Figure 3-8. Despite this, the lower required flow rate more than compensates for the general pressure increase as the pressure drop corresponding to the critical flow rate decreases by about sixty percent to approximately $80 Pa$.

Further increasing the number of cells to an 8x8 array ($A_{cell} = 158 mm^2$ which is a diameter of about $14 mm$ and 64 cells total) continues to decrease the average temperature of the pack for each flow rate but causes an increase in ΔT as shown by the wide spread of the temperature curves in Figure 3-9. This causes the ΔT constraint to overtake the average temperature constraint as the driving flow rate requirement. The flow rate at which the average temperature constraint is satisfied may have decreased with respect to the 5x5 array, but because it is no longer driving, this has no effect on the required airflow.




Figure 3-7: Temp vs. CFM for $n=5$ ($A_{cell} = 403 \text{ mm}^2, \varepsilon_{mod} = 278 \frac{Wh}{L}$)




Figure 3-8: Pressure vs. CFM for $n=5$ ($A_{cell} = 403 \text{ mm}^2, \varepsilon_{mod} = 278 \frac{Wh}{L}$)

From Figure 3-10, it becomes clear that the pressure drop has increased slightly as compared to the 5x5 array. This small increase, in combination with the increased




Figure 3-9: Temp vs. CFM for $n=8$ ($A_{cell} = 158 \text{ mm}^2, \varepsilon_{mod} = 278 \frac{Wh}{L}$)




Figure 3-10: Pressure vs. CFM for $n=8$ ($A_{cell} = 158 \text{ mm}^2, \varepsilon_{mod} = 278 \frac{Wh}{L}$)

flow rate requirement due to the ΔT constraint, causes a large increase in required pressure over the 5x5 array. This transition between the two constraints highlights

how they are competing and suggests that some balance must be struck between them.

3.3.4 Overall Cylindrical Results

To reveal the bigger picture, a plot of airflow and pressure drop verses cell cross-sectional area is shown in Figure 3-11. This plot includes every array size from $n=2$ to $n=8$ for $\varepsilon_{mod} = 278 \frac{Wh}{L}$ and illustrates the formation of a clear minimum. This valley can be attributed to the two competing constraints; **average temperature** and ΔT . As the size of the cells decreases, the surface area to volume ratio increases, and the average heat transfer improves. This is captured in the right side of the plot as both the required airflow and pressure drop decrease as cell cross-sectional area decreases. The downward trend continues until the ΔT spikes for the smaller cells (as it did in Figure 3-9). The minimum for both the airflow and resulting pressure drop curves occurs at the critical point where the airflow required for average temperature is approximately equal to the airflow required for ΔT .

By evaluating each solution using the cost model described previously, a relative dollar amount can be assigned to each cell size as depicted in Figure 3-12. This plot accounts for both the flow rate and pressure drop, but since both curves have a similar shape and fan cost is more sensitive to pressure capability, the shape closely resembles the pressure drop curve. This said, it is interesting to see that the minimum cost blower is possible when a cell between an 18650 and a 26650 is selected. While the magnitude of the fan cost for both the 18650 and 26650 are similar in this case, the slope of the cost curve is much more drastic for the smaller cells than the larger cells.

Increasing the module energy density to $\varepsilon_{mod} = 306 \frac{Wh}{L}$ significantly increases the resulting pressure drop as shown in Figure 3-13. This is intuitive since increasing the module energy density decreases the gaps between the cells. The increase in pressure is especially pronounced for the smaller cells. The necessary airflow increases for the smaller cells and decreases for the larger cells. Both airflow and pressure drop still have clear minimums, but they have shifted slightly toward the right (as compared




Figure 3-11: CFM and Pressure vs. Cell Cross-sectional Area, $\epsilon_{mod} = 278 \frac{Wh}{L}$




Figure 3-12: Fan Cost vs. Cell Cross-sectional Area, $\epsilon_{mod} = 278 \frac{Wh}{L}$

to the lower energy density packs).

As shown in Figure 3-14, estimated fan cost for the $306 \frac{Wh}{L}$ pack has increased




Figure 3-13: CFM and Pressure vs. Cell Cross-sectional Area, $\epsilon_{mod} = 306 \frac{Wh}{L}$




Figure 3-14: Fan Cost vs. Cell Cross-sectional Area, $\epsilon_{mod} = 306 \frac{Wh}{L}$

across the board. The minimum cost fan now occurs for cells slightly larger than the 26650 form factor. Selecting an 18650 over a 26650 would increase the cost of the




Figure 3-15: CFM and Pressure vs. Cell Cross-sectional Area, $\epsilon_{mod} = 333 \frac{Wh}{L}$




Figure 3-16: Fan Cost vs. Cell Cross-sectional Area, $\epsilon_{mod} = 333 \frac{Wh}{L}$

blower by more than 50%. The fan cost for cells smaller than 500 mm^2 is even more sensitive to changes than in the lower energy density pack.

For the highest energy density packs ($\varepsilon_{mod} = 333 \frac{Wh}{L}$), the required pressure increases even further over the $306 \frac{Wh}{L}$ pack as illustrated in Figure 3-15. Once again, this makes sense because the cells are being squeezed closer together. This also causes a decrease in the required inlet volumetric airflow, since air going through smaller gaps accelerates faster and causes more effective convection. The minimums are still present, and have only shifted slightly with respect to the middle energy density packs.

The estimated fan cost curve once again increases across the board as shown in Figure 3-16, although the increase is less pronounced. The added cost of selecting an 18650 over a 26650 has now jumped to over double. Since the slope of the fan cost curve is even steeper, the cost sensitivity has become even greater for the smaller cells. A few millimeter change in diameter for smaller cells can cause double digit percent changes to fan cost.

3.3.5 Cylindrical Cell Sizing Conclusions

When this process is repeated for each pack from a 2x2 to an 8x8, several patterns become clear. Naturally, increasing flow rate decreases the average and delta temperature of the cells. At the same time, the pressure drop increases as flow rate increases. For larger cells, average temperature is the limiting factor and requires more airflow as shown in Figures 3-5 and 3-6. For small cells, ΔT is the limiting constraint as Figures 3-9 and 3-10 illustrate. The result is a cost function with a clear minimum at the cell size where the two constraints intersect.

As module energy density is increased, the pressure drop required for a given airflow also increases. The change in the airflow based on energy density depends on the cell size, but regardless, the increase in necessary pressure drop causes the cost of the fan to increase with module energy density. For instance, while only one cell size results in a cost of \$20 when $\varepsilon_{mod} = 306 \frac{Wh}{L}$, a whole range of cells yield costs of less than \$20 when module energy density is reduced to $278 \frac{Wh}{L}$.

This suggests that for air-cooled packs, cell size should be carefully considered as a driving factor in cooling cost. Utilizing a non-optimal cell size can easily double the

cost of the blower for a given energy density. This cost is particularly sensitive with small cells. While both the 18650 and 26650 cell sizes yield similar fan costs when $\varepsilon_{mod} = 278 \frac{Wh}{L}$, the sensitivity to change is much higher for the 18650, and thus a 26650 would be preferable.

Every battery pack is unique, and the optimal cell size for minimizing the cost of thermal management will depend heavily on the particular parameters of that battery pack. While there is no single battery size that will minimize the cost of cooling in *every* pack, these results show that there is a cell size that will minimize the cost of cooling for that *particular* battery pack. Considering the effects of cell size on thermal management early on in the design process can help to cut costs and maximize attainable energy density.

3.4 Effects of Geometry in Air-Cooled Battery Packs

3.4.1 Design Constraints

Based on the parameters above, models were created, and forced convection was simulated for each scenario with an inlet at the front of the module and an outlet at the back (cells in cross flow) [56]. A target average cell temperature (\bar{T}) of $40^\circ C$ and ΔT (where $\Delta T = \max(T_{cell}) - \min(T_{cell})$) of under $5^\circ C$ were selected as constraints. These values depend on the requirements of the particular application and cell chemistry [57] but were selected to be values typically used in industry. Inlet air temp was chosen to be $27^\circ C$. For each flow rate in a range of inlet volumetric airflows, the resulting temperatures of the cells and the pressure drop from the front to the back were evaluated using COMSOL with LiveLink for MATLAB, a multiphysics FEA simulation software package scripted through MATLAB. Relative fan cost was estimated using a planar fit based on CFM and pressure requirements as discussed in Section 3.2.4 [54].

3.4.2 Input Parameters

For this analysis, the model input parameters are cell energy density $\varepsilon_{cell} = 474 \left[\frac{Wh}{L}\right]$, total module energy $E_{mod} = 309 Wh$, cell height $H = 64.8 mm$, total heat output $\dot{Q}_{tot} = 42 W$, $\alpha = 0.96$, and outer gap $G_{out} = 9.8 mm$. The ranges of variables are number of cells per side $n = [2 : 8]$ and desired module energy density $\varepsilon_{mod} = [278, 333] \left[\frac{Wh}{L}\right]$.

3.4.3 Geometry Simulation Results

Figures 3-17 through 3-28 show simulation results for each cell geometry. Every geometry has a set of four figures that show the behavior of ΔT , average temperature, pressure drop, and required airflow with respect to the cross-sectional area of that cell shape. Contour plots are used for the temperature and pressure plots to display a large number of simulations in each plot. They also highlight large differences in performance between large cells and small cells for each geometry. The $\Delta T \leq 5^\circ$ and $\bar{T} \leq 40^\circ$ constraints are overlaid as contours on the plots. Every flow rate above both constraint contours is acceptable for the corresponding cross-sectional area. The required flow rate is the minimum without violating either constraint (*min(CFM) such that $\bar{T} \leq 40^\circ C$ and $\Delta T \leq 5^\circ \Delta C$*).

Cylindrical Cell Contour Plots

Figure 3-17 shows ΔT as a function of airflow and cell cross-sectional area for the full range of cylindrical cells in this study. As with each of the other ΔT plots, there is a hot spot at the bottom left corner. This corresponds to smaller cells at the lowest flow rates and, at over $22^\circ C$, also represents the highest ΔT out of all the geometries. Although it is roughly half the temperature, there is another hot spot at the bottom right corner, which corresponds to the larger cells.

The contour plot of average temperature verses cell cross-sectional area and flow rate for cylindrical cells is shown in Figure 3-18. The hottest average temperature is immediately evident in the bottom right corner of the plot, which represents larger




Figure 3-17: Delta Temperature [C] Contour Plot for Cylindrical Cells




Figure 3-18: Average Temperature [C] Contour Plot for Cylindrical Cells




Figure 3-19: Pressure Drop [Pa] Contour Plot for Cylindrical Cells




Figure 3-20: CFM and Pressure vs. Cell Cross-sectional Area for Cylindrical Cells

cells at low flow rates. The average temperature slopes down as cell cross-sectional area is reduced and as flow rate is increased. At over $90^{\circ}C$, the large cylindrical cells yield the highest average temperature in the study.

Figure 3-19 shows the contour plot of pressure drop verses cell size and airflow rate. The maximum pressure occurs for the smallest cells at the highest flow rates. This pressure decreases as cell cross-sectional area increases and flow rates drop. This represents the opposite trend as the plot of average temperature. At over $300 Pa$, the cylindrical cells have the lowest peak pressure drop out of all the geometries. As with each of the pressure plots, both the $\Delta T \leq 5^{\circ}$ and $\bar{T} \leq 40^{\circ}$ constraints are overlaid on top to illustrate how the final pressure curve is derived.

The final plot devoted to the cylindrical cells isolates the required airflow and resulting pressure drop verses cell size and is shown in Figure 3-20. The airflow curve (in blue) represents the minimum airflow where both constraints are satisfied. The cylindrical cells require the highest airflow in the study. The pressure drop curve (in green) depicts the pressure that must be developed across the pack to achieve the required airflow. Despite having the lowest pressure drop for a given airflow and cell size, the relatively high required airflows result in the second highest pressure curve.

Rounded Prismatic Cell Contour Plots

Figures 3-21 through 3-24 show the simulation results for rounded prismatic cells. Figure 3-21 illustrates ΔT as a function of airflow and cross-sectional area for these cells. As compared to the cylindrical cells, the rounded prismatic cells yield a substantially lower ΔT ($11^{\circ}C$ peak verses $22^{\circ}C$).

The average temperature plot shown in Figure 3-22 closely resembles that for cylindrical cells. Both have the same peak temperature location in the bottom right corner. However, the temperature range is shifted to lower temperatures for the rounded prismatic cell, i.e., from $30^{\circ}C - 90^{\circ}C$ down to $30^{\circ}C - 75^{\circ}C$. The $40^{\circ}C$ contour is also shifted to lower flow rates ($\sim 20 - 35 CFM$) for the rounded prismatic geometry.

Similar to the pressure plot for cylindrical cells, the pressure contour plot for




Figure 3-21: Delta Temperature [C] Contour Plot for Rounded Prismatic Cells




Figure 3-22: Average Temperature [C] Contour Plot for Rounded Prismatic Cells




Figure 3-23: Pressure Drop [Pa] Contour Plot for Rounded Prismatic Cells




Figure 3-24: CFM and Pressure vs. Cell Cross-sectional Area for Rounded Prismatic Cells




Figure 3-25: Delta Temperature [C] Contour Plot for Prismatic Cells

rounded prismatic cells illustrated in Figure 3-23 shows the highest pressures in the top left corner, which corresponds to smaller cells at high flow rates. While the trend is the same, the magnitude of the pressure drops has increased by about a factor of two (from over 300 Pa to over 600 Pa). The $\Delta T \leq 5^\circ$ and $\bar{T} \leq 40^\circ$ constraints, on the other hand, have both dropped with respect to the cylindrical cells.

Figure 3-24 shows the required airflow (in blue) and the resulting pressure drop (in green) for the rounded prismatic cells. Despite having a much larger pressure drop for a given flow rate and cell size, the required pressure curve is significantly lower than the same curve for the cylindrical cells (Figure 3-20). This can be attributed to the lower overall airflow rates required to maintain the temperature constraints.

Prismatic Cell Contour Plots

The simulation results for the prismatic cells are shown in Figures 3-25 through 3-28. The ΔT is depicted in Figure 3-25. From this plot it is clear that, in this study, prismatic cells perform better than cylindrical but worse than rounded prismatic cells




Figure 3-26: Average Temperature [C] Contour Plot for Prismatic Cells




Figure 3-27: Pressure Drop [Pa] Contour Plot for Prismatic Cells




Figure 3-28: CFM and Pressure vs. Cell Cross-sectional Area for Prismatic Cells

with respect to ΔT . Regardless, the same hot spot location can be seen for smaller cells at low flow rates in the bottom left corner. The hot spot for larger cells (bottom right corner) disappears for prismatic cells.

Figure 3-26 shows the average temperature verses cell size and airflow rate for prismatic cells. Once again, the trend of higher average temperatures for larger cells at low flow rates holds true. Relative to the cylindrical and rounded prismatic cells, this plot illustrates that prismatic cells require the lowest flow rates to stay within the average temperature constraint.

While the pressure contour plot for prismatic cells (Figure 3-27) looks almost identical to the same plot for rounded prismatic cells (Figure 3-23), the magnitude of the pressure drops has increased nearly threefold from over 600 Pa to over 1800 Pa . The $\Delta T \leq 5^\circ$ and $\bar{T} \leq 40^\circ$ contours overlaid on top of the pressure plot represent the lowest required airflows on average.

Figure 3-28 shows the summary of required airflow and pressure drop verses cell cross-sectional area for prismatic cells. Despite having the lowest average required




Figure 3-29: Fan Cost vs. Cell Area for $\varepsilon_{mod} = 278 \left[\frac{Wh}{L} \right]$

airflow, the pressure drop curve is the highest in the study. This can be attributed to the large pressure drop required to develop a particular airflow for the prismatic pack. Unlike the other two cell geometries, where the minimum pressures and airflows fall in between an 18650 and 26650 cell, the minimum pressure drop and airflow for prismatic cells rest with cells slightly larger in cross-sectional area than a 26650.

Overall Geometry Comparison Results

Using the planar fit method described in [54], the airflow and pressure drop curves for each geometry can be used to estimate the dollar cost of a blower or fan for any given pack. Figures 3-29 and 3-30 show these estimated costs for each geometry at $\varepsilon_{mod} = 278 \left[\frac{Wh}{L} \right]$ and $\varepsilon_{mod} = 333 \left[\frac{Wh}{L} \right]$ respectively.

These overall results illustrate the differences between the geometries with respect to cell cross-sectional area. A clear minimum cost can be obtained by choosing the appropriate cross-sectional area for each geometry. Depending on the energy density,




Figure 3-30: Fan Cost vs. Cell Area for $\varepsilon_{mod} = 333 \left[\frac{Wh}{L} \right]$

some geometries are better suited for air-cooling than others.

At $278 \left[\frac{Wh}{L} \right]$, the rounded prismatic cells perform substantially better than cylindrical and prismatic cells across the board. When the module energy density is increased to $333 \left[\frac{Wh}{L} \right]$, the gap between the rounded prismatic and cylindrical cells almost disappears. This is most likely because the spaces between the cells have become small enough that the lines of cylindrical cells become more streamlined, in which case the two geometries are not so different.

For each geometry, increasing the energy density from $278 \left[\frac{Wh}{L} \right]$ to $333 \left[\frac{Wh}{L} \right]$ causes a large increase in the cost of the blower (from 50–100%). It also causes the minimum cost cell size to shift to the right (toward larger cells) in each case.

3.4.4 Geometry Comparison Conclusions

The results show that small differences in cell shape and size can have large impacts on the necessary air flow and resulting pressure drop for a given module. For each cell

geometry and energy density studied, a minimum cost cell size was found. Straying from this cell size translates directly to additional cost of cooling. Furthermore, the differences between the cell geometries were shown to be large. Utilizing a non-optimal cell size or geometry can easily double or triple the cost of cooling. Due to the relative slope of the cost function on either side of the minimum, this cost is particularly sensitive for smaller cells. The results also illustrate that cooling costs can be reduced by increasing the space between cells, but for mobile applications, extra space comes at a premium.

During cell selection, both the geometry and shape should be carefully considered given the pack requirements. For air cooling, these results show that while cylindrical cells minimize pressure drop for a given flow rate and prismatic cells minimize the required flow rate, rounded prismatic cells minimize the pressure drop resulting from the required flow rate, which in turn minimizes the cost of cooling the pack.

3.5 Analytical Model Based on Tube Bank Theory

While understanding which cell can minimize the cost of air cooling for a given pack is important, the multiphysics simulation methods from the previous two sections require long computation times and many iterations to converge, making it relatively cumbersome as a design tool early on in the design process.

Here we consider an analytical method derived from well-established heat transfer equations empirically shown to predict the behavior of tube banks in crossflow [58]. While battery packs and tube banks are similar in many ways, there are several key differences that necessitate a number of changes to the model, which are outlined in Section 3.5.2. After the model is explained, the analytical results are compared side-by-side with the simulations in order to understand how well they line up and when and why they diverge.

The analytical method allows for very quick calculations that link the simulation results to thermal intuition and provide a means of external validation. The results indicate that our analytical method has the potential to serve as a useful tool when

pack designers consider the ramifications of one cell over another early on in the design process.

3.5.1 Inlet Turbulence

It should be noted that the simulations covered in Sections 3.3 and 3.4 utilized an inlet turbulent intensity of 0.1. The cylindrical simulations that are compared with the analytical results were run with an inlet turbulent intensity of 0.05. This boundary condition was updated in order to more closely parallel the conditions of the original experiments used to derive the tube bank heat transfer equations.

The analytical model accounts for turbulence within the Reynolds number dependent factors along with correction factors to adjust for developing turbulence as the air travels through the first few rows. These factors will be discussed further in Section 3.5.2.

3.5.2 Analytical Formulation for Cylindrical Packs

Our framework builds upon a foundation of empirically derived heat transfer correlations for aligned tube banks in crossflow [58]. While battery packs are similar to tube banks in many ways, there are several key differences. Tube banks are generally used in heat exchangers where the tubes are long and often filled with a phase change material that maintains a constant temperature. In battery packs, the battery cells (or tubes) cannot be considered to be infinitely long and are not held at a constant temperature. To recreate the simulated scenario, two parallel flow paths must be considered: through the bank, and around the outside as shown in Figure 3-31. Both paths must be considered because the gaps around the bank will differ in size from the gaps between the cells (especially as the inner gaps vary with the number of cells). For this reason, we incorporate the heat transfer through the tops and bottoms of the cells. In addition, each row of cells must be considered sequentially in order to find the temperature gradient across the pack, unlike tube banks, where the whole bank can be treated as a black box with an inlet temperature, an exit temperature, and an




Figure 3-31: Top View of Flow Paths for Scalable Cylindrical Battery Pack

average heat transfer coefficient.

Assuming the bank of battery cells exhibits a coefficient of drag c_d with respect to the air flowing around, the pressure drop can be calculated as

$$\Delta P_{out} = \frac{1}{2} c_d v_{out}^2 \rho, \quad (3.23)$$

where ρ is the density of air, and v_{out} is the velocity of air. This defines the pressure drop for the outer path by approximating the cell bank as a bluff body.

The resulting pressure drop [59] from traveling through the center of a tube bank is

$$\Delta P_{thru} = \frac{1}{2} N \chi \rho f v_{thru}^2 \left(\frac{S_T}{S_T - D} \right)^2, \quad (3.24)$$

where N is the number of rows and friction factor (f), and correction factor (χ), can be determined from the reference chart found in [58]. For the range of parameters in this study, these factors were taken to be constants.

Since the two flow paths are parallel, we can solve for the ratio of velocities that equates the pressure drops.

$$\Delta P_{out} = \frac{1}{2} c_d v_{out}^2 \rho = \frac{1}{2} N \chi \rho f v_{thru}^2 \left(\frac{S_T}{S_T - D} \right)^2 = \Delta P_{thru} \quad (3.25)$$

Eliminating and solving for $\frac{v_{out}}{v_{thru}}$ yields ϕ ,

$$\frac{v_{out}}{v_{thru}} = \left(\frac{S_T}{S_T - D} \right) \sqrt{\frac{N \chi f}{c_d}} = \phi. \quad (3.26)$$

By continuity, the total volumetric flow must be

$$\dot{V} = v_{out} A_{out} + v_{thru} A_{thru}, \quad (3.27)$$

where A_{out} and A_{thru} are the outside and inside flow areas.

Substituting in for v_{out} , v_{thru} can be determined

$$v_{thru} = \frac{\dot{V}}{A_{out} \phi + A_{thru}}. \quad (3.28)$$

So v_{out} is

$$v_{out} = \phi v_{thru} = \frac{\dot{V}}{A_{out} + \frac{A_{thru}}{\phi}}, \quad (3.29)$$

where $\phi = \left(\frac{S_T}{S_T - D} \right) \sqrt{\frac{N \chi f}{c_d}}$.

We also know that the Reynolds number for a tube bank is

$$Re_D = \frac{u_{thru} D}{\nu}, \quad (3.30)$$

where $u_{thru} = \frac{v_{thru}}{\pi} \int_0^\pi \left(\frac{S_T}{S_T - D \sin \theta} \right) d\theta$ for row 1, and $u_{thru} = v_{thru} \left(\frac{S_T}{S_T - D} \right)$ for rows 2 through N , and ν is kinematic viscosity [59].

For tube banks, we know that the average Nusselt number [60] for the surface of the tube is

$$\overline{Nu_D} = Pr^{0.36} fn(Re_D); \quad (3.31)$$

if $10^2 < Re_D < 2 \cdot 10^3$

$$fn(Re_D) = 0.52 Re_D^{0.5}, \quad (3.32)$$

if $10^3 < Re_D < 2 \cdot 10^5$

$$fn(Re_D) = 0.27 Re_D^{0.63}, \quad (3.33)$$

if $10^5 < Re_D$

$$fn(Re_D) = 0.033 Re_D^{0.8}. \quad (3.34)$$

This Nusselt number, $\overline{Nu_D}$, only applies to the inner rows of a tube bank because the heat transfer coefficient is smaller for the rows toward the front of the bank [60]. This is due to the generation of turbulence by the first few rows, after which the intensity of turbulence levels off. As the turbulence approaches the steady state value, the heat transfer correction factor approaches one. This correction factor (c_n) for each row (n), can be approximated as

$$\overline{Nu_n} = c_n \overline{Nu_D} \quad (3.35)$$

where

$$c_1 = 0.7$$

$$c_2 = 0.9$$

$$c_3 = 0.98$$

$$c_{n>3} = 1.$$

To find the Nusselt number for the top and bottom surfaces of the cells, we model

the flow over the flat surface as a modified flow over a flat plate [60].

Reynolds number can be calculated as

$$Re_x = \frac{v_{out}x}{\nu}, \quad (3.36)$$

where x is the distance from the inlet of the module, and ν is kinematic viscosity. For a given row, n , the Reynolds number is then

$$Re_n = \frac{v_{out}(nD + (S_T - D)(n - 1))}{\nu}. \quad (3.37)$$

If $Re_n < 5 \cdot 10^5$, then the Nusselt number for the top and bottom surfaces is

$$\overline{Nu}_t = 0.665Re_n^{0.68}Pr^{0.33}. \quad (3.38)$$

If $Re_n \geq 5 \cdot 10^5$, then

$$\overline{Nu}_t = 0.037Re_n^{0.8}Pr^{0.33}, \quad (3.39)$$

where $Pr = \frac{\nu}{\alpha}$.

The total heat flux for a given cell is the sum of all the cell surfaces

$$\dot{Q}_{tot} = 2\dot{Q}_t + \dot{Q}_n, \quad (3.40)$$

where \dot{Q}_t is the total heat flux through the top or bottom surface, and \dot{Q}_n is the heat flux through the battery tube surface.

From Newton's law of cooling we know that

$$\dot{Q} = \bar{h}A(T_\infty - T_s). \quad (3.41)$$

Assuming the entire surface of each cell maintains the average cell temperature, the total heat flux is

$$\dot{Q}_{cell} = 2\bar{h}_tA_t(T_\infty - \bar{T}_{cell}) + \bar{h}_nA_n(T_\infty - \bar{T}_{cell}), \quad (3.42)$$

and the ratio of heat flux through the flat surface versus the tube surface is $\frac{2\dot{Q}_t}{\dot{Q}_n} = \frac{\overline{Nu}_t D}{2\overline{Nu}_n H}$.

We can substitute in $A_t = \frac{\pi D^2}{4}$, $A_n = H\pi D$, and $\bar{h} = \frac{\overline{Nu}k}{D}$ to yield

$$\dot{Q}_{cell} = \left(2\frac{\overline{Nu}_t k \pi D^2}{D \cdot 4} + \frac{\overline{Nu}_n k}{D} H\pi D \right) (T_\infty - \bar{T}_{cell}), \quad (3.43)$$

where H is the height of the cell, and k is the thermal conductivity.

Simplifying and solving for T_{cell} , we can find

$$\bar{T}_{cell} = T_\infty - \frac{\dot{Q}_{cell}}{\pi k \left(\frac{\overline{Nu}_t D}{2} + \overline{Nu}_n H \right)}. \quad (3.44)$$

So the surface temperature for a single row can be determined where \dot{Q}_{cell} out of the cell is negative.

In order to determine the temperatures of the next $N - 1$ rows, the exit air temperature of each row must be found in sequence (since each exit temperature is the inlet temperature for the next row).

The change in temperature of the air can be found as

$$\Delta T_{air} = \frac{\dot{Q}_{cell}}{\dot{m}_{cell} c_p}, \quad (3.45)$$

where \dot{m}_{cell} is the mass flow rate of air heated by the cell, and c_p is the specific heat capacity of the air.

The mass flow rate of air heated by the cell can be estimated as

$$\dot{m}_{cell} = \eta \dot{V} \rho \frac{S_V S_H}{A_{tot}}, \quad (3.46)$$

where η is a discount factor between 0 and 1 (due to the shadowing effect [58]), $S_V = S_T$ for inner cells and $S_V = \frac{S_T + D}{2} + S_O$ for cells along the outer edge, S_H is the height of the module, and A_{tot} is the total flow area.

By definition,

$$A_{tot} = A_{out} + A_{thru}. \quad (3.47)$$

The inlet temperature for the next cell (in row n) is then

$$T_{\infty_n} = T_{\infty_{n-1}} + \Delta T_{air_{n-1}}, \quad (3.48)$$

and so the surface temperature of a cell in row n is

$$\bar{T}_{cell_n} = T_{\infty_{n-1}} + \Delta T_{air_{n-1}} - \frac{\dot{Q}_{cell}}{\pi k \left(\frac{Nu_t D}{2} + Nu_n H \right)}, \quad (3.49)$$

where T_{∞_0} is the ambient air temperature, and ΔT_{air_0} is 0.

The overall average cell temperature can be found by adding up all the cell row temperatures and dividing by the number of rows

$$\bar{T}_{pack} = \frac{1}{N} \sum_{n=1}^N \bar{T}_{cell_n}. \quad (3.50)$$

The overall pack ΔT_{pack} can be found as the maximum cell temperature (usually row 1 or N) minus the minimum cell temperature (usually row 2)

$$\Delta T_{pack} = \bar{T}_{cell_{max}} - \bar{T}_{cell_{min}}. \quad (3.51)$$

The overall pressure drop of the pack is

$$\Delta P_{pack} = \Delta P_{out} = \Delta P_{thru}. \quad (3.52)$$

3.5.3 Comparison of Analytical and Simulated Results for Cylindrical Packs

Using the analytical formulation in the above section, identical parameters as [54] (listed in Table 3.1) were used to make a side-by-side comparison of the results. The comparison shows a relatively good match with a few exceptions that are explained by the assumptions of the analytical model.




Figure 3-32: Mesh Sensitivity for (a) Airflow and (b) Pressure Drop when $\varepsilon_{mod} = 278 [Wh L^{-1}]$

Mesh Sensitivity Analysis

For each of the simulated packs, the meshing granularity was selected to provide a convergent result while not exceeding the computational capacity of the resources available for this study. In a similar manner to [61], the mesh count is increased until the marginal improvement of the result becomes negligible. The meshes were calibrated for the fluid dynamics solver using COMSOL’s predefined element size parameters that range from extremely coarse to extremely fine. As illustrated in Figure 3-32, both airflow and pressure drop were shown to have converged for every pack from $N = 2$ to $N = 8$.

Average Temperature

Starting with the average cell temperature plots in Figure 3-33, the analytical and simulated results show a very good resemblance. The normalized root-mean-square (NRMS) error between the two contour plots is 2.05%, where NRMS is the root-mean-square error divided by the difference between the maximum and minimum average temperature values. Both exhibit an increasing average temperature trend as the cell cross-sectional area is increased and decreasing temperature as flow is increased. The hottest average temperature occurs with the largest cells and the lowest flow in the




Figure 3-33: (a) Analytical and (b) Simulated Average Cell Temp [°C] for Cylindrical Cells $\varepsilon_{mod} = 278 [Wh L^{-1}]$

bottom right corner as should be expected. The 40°C average temperature constraint follows an almost identical path in both cases. Additionally, both plots show a very similar peak temperature of 84°C and 86°C.

Delta Temperature

The delta temperature plots shown in Figure 3-34 tell a very different story than the average temperature plots. The left sides of the two plots show very similar trends but the right sides are significantly different. This results in a normalized root-mean-square error of 8.32% between the two. In both plots, the highest delta temperature occurs with the smallest cells at the lowest flows and decreases as flow is increased as well as when the cell’s size is increased. Once the cell size becomes large enough (greater than 1000 mm²), the two plots begin to differ. The analytical model predicts that the delta temperature should increase gradually for the larger cells while the simulation predicts that it should increase abruptly for larger cells at lower flow rates.

This divergence can be accounted for in the analytical assumptions that revolve around the cells being a part of a “tube bank.” As the number of cells decreases beyond $N = 4$, the pack in question behaves less and less like a tube bank when

it comes to ΔT . This can be traced back to the driving force behind the ΔT in each case. When N is small, the main contributor to ΔT is the change of the heat transfer coefficient due to turbulence that is developing more with each row of cells. Once N becomes large enough, the turbulence reaches steady state before the last row. At that point, it becomes the increasing temperature of the air that drives the temperature gradient. To illustrate these different modes, simulated cell temperature is shown in Figure 3-35 for $N = 2$ and $N = 5$. As Figure 3-35a illustrates, the first row ($n = 1$) is significantly hotter than the last row ($n = 2$). On the other hand, Figure 3-35b demonstrates that the first row ($n = 1$) is cooler than the last row ($n = 5$), but still hotter than the second row ($n = 2$). The upshot of this is that the analytically modeled ΔT for small numbers of cells hinges on the correction factor c_n for the first three rows, and these correction factors were derived for relatively large tube banks. According to [58], the heat transfer of inner rows can range from 30 – 100% more than the first row and 10 – 30% higher than the second depending on the longitudinal pitch and the flow turbulence. This is an extremely wide range, and demonstrates that if N becomes small enough that c_n is the main contributor to ΔT , further investigation could be required. With a better understanding of the specific turbulence characteristics of a given application, more accurate c_n values could be determined.

Pressure Drop

While the average and delta temperature plots reveal a great deal about the necessary flow rate, the entire picture is incomplete without the pressure drop. The analytical and simulated pressure drop plots are shown in Figure 3-36. Looking first at the contour lines that represent the pressure drop across the packs, a close resemblance can be observed in the shapes of the contours. Over the whole range of flow rates and areas, the normalized root-mean-square error is 5.16%. In each case, the pressure drop for a given flow rate decreases as the cells become larger. This seems to indicate that larger cells yield lower pressure drops, but because the required flow rate increases based on the temperature constraints, the necessary pressure drop actually increases.




Figure 3-34: (a) Analytical and (b) Simulated Delta Temp [°C] for Cylindrical Cells $\varepsilon_{mod} = 278 [Wh L^{-1}]$




Figure 3-35: Simulated Temperature for (a) $N = 2$ and (b) $N = 5$ in [°C] for 20[CFM] and $\varepsilon_{mod} = 278 [Wh L^{-1}]$




Figure 3-36: (a) Analytical and (b) Simulated Pressure Drop with Temperature Constraint Contour Overlay

On each pressure plot, the 40°C average temperature and 5°C ΔT contours are overlaid on top. These contours show the minimum airflow required to satisfy each constraint, and since both constraints must be withheld, the necessary airflow is the maximum of the two contour lines for any given cell size. Looking at only the maximum of the two constraint contours on Figure 3-36 yields very similar v-shape profiles. Because the average temperature constraint becomes the driving factor for larger cells (after the two constraints intersect) the discrepancy in the delta temperature contour for larger cells has no effect on the final result. This is a convenient byproduct of the fact that packs composed of larger cells convect heat away less effectively.

Required Airflow and Pressure Drop

Figure 3-36 can be used to determine the required airflow and pressure drop for a given cell cross-sectional area. The necessary airflow for each cell size is just the maximum airflow required by either the average temperature or ΔT constraint. The required pressure can then be interpolated from the pressure contour plot. To show a clear comparison between the analytical and simulated predictions for both airflow and pressure drop across the range, Figure 3-37a illustrates these values with respect to cell size. Both the airflow and pressure drop show a reasonably good fit across




Figure 3-37: (a) Airflow and Pressure Drop Comparison and (b) Estimated Fan Cost Comparison

most of the cell size range.

Relative Cost of Cooling

Using the method described in Section 3.2.4, we estimated the relative cost of a fan or blower that develops the required airflow at the necessary pressure. Figure 3-37b shows a comparison between the analytical and simulated methods of this cost estimation. When $\varepsilon_{mod} = 278 \text{ Wh L}^{-1}$, the analytical and simulated methods yield very similar cost estimations. Most importantly, the minimum cost point predictions coincide at $A_{cell} = 403 \text{ mm}^2$ which corresponds to a cell diameter of $D = 22.7 \text{ mm}$.

3.6 Conclusion

In this chapter, I showed the importance of considering the implications of the size and shape of a lithium-ion cell on the cost of air-cooling at the onset of a battery pack design. Through simulation, I illustrated that a minimum cost cell can be found and this choice can have large impacts on the necessary airflow rate and pressure drop across that pack. I also developed an analytical method based on tube bank theory that confirms the results of the simulations for cylindrical cells.

Despite the limitations of the analytical method, the advantage of not requiring lengthy simulations through a multiphysics solver make it much more feasible for use as an early stage design tool. The results of both the simulations and analytical method highlight the need for this type of an analysis at the onset of a new battery pack design. Selecting a cell that is larger or smaller than the optimum can result in much higher required flow rates and pressure drops, which directly translates to additional cost of the fan or blower. For the parameters utilized in this study, we found that a non-optimal cell size could double or triple the cost of the blower. Additionally, the results show that fan cost is more sensitive for cells that are on the smaller side of the optimal point. This suggests that a designer might opt to err on the larger side to decrease sensitivity.

In the end, this analytical method cannot provide as much detail as a multiphysics simulation, but when the ultimate goal is to predict the overall pressure drop, average temperature, delta temperature, and cost of cooling for a given cylindrical pack at a variety of flow rates, it gives comparable results in a fraction of the time. For future work, a better understanding of c_n should be developed for small numbers of N . Furthermore, this analytical method can be extended to other form factors such as prismatic and rounded prismatic cells.

Chapter 4

Energy Management of Blended Battery Packs for Weight Minimization

In this chapter, I start by giving an overview of different types of energy storage that are currently in use in hybrid and electric vehicles. After comparing the inherent tradeoffs that are present in these energy storage technologies, I introduce the idea of tribridization through the use of blended battery packs and define the requirements in terms of drive cycles. Next, I define seven different automotive use cases that are representative of the applications where tribridization might be beneficial.

After showing how uniform battery packs are sized based on power and energy requirements, I discuss how the additional degree of freedom of controlling the power split in blended battery packs complicates their design. Next, I introduce three methods of sizing blended battery packs for minimum overall pack mass. For each of the use cases over a range of parameters, I compare the best available uniform battery pack with the blended battery packs found using the three different methods. By examining the results of this experiment, I identify which general applications are best suited for tribridization, and quantify how much potential there is to reduce the mass of conventional battery packs. Finally, I discuss the implications of these results on energy management in hybrid vehicles, and how more information is required to

actually achieve the minimum size pack.

4.1 Energy Storage Types

There are a plethora of energy storage solutions available on the market today. Each one has unique strengths and weaknesses that make them more or less suited for different applications. For hybrid and electric vehicles, Lithium-ion batteries have supplanted nickel-metal hydride to become the industry standard mostly due to their relatively high energy density, power density, and longevity of life [62]. Ultra-capacitors offer even higher power density and life [63]-[64], but have seen limited adoption in the automotive space due to extremely reduced energy density [65].

For the sake of comparison, Tables 4.1 and 4.2 show 15 different production hybrid and electric vehicles and their associated energy storage types. Additionally, three different ultra-capacitor modules have been included for reference. Table 4.2 provides details on each energy storage unit. It is important to note that these tabulated numbers are derived from pack-level attributes. In other words, each value reflects the effective cell properties as a part of a battery pack. So the tabulated masses and densities are not representative of a single cell, but of that cell and the associated packaging materials, battery management circuitry, bus bars, cooling materials, and so forth. All production vehicle storage data was supplied by the Idaho National Lab's Advanced Vehicle Testing project.

Figure 4-1 gives a visual representation of the energy storage information in Table 4.2. By plotting gravimetric energy density versus gravimetric power density, we see that there is a wide spectrum of different options. It also becomes clear there is a trade-off between energy and power density [66]. As energy density increases along the x-axis, the power density trends downward along the y-axis. This is a byproduct of the fact that the materials that increase power density (current collectors) are independent from the materials that increase energy density (active material), and both of these materials are competing for the same space within the cell [67]. This power-energy tradeoff is a fundamental barrier that stands in the way of producing

Table 4.1: Vehicle Energy Storage Information with Associated ID Numbers

ID #	Model	Class	Cell Man.	Type	Cooling
1	Tesla Model S '12	EV	Panasonic	Li-ion	Active-Liquid
2	Kia Soul '15	EV	SK Innovation	Li-ion	Active-Air
3	Mitsubishi iMiev '12	EV	GS Yuasa	Li-ion	Active-Air
4	Smart EV '14	EV	D. ACCUmotive	Li-ion	Active-Liquid
5	Chevy Spark '15	EV	LG Chem	Li-ion	Active-Liquid
6	Chevy Volt '13	PHEV	LG Chem	Li-ion	Active-Liquid
7	Nissan Leaf S '13	EV	AESC	Li-ion	Passive
8	BMW i3 EV '14	EV	Samsung SDI	Li-ion	Rerfrigerant
9	Ford Focus '13	EV	LG Chem	Li-ion	Active-Liquid
10	Ford Fusion Energi '13	PHEV	Panasonic	Li-ion	Active-Air
11	Toyota Prius '13	PHEV	Primearth	Li-ion	Active-Air
12	Ford Fusion SE '13	HEV	Panasonic	Li-ion	Active-Air
13	Honda Civic '13	HEV	Blue Energy	Li-ion	Active-Air
14	Volkswagen Jetta '14	HEV	Sanyo	Li-ion	Active-Air
15	Toyota Prius '10	HEV	Panasonic	NiMH	Active-Air
16	BMOD0006 E160 B02	NA	Maxwell	U-cap	Passive
17	GTSM-48V165FUS	NA	GreenTech	U-cap	Passive
18	BMOD0063 P125 B08	NA	Maxwell	U-cap	Active-Air

an ideal energy storage device that would sit at the upper right hand corner of Figure 4-1. It is this limitation that motivates the following sections.

4.2 Blended Battery Packs for Tribridization

A blended battery pack is one composed of multiple different energy storage types. There are several different topologies for combining these different sources which are discussed at length in Section 2.4. A tribrid is a hybrid with a blended battery pack instead of a uniform pack. The benefit of hybridizing the battery pack itself is to leverage the strengths of the different energy storage types while downplaying the weaknesses of each. For example, a high energy cell can supply the average power

Table 4.2: Energy Storage Parameters for Each ID Number

ID #	V_{nom}	$\frac{kW}{kg}$	$\frac{Wh}{kg}$	Ohms	$I_{max}[A]$	$I_{min}[A]$	$E [Wh]$	Mass [g]
1	3.6	0.57	156	0.1021	12.1	5.9	11.97	77
2	3.7	1.374	133	0.00156	392.2	233.4	140.63	1056
3	3.7	0.599	99	0.00148	303.1	141.6	185.23	1872
4	3.7	0.359	92	0.00172	199.1	234.5	189.25	2054
5	3.7	0.538	86	0.00229	162.7	101.1	95.83	1120
6	3.7	0.922	84	0.00281	170.8	57.7	57.29	685
7	3.8	0.824	83	0.00229	332.2	89.3	125	1512
8	3.7	0.656	80	0.00125	433.8	111.2	195.83	2448
9	3.7	0.498	76	0.00407	94.8	62.8	53.49	704
10	3.7	0.656	62	0.00167	260.6	167.3	90.48	1469
11	3.7	0.689	55	0.00196	265.4	188.2	78.57	1425
12	3.7	1.916	41	0.00342	235.1	173.9	18.42	454
13	3.6	1.142	31	0.0035	172.9	153.5	17	545
14	3.7	1.096	30	0.00333	180.9	158.2	18.33	605
15	1.2	0.379	25	0.00268	100.2	82.8	7.8	317
16	2.7	2.5	4	0.004	81.3	81.3	0.35	87
17	2.7	3.85	4	0.00028	1163	1163	2.93	806
18	2.7	1.7	2	0.00038	829.6	829.6	2.92	1271

while an ultra-capacitor can absorb the high power transients. The feasibility of implementing blended battery packs has been explored in several examples [25]-[38]. While it is clearly achievable to produce a pack made up of different devices, the question of which applications would benefit most from tribrirization remains. The rest of this chapter is devoted to answering this question.

4.3 Drive Cycles as a Design Requirement

Before we can answer the question of which applications are well suited for tribridization, we must be able to clearly define an application. The most common way to represent an application is to capture a snapshot of it in the form of a drive cycle. At




Figure 4-1: Gravimetric Power Density vs. Gravimetric Energy Density for Surveyed Cells




Figure 4-2: Urban Dynamometer Drive Cycle (UDDS)




Figure 4-3: Heavy Duty Vehicle Urban Dynamometer Drive Cycle (HDVUDDS)

the surface, drive cycles are very simple. They are just a speed versus time profile, such as the Urban Dynamometer Drive Schedule (UDDS) shown in Figures 4-2 and the Heavy Duty Vehicle Urban Dynamometer Drive Schedule in Figure 4-3. This profile is represented as a function of time by

$$v = \frac{dr}{dt} = v(t), \quad (4.1)$$

where r is position, t is time, and v is speed.

But when we dig a little bit deeper, we find that each drive cycle actually encapsulates a large number of automotive requirements. From the speed profile, the acceleration and linear position of the vehicle can be derived, and in combination with a relatively simple physical model of a vehicle, a power curve can be calculated that will allow a particular vehicle to follow the speed trace. We represent the power profile as

$$p = f(v(t)) = p(t), \quad (4.2)$$

where p is power and f is a function that estimates the required power based on vehicle parameters and speed (assuming level ground). For the hybrid and electric vehicles in this thesis, we utilized the ADVISOR Advanced Vehicle Simulator to generate these power curves.

The simultaneous simplicity and underlying complexity of drive cycles are why the Environmental Protection Agency (EPA) uses them to set emissions requirements and measure fuel economy. With a finite set of cycles that define "city driving" versus "highway driving", the EPA can ensure that every car they test is held to the same standards. One unfortunate side effect of using drive cycles in this way is that automotive manufacturers can tune the vehicle to perform best on the test, while real world performance may suffer. To avoid too much overfitting, automotive manufacturers also use their own internal drive cycles for many validation tasks such as endurance and performance testing.

Because of their ubiquitous usage throughout the automotive industry, drive cycles are used as the main requirements for sizing the different battery packs in this chapter. It's true that the drive cycles selected for this study may not be aggressive enough to size a pack for an actual vehicle platform (since the EPA cycles are relatively conservative), but they will nonetheless provide a good means for comparing and contrasting various control strategies and energy storage devices.

4.4 Automotive Use Cases

Before a drive cycle can be translated into a power curve, the physical parameters of a vehicle must be applied. For this reason, seven different use cases are considered in this chapter. These cases were chosen to represent a range of hybrid and electric applications from light duty to heavy duty, with special emphasis on hybrid electric vehicles. As Table 4.3 summarizes, six different vehicles and two different drive cycles are used for comparison. For all the use cases, required energy storage power versus time was simulated using the ADVISOR Advanced Vehicle Simulator developed by the National Renewable Energy Laboratory (NREL). These power curves are utilized

Table 4.3: Use Case Parameters

ID	Name	Cycle	C_d	kg	Dis. [kW]	Ch. [kW]	[Wh]
C1	Toyota Prius	UDDS	0.3	1368	21.74	-8.87	338
C2	Honda Insight	UDDS	0.25	1000	6.31	-6.42	118
C3	Hybrid Bus	HDVUDDS	0.79	15940	272.21	-160.36	12111
C4	Hybrid Bus	UDDS	0.79	15940	260.42	-136.48	15436
C5	Hybrid SUV	UDDS	0.44	2000	21.49	-14.74	404
C6	Small PHEV	UDDS	0.34	1254	34.09	-11.01	1157
C7	Ford Focus EV	UDDS	0.31	1445	39.98	-8.99	2021

as the design requirement for each pack in question.

4.5 Uniform Battery Packs

Many factors must be considered to successfully design an automotive battery pack. Simply put, the battery pack must be able to provide enough current at a high enough voltage for the vehicle to meet the worst case drive cycle over the entire life of the vehicle in the most extreme operating environments. Designing such a battery pack and then validating that it can actually meet all of these requirements is a venture that automotive manufacturers spend millions of dollars and countless engineering man-hours completing. Because of this complexity, early stage design studies are limited in scope and are used to eliminate certain choices while spending the fewest resources possible to properly evaluate them. This chapter focuses exclusively on this type of early stage design down-selection. In particular, the main focus is on power and energy requirements since a pack that does not meet both can be immediately ruled out as an option.

4.5.1 Uniform Pack Modeling and Constraints

Although there are several layout options for a battery pack, we will use the assumption that each battery pack is wired first in parallel, and then in series as shown in Figure 4-4. Additionally, every cell within that battery is modeled using an internal resistance and a ideal voltage source in series [68]. The voltage source (V_{oc}) provides




Figure 4-4: Uniform Battery Pack Layout and Internal Resistance

the open circuit voltage while the resistor R_i models the current dependent voltage drop and subsequent power loss. Both of these are illustrated with the schematics embedded in each cell in Figure 4-4.

Summing the internal voltages yields the cell output voltage

$$V_{cell} = V_{oc} - iR_i, \quad (4.3)$$

where i is the cell current.

To find the voltage of the entire pack, we can simply sum up all the series cell voltages

$$V_{pack} = SV_{oc} - SiR_i, \quad (4.4)$$

where S is the number of series cells.

The pack current is the cell current times the number of parallel cells

$$i_{pack} = iP, \quad (4.5)$$

where P is the number of parallel cells.

From Equations 4.4 and 4.5, the pack level power is just the pack current times the pack voltage

$$p = (SV_{oc} - SiR_i)iP. \quad (4.6)$$

Multiplying through yields

$$p = -SPR_i i^2 + SPV_{oc}i, \quad (4.7)$$

or

$$p = -NR_i i^2 + NV_{oc}i, \quad (4.8)$$

where N is the total number of cells in the pack and is equal to SP .

The pack level energy capacity is just the total number of cells times the cell energy capacity,

$$E_{pack} = NE, \quad (4.9)$$

or

$$E_{pack} = SPE, \quad (4.10)$$

where E_{pack} is the pack energy capacity, and E is the cell energy capacity.

Equations 4.7 and 4.10 give us the ability to calculate the pack level power and energy based on the cell level properties and the number of series and parallel cells. Using these equations in combination with the power cycle extracted from the appli-

cation drive cycle provides a means for constraining the minimum size of the battery pack.

First, we must ensure that the battery pack is capable of supplying the maximum and minimum required power. We can calculate the current at the maximum and minimum cycle power by rearranging Equation 4.7

$$0 = -SPR_i i_{max}^2 + SPV_{oc} i_{max} - p_{max}, \quad (4.11)$$

where $p_{max} = \max(p(t))$, and solving for the quadratic roots as follows

$$i_{max} = -SPV_{oc} \pm \frac{\sqrt{SP(V_{oc})^2 - 4R_i p_{max}}}{-2R_i}, \quad (4.12)$$

and similarly

$$i_{min} = -SPV_{oc} \pm \frac{\sqrt{SP(V_{oc})^2 - 4R_i p_{min}}}{-2R_i}, \quad (4.13)$$

where $p_{min} = \min(p(t))$.

For a particular cell to be able to meet the power requirements, we must therefore ensure that

$$I_{max} \geq i_{max} \text{ and } I_{min} \leq i_{min}, \quad (4.14)$$

where I_{max} is the maximum allowed cell current, I_{min} is the minimum allowed cell current, and i_{max} and i_{min} are the maximum and minimum required currents as defined in Equations 4.12 and 4.13.

We can calculate the amount of energy consumed during the cycle by taking the integral of the power profile

$$e(t) = \int p(t) dt. \quad (4.15)$$

Then, we can find the maximum range of capacity required by taking the maximum minus the minimum point of the energy cycle

$$e_{range} = e_{max} - e_{min}, \quad (4.16)$$

where $e_{max} = \max(e(t))$ and $e_{min} = \min(e(t))$.

To make sure the pack can complete the entire duration of the cycle, we must require that

$$E_{pack} \geq e_{max} - e_{min}. \quad (4.17)$$

Plugging in from Equation 4.10, we find the constraint on cell energy capacity as follows

$$E \geq \frac{e_{max} - e_{min}}{PS}. \quad (4.18)$$

4.5.2 Uniform Battery Pack Sizing

With the power constraints (Equations 4.12, 4.13, and 4.14) and energy constraint (Equation 4.18) defined, we can now begin to size uniform battery packs. In order to develop an understanding of the minimum possible battery pack size given that only one battery type can be used, it is useful to consider a number of different use cases over a range of requirements.

The best uniform battery pack for a given application can be considered the trivial solution as compared to a blended solution. There are only as many trivial solutions per application as there are possible energy storage types. Because of this, it is straightforward to iterate through each possible cell type, find the minimum mass pack that satisfies the energy and power constraints, and then choose the cell type with the overall minimum mass as the winning pack.

It is also easy to study the effect of cycle distance (total miles traveled in a given cycle) by repeating the drive cycle some number of times. In the case of pure electric vehicles, this translates directly to operating range. For a PHEV, this translates to a longer charge depleting range (resulting in greater fuel economy) until the control strategy is switched to charge-sustaining mode. The effect of repeating the cycle for




Figure 4-5: Minimum Pack Mass for Uniform Solutions for Light Duty Vehicles on a UDDS Cycle, Numbers Indicate Energy Storage ID as Defined in Table 4.1

a full hybrid vehicle will have less of an effect since the hybrid control strategies are always charge-sustaining over a long enough average.

Using this simple “brute-force” method, we found the minimal mass uniform packs for all use cases listed in Table 4.3 over a range of repeated cycles. We evaluated each of the eighteen different energy storage units listed in Table 4.2 for every use case for every number of repeated cycles. For every pack, the number of series cells (S) was held constant at 100 and the open circuit voltage is assumed to be V_{nom} . In the case of the PHEV, we assumed that charge depleting mode was maintained throughout.

Figure 4-5 shows the minimum mass uniform pack for all the light duty use cases over a range from 0.1 to 15 repeated UDDS cycles. Each cycle covers 7.45 miles, so this extends from less than a mile to over 110 miles of driving range. At each point, the ID number (from Table 4.2) of the selected energy storage unit is overlaid on top of the plot for reference. The light duty vehicle results for each case are summarized in Tables 4.4 and 4.5.

From Figure 4-5, it is clear that increasing the number of drive cycles tends to

Table 4.4: Minimum Uniform Pack Mass [kg] for Light Duty Vehicles on a UDDS Cycle

Cycles	Prius	Insight	SUV	PHEV	Focus EV
0.1	26.1	8.7	26.1	43.5	45.4
0.5	45.4	17.4	45.4	45.4	45.4
1	45.4	23.1	45.4	45.4	68.5
2	45.4	23.1	45.4	68.5	68.5
3	45.4	23.1	45.4	68.5	105.6
4	45.4	23.1	45.4	68.5	105.6
5	45.4	23.1	68.5	105.6	105.6
6	68.5	23.1	68.5	105.6	105.6
7	68.5	23.1	68.5	105.6	123.2
8	68.5	23.1	68.5	105.6	123.2
9	68.5	23.1	68.5	105.6	123.2
10	68.5	23.1	68.5	105.6	130.9
11	68.5	23.1	68.5	105.6	146.3
12	68.5	23.1	68.5	105.6	161.7
13	68.5	23.1	68.5	107.8	169.4
14	68.5	23.1	68.5	107.8	184.8
15	68.5	23.1	69.3	115.5	200.2

increase the mass of the minimally sized uniform pack. For the fewest numbers of cycles, high power density cells such as the Ford Fusion SE Panasonic cell (12) and the Maxwell E160 ultra-capacitor (16) are the optimal uniform cell choice. As the number of repeated cycles increases, the optimal cell type transitions to more and more energy-dense cells. This is because the energy requirement continues to increase while the maximum and minimum power requirements stays the same. In other words, the required C-rate decreases, and therefore, higher energy density cells can be used without running into power limitations.

In Figure 4-5, we see that the Chevy Volt LG Chem cell happens to be the best choice for the small PHEV with charge depleting ranges between 15 and 37.25 miles. This confirms that GM selected a great cell for their uniform PHEV pack. After this point, a more energy dense EV cell (2) wins out. It is also logical that for the Focus EV with ranges greater than 50 miles, the best uniform cell choice is the most energy dense cell (1) from the Tesla Model S.

Similar trends can be observed in Figure 4-6 for the hybrid bus for repeated UDDS

Table 4.5: Minimum Uniform Pack Types for Light Duty Vehicles on a UDDS Cycle

Cycles	Prius	Insight	SUV	PHEV	Focus EV
0.1	16	16	16	16	12
0.5	12	16	12	12	12
1	12	1	12	12	6
2	12	1	12	6	6
3	12	1	12	6	2
4	12	1	12	6	2
5	12	1	6	2	2
6	6	1	6	2	2
7	6	1	6	2	1
8	6	1	6	2	1
9	6	1	6	2	1
10	6	1	6	2	1
11	6	1	6	2	1
12	6	1	6	2	1
13	6	1	6	1	1
14	6	1	6	1	1
15	6	1	1	1	1




Figure 4-6: Minimum Pack Mass for Uniform Solutions for Heavy Duty Vehicles on a UDDS and HDVUDDS Cycle

Table 4.6: Minimum Uniform Pack Mass [kg] for the Hybrid Bus on the UDDS and HDVUDDS Cycles

Cycles	UDDS	HDVUDDS
0.5	182	272
1	317	317
2	317	317
3	317	422
4	422	528
5	528	634
6	634	739
7	634	778
8	739	793
9	809	893
10	809	993
11	855	1093
12	932	1194
13	1009	1294
14	1086	1394
15	1163	1486

and HDVUDDS cycles. It is important to note that the repeated hybrid bus cycle was operating in charge depleting mode, so it is more indicative of a PHEV than a full hybrid. Just like the light duty vehicles, the optimal cell choice starts with a highly power dense cell, and then transitions to more and more energy dense cells as the cycle is repeated. The full results are shown in Tables 4.6 and 4.7

4.6 Blended Battery Packs

While blended battery packs have the potential to leverage the strengths of multiple energy storage types, the expanded flexibility also complicates the design process. With the addition of an extra energy source to the battery pack, a whole new degree of freedom emerges: how to control the power split between the two devices.

During the early stages of design, it is costly to get bogged down in the specifics of how something should be controlled when the most consequential consideration is the layout and mass of the pack in question. As illustrated in Section 4.5, determining an estimated pack size for a given energy storage type is relatively straightforward

Table 4.7: Minimum Uniform Pack Types for the Hybrid Bus on the UDDS and HDVUDDS Cycles

Cycles	UDDS	HDVUDDS
0.5	12	12
1	2	2
2	2	2
3	2	2
4	2	2
5	2	2
6	2	2
7	2	1
8	2	1
9	1	1
10	1	1
11	1	1
12	1	1
13	1	1
14	1	1
15	1	1

and easy to brute force when dealing with uniform battery packs.

Optimal blended battery pack solutions, on the other hand, are practically impossible to brute force. As the length of the drive cycle increases, the time complexity grows exponentially as $O(c^N)$ where N is the length of the cycle, and c is the number of discrete power levels to be considered.

Consider the Toyota Prius on the UDDS cycle (C1 from Table 4.1) for instance. Suppose we discretize the space into 30 different possible power levels (1 per kW) per unit time. At 1369 seconds long, this translates to 30^{1369} different possible ways to control the power split. This number is astronomical, so every possibility cannot be individually examined.

In addition to the sheer number of possible power split strategies, the specific control strategy will completely change the optimal size of the blended battery pack. The end result is extremely sensitive to the selected control control strategy. These two complexities motivate the need for control strategy agnostic sizing methodologies.

4.6.1 Linear Integer Programming for Lower Bound

Perhaps the easiest, and also most trivial, method to size a blended battery pack without considering control strategy is to ignore all time dependence of the power cycle. Without time dependence, the blended pack cannot be guaranteed to successfully complete the power cycle, but this does not mean the analysis is worthless. To understand if tribridization is plausibly beneficial, it is useful to first understand the lower bound of battery pack size. This is analogous to the Carnot efficiency of a heat cycle. While it might not always be practically realizable, it's still useful to understand the best possible scenario of efficiency for the purpose of benchmarking.

Linear Integer Programming Modeling and Constraints

By ignoring time dependence, we do not need to iterate through the entire cycle. We must simply ensure that the combined power and energy capability of the blended battery pack is greater than the overall demand of the cycle. The first requirement is that the two energy sources must be able to supply enough combined power to fulfill the maximum discharge and charge power experienced during the cycle.

From Equation 4.7, we can represent the maximum available baseload power as

$$p_{max_{base}} = -S_{base}P_{base}R_{i_{base}}I_{max_{base}}^2 + S_{base}P_{base}V_{oc_{base}}I_{max_{base}}, \quad (4.19)$$

where S_{base} is the number of series baseload cells, P_{base} is the number of parallel baseload cells, $R_{i_{base}}$ is the internal resistance of a baseload cell, $V_{oc_{base}}$ is the open circuit voltage of a baseload cell, and $I_{max_{base}}$ is the maximum available current from a baseload cell.

The minimum available baseload power is

$$p_{min_{base}} = -S_{base}P_{base}R_{i_{base}}I_{min_{base}}^2 + S_{base}P_{base}V_{oc_{base}}I_{min_{base}}. \quad (4.20)$$

where $I_{min_{base}}$ is the minimum available current from a baseload cell.

Similarly, the maximum available buffer power is

$$p_{max_{buffer}} = -S_{buffer}P_{buffer}R_{i_{buffer}}I_{max_{buffer}}^2 + S_{buffer}P_{buffer}V_{oc_{buffer}}I_{max_{buffer}}, \quad (4.21)$$

where S_{buffer} is the number of series buffer cells, P_{buffer} is the number of parallel buffer cells, $R_{i_{buffer}}$ is the internal resistance of a buffer cell, $V_{oc_{buffer}}$ is the open circuit voltage of a buffer cell, and $I_{max_{buffer}}$ is the maximum available current from a buffer cell.

The minimum available buffer power is

$$p_{min_{buffer}} = -S_{buffer}P_{buffer}R_{i_{buffer}}I_{min_{buffer}}^2 + S_{buffer}P_{buffer}V_{oc_{buffer}}I_{min_{buffer}}. \quad (4.22)$$

where $I_{min_{base}}$ is the minimum available current from a buffer cell.

Summing the maximum power from the baseload and buffer yields the total power capability of the blended pack

$$p_{max_{blended}} = p_{max_{base}} + p_{max_{buffer}}. \quad (4.23)$$

Plugging in from Equations 4.19 and 4.21 and factoring yields

$$p_{max_{blended}} = S(-P_{base}(R_{i_{base}}I_{max_{base}}^2 + V_{oc_{base}}I_{max_{base}}) - P_{buffer}(R_{i_{buffer}}I_{max_{buffer}}^2 + V_{oc_{buffer}}I_{max_{buffer}})), \quad (4.24)$$

where $S_{base} = S_{buffer} = S$.

Summing the minimum power from the baseload and buffer yields the total regenerative power capability of the blended pack

$$p_{min_{blended}} = p_{min_{base}} + p_{min_{buffer}}. \quad (4.25)$$

Plugging in for the minimum powers from Equations 4.20 and 4.22 yields the minimum blended power

$$p_{min_{blended}} = S(-P_{base}(R_{i_{base}} I_{min_{base}}^2 + V_{oc_{base}} I_{min_{base}}) - P_{buff}(R_{i_{buff}} I_{min_{buff}}^2 + V_{oc_{buff}} I_{min_{buff}})), \quad (4.26)$$

where $S_{base} = S_{buff} = S$.

The second constraint is that there must be enough energy capacity in the entire pack to complete the cycle. Starting with the baseload and according to Equation 4.10, the total available energy is

$$E_{pack_{base}} = S_{base} P_{base} E_{base}. \quad (4.27)$$

For the buffer, the same is true

$$E_{pack_{buff}} = S_{buff} P_{buff} E_{buff}. \quad (4.28)$$

Summing the available baseload and buffer energy yields the total available energy from the blended pack

$$E_{pack_{blended}} = E_{pack_{base}} + E_{pack_{buff}}. \quad (4.29)$$

Plugging in from Equations 4.27 and 4.28 and factoring gives the expanded form

$$E_{pack_{blended}} = S(P_{base} E_{base} + P_{buff} E_{buff}), \quad (4.30)$$

where $S_{base} = S_{buff} = S$.

Before a linear integer programming (IP) algorithm can be applied, it is useful to phrase the problem in the canonical form:

$$\text{maximize } c^T x \begin{cases} Ax \leq b \\ x \geq 0 \\ x \in \mathbb{Z}^n \end{cases} \quad (4.31)$$

In order to minimize the overall blended battery pack mass, we must maximize




Figure 4-7: Minimum Pack Mass for Integer Programming and Uniform Solutions for the Toyota Prius on a UDDS Cycle

the negative masses

$$c = \begin{bmatrix} -M_{base} \\ -M_{buff} \end{bmatrix}, \quad (4.32)$$

where M_{base} is the mass of a baseload cell, and M_{buff} is the mass of a buffer cell.

The working variables are the number of parallel cells in the baseload and buffer

$$x = \begin{bmatrix} P_{base} \\ P_{buff} \end{bmatrix}, \quad (4.33)$$

where M_{base} is the mass of a baseload cell, and M_{buff} is the mass of a buffer cell.

To ensure that the power and energy requirements are met, there are three constraints:

$$p_{maxblended} \geq p_{max}, \quad (4.34)$$




Figure 4-8: Minimum Pack Mass for Integer Programming and Uniform Solutions for the Honda Insight on a UDDS Cycle

where $p_{max} = \max(p(t))$,

$$p_{minblended} \leq p_{min}, \tag{4.35}$$

where $p_{min} = \min(p(t))$,

and

$$E_{packblended} \geq e_{max} - e_{min}, \tag{4.36}$$

where $e_{max} = \max(e(t))$ and $e_{min} = \min(e(t))$.

From Relations 4.34, 4.35, 4.36, and Equations 4.24, 4.26, and 4.30, the constraint matrix can be developed as




Figure 4-9: Minimum Pack Mass for Integer Programming and Uniform Solutions for the Hybrid Bus on a HDVUDDS Cycle

$$A = S \begin{bmatrix} R_{i_{base}} I_{max_{base}}^2 - V_{oc_{base}} I_{max_{base}} & R_{i_{buff}} I_{max_{buff}}^2 - V_{oc_{buff}} I_{max_{buff}} \\ -R_{i_{base}} I_{min_{base}}^2 + V_{oc_{base}} I_{min_{base}} & -R_{i_{buff}} I_{min_{buff}}^2 + V_{oc_{buff}} I_{min_{buff}} \\ -E_{base} & -E_{buff} \end{bmatrix}, \quad (4.37)$$

Finally, the upper bound matrix is just the other sides of Relations 4.34, 4.35, and 4.36

$$b = \begin{bmatrix} -p_{max} \\ p_{min} \\ -e_{max} + e_{min} \end{bmatrix}. \quad (4.38)$$

With the objective and constraints in the canonical form, the matrices can be plugged into a solver for any given application.




Figure 4-10: Minimum Pack Mass for Integer Programming and Uniform Solutions for the Hybrid Bus on a UDDS Cycle

Linear Integer Programming for Blended Battery Pack Sizing

Figures 4-7, 4-8, 4-9, 4-10, 4-11, 4-12, and 4-13 show the results of the linear integer programming optimization over a range of cycle counts for all the use cases (see Table 4.3 for use case definitions). Each figure plots the integer programming result alongside the best uniform pack solution for that use case. The wider the gap between the two traces, the more potential there is for improving over a uniform battery pack through tribridization.

Perhaps the most useful insight to be gained from the integer programming solution is at what point it converges with the uniform battery pack solution. Once this occurs, there is no reason to explore tribridization. For the Toyota Prius (Figure 4-7) and Hybrid SUV (Figure 4-11) use cases, the IP solution does not converge to the uniform solution within the range of the plot. This suggests that tribridization could be very beneficial for these applications.

For the Ford Focus EV (Figure 4-13), on the other hand, the two solutions converge




Figure 4-11: Minimum Pack Mass for Integer Programming and Uniform Solutions for the Hybrid SUV on a UDDS Cycle

at the ninth repeated UDDS cycle, which corresponds to approximately 67 miles. Since the Focus boasts a range of over 70 miles, this suggests that a uniform pack is the best solution for this use case. This is not surprising since the Focus demands approximately five times more energy per cycle than the similarly sized hybrids while the power is only about double. This translates to significantly reduced C rates, which can more easily be supported by high energy density cells. In other words, the energy demand is the dominant constraint. This conclusion is transferrable to most full EV applications, and suggests that EVs are inherently less suited for blended battery packs than hybrids from a pure mass reduction standpoint.

The IP and uniform solutions also converge in the small PHEV use case (Figure 4-12). In this case, the intersection occurs on the thirteenth repeated UDDS cycle, which corresponds to approximately 97 miles of driving range. Unlike EVs, PHEVs usually only have charge depleting ranges of less than 50 miles. This means that as long as the desired charge depleting range is below the 97 mile threshold, tribridization has the potential to offer significant weight reduction.




Figure 4-12: Minimum Pack Mass for Integer Programming and Uniform Solutions for the PHEV on a UDDS Cycle

The hybrid bus use cases (Figures 4-9 and 4-10) indicate that there is potential for improvement over the uniform solution if the cycle is repeated fewer than 11 times for the HDVUDDS cycle (61 miles), and fewer than 8 times for the UDDS cycle (or 60 miles). This indicates that there is likely room for mass reduction since a charge depleting cycle is not expected to continue for more than the 60 miles of range on hybrid electric buses.

It is very important to note that the integer programming solution may meet the constraints that were imposed, but these do not guarantee that the pack can successfully complete the cycle. It just means that the blended pack can output the maximum power, absorb the minimum power, and also has more energy capacity than the cycle requires. The main failure mode of these assumptions occurs when the required power exceeds the baseload power capacity for a long enough time that the buffer runs out of usable energy. It is for this reason that further investigation is required to confirm that a blended battery pack can offer reduced mass with respect to a uniform pack.




Figure 4-13: Minimum Pack Mass for Integer Programming and Uniform Solutions for the Focus EV on a UDDS Cycle

4.6.2 Dynamic Programming for Minimum Achievable Pack Size

With the lower bound understood, we can now expend extra resources on the applications with room for improvement through tribridization. To see what the minimum achievable pack size is, we need to find the control strategy that minimizes the mass of the pack. Conversely, to fully optimize the control strategy, we need to know what the pack layout is. In light of this “chicken-and-the-egg” predicament, we must somehow simultaneously size the blended battery pack while controlling it in such a way as to minimize the resulting mass.

Objective Function

Before discussing the structure of the algorithm, it is important to understand the objective function. This function effectively controls the power split between the baseload and the buffer throughout the cycle. By splitting the overall function into

three different objectives, we can ensure that they are prioritized appropriately. The second objective acts as a tie breaker for the first objective, and the third objective function acts as a tie breaker for the second.

Objective 1 - Total Mass of the Pack The number one priority in controlling the power split is to minimize the overall mass of the blended battery pack. This function takes the form of

$$f_1(P_{base}, P_{buff}) = SP_{base}M_{base} + SP_{buff}M_{buff}. \quad (4.39)$$

Objective 2 - Effective Mass of Energy To act as a tiebreaker between paths that allow for the same mass battery pack, two different objectives are used. The first is applied during charge-sustaining mode, and the second applies during charge-depletion mode. Charge-depletion mode is used when the energy available in the baseload and buffer is greater than the energy remaining in the cycle.

For charge sustaining mode, we have

$$f_{2a}(e_{base}, e_{buff}) = \frac{e_{base}M_{base}}{E_{base}} + \frac{e_{buff}M_{buff}}{E_{base}}, \quad (4.40)$$

and for charge depletion mode, we use

$$f_{2b}(e_{base}, e_{buff_{CD}}) = \frac{e_{base}M_{base}}{E_{base}} + \frac{e_{buff_{CD}}M_{buff}}{E_{base}}. \quad (4.41)$$

where $e_{buff_{CD}}$ is the absolute difference between the available buffer energy divided by the remaining number of steps and the energy added to the buffer in the current time step.

Objective 3 - Minimize Inefficiency The third objective and final tiebreaker is to choose the most efficient path with respect to I^2R losses.

$$f_3(e_{R_{base}}, e_{R_{buff}}) = e_{R_{base}} + e_{R_{buff}} \quad (4.42)$$

Deterministic Dynamic Programming Based Algorithm

Deterministic dynamic programming (DDP) is commonly used to optimize power split between two sources such as an engine and a motor in a hybrid vehicle [69]-[70]. However, these sources are fixed in size throughout the cycle. The fundamental assumption of dynamic programming is that a very complicated problem can be solved by breaking that problem down into a finite number of simpler subproblems, solving them, and storing the answers. Combining the solutions to all the subproblems ultimately yields the overall solution. This memoization technique cuts out a vast number of steps that would otherwise be repeated in a brute force approach.

The assumption that the optimally solved subproblems will cumulatively solve the overall problem breaks down when the size of the battery pack is changing over the course of the cycle. For example, if the baseload grows steadily as the algorithm iterates backward through the cycle, the power capability will increase as it grows. It's possible that once the algorithm reaches the start of the cycle, the baseload will have enough power capability to deliver the entirety of the peak power demand. In this case, it might improve the solution if the baseload delivers more or all of the demand throughout the previously solved subproblems. For this reason, the algorithm cannot be a direct application of DDP and we must use a heuristic approach that solves the problem iteratively [71].

DDP Algorithm Step 1 Figure 4-14 shows a simplified flow chart of the utilized dynamic programming based method. The first step in the iterative algorithm is the same as all the other approaches: translate the drive cycle into a power cycle using the Advisor Advanced Vehicle Simulator. An example section of a power cycle is shown in Figure 4-15.

DDP Algorithm Step 2 The next step (number 2 in Figure 4-14) is to discretize the domain to the desired granularity. This is represented by the black grid overlaid on top of the power cycle in Figure 4-16. The grid represents the possible baseload power levels (p_{base}) throughout the analysis. Assuming no loss of power in the conversion,




Figure 4-14: Deterministic Dynamic Programming Based Algorithm for Blended Battery Pack Sizing




Figure 4-15: Example Portion of Power Cycle for UDDS Cycle

the buffer power (p_{buffer}) is simply found by subtracting the baseload power from the total power

$$p_{buffer}(t, g) = p(t) - p_{base}(t, g), \quad (4.43)$$

where t is time (or horizontal grid position) and g is vertical grid position.

The number of vertical elements should be sized based on the necessary level of accuracy and the available computational resources. The finer the grid, the higher the accuracy, but computation times will increase as $O(NG^2)$ where G is the vertical grid size, and N is the number of horizontal elements. The grid must include 0 as one of the vertical elements or the all-buffer case will not be a possible solution. It should also be noted that if the precision of the power cycle is finer than the resolution of the grid, the all-baseload case cannot be found as a possible solution (as the value of $p_{buffer}(t)$ will always be non-zero). This is not generally a problem since it is relatively trivial to check the all-baseload case manually (as outlined in Section 4.5). The




Figure 4-16: Discretized Example Portion of Power Cycle for UDDS Cycle

number of horizontal elements should be equal to the total number of sample points in the drive cycle (usually sampled at $1Hz$). Finally, the mass of the minimum pack (M) is set to infinity.

DDP Algorithm Step 3, $t = N$ Zooming in on the grid section outlined in red in Figure 4-16 brings about the next step (the first iteration of step number 3 in Figure 4-14). The zoomed view is shown in Figure 4-17. The algorithm enters the grid once for each vertical level at $t = N$. Each blue circle represents an instantaneous baseload power level ($p_{base}(N)$). The buffer power level at position g is found from equation 4.43 as

$$p_{buff}(N, g) = p(N) - p_{base}(N, g). \quad (4.44)$$

Since this is the first step, the maximum and minimum power levels are both set equal to their respective instantaneous powers at N . The maximum and minimum baseload powers are

$$p_{max_{base}}(N, g) = p_{min_{base}}(N, g) = p_{base}(N, g), \quad (4.45)$$

and the maximum and minimum buffer power is

$$p_{max_{buff}}(N, g) = p_{min_{buff}}(N, g) = p_{buff}(N, g). \quad (4.46)$$

In this case, the maximum, minimum, and current energy states are set to zero, since no time has passed. So the maximum ($e_{max_{base}}$), minimum ($e_{min_{base}}$), and current (e_{base}) baseload energies are

$$e_{max_{base}}(N, g) = e_{min_{base}}(N, g) = e_{base} = 0, \quad (4.47)$$

and the maximum, minimum, and current buffer energies are

$$e_{max_{buff}}(N, g) = e_{min_{buff}}(N, g) = e_{buff} = 0. \quad (4.48)$$

Initially, all parallel cell counts are also set to zero

$$P_{buff}(N, g) = P_{base}(N, g) = 0, \quad (4.49)$$

and the number of series cells is assumed to be a constant S throughout.

With $p_{max_{base}}(N, g)$, $p_{min_{base}}(N, g)$, $p_{max_{buff}}(N, g)$, $p_{min_{buff}}(N, g)$, $e_{max_{base}}(N, g)$, $e_{min_{base}}(N, g)$, $e_{max_{buff}}(N, g)$, and $e_{min_{buff}}(N, g)$ known for all grid positions (g), the uniform pack sizing methodology described in Section 4.5 can be applied to both the baseload and the buffer individually. Using Relations 4.14 and 4.18, the number of parallel cells ($P_{base}(N, g)$ and $P_{buff}(N, g)$) are incremented until all constraints are satisfied.

Since this is the first time step, and no time has passed, the energy lost due to the internal resistance of the packs is set to zero for each position

$$e_{R_{base}}(N, g) = e_{R_{buff}}(N, g) = 0. \quad (4.50)$$

DDP Algorithm Step 3, $t < N$ As Figure 4-14 illustrates, Step 3 is repeated for every time step moving backward in time until the first point in the cycle ($t = 1$) is finally reached. Each of the steps from $t = N - 1$ to $t = 1$ follow the process represented in Figures 4-18 and 4-19 and described below.

Each point in the grid at time step t is evaluated from each point in the grid at




Figure 4-17: Discretized Example Portion of Power Cycle for UDDS Cycle Zoomed to Red Region from Figure 4-16




Figure 4-18: DDP Algorithm Substep Evaluation Technique for UDDS Power Cycle Example

time step $t + 1$. Figure 4-18 illustrates the path lines from each grid point $(g_{t+1})_1^G$ at $t = N$ to the first grid point $(g_t = 1)$ at $t = N - 1$.

The blue circles at $t = N - 1$ represent an instantaneous baseload power level $(p_{base}(t, g_t))$. The buffer power level at position g_t is found from equation 4.43 as

$$p_{buff}(t, g_t) = p(t) - p_{base}(t, g_t). \quad (4.51)$$

The maximum and minimum power levels are both updated to reflect the new maximum and minimum powers. Temporary values are denoted with an asterisk and stored for each possible path. All but the best path according to the objective function will be discarded. The temporary maximum and minimum baseload powers for a path from g_t to g_{t+1} are

$$p_{max_{base}}^*(t, g_t, g_{t+1}) = \max [p_{base}(t, g_t), p_{max_{base}}(t + 1, g_{t+1})], \quad (4.52)$$

and

$$p_{min_{base}}^*(t, g_t, g_{t+1}) = \min [p_{base}(t, g_t), p_{min_{base}}(t + 1, g_{t+1})]. \quad (4.53)$$

The temporary maximum and minimum buffer powers for a path from g_t to g_{t+1} are

$$p_{max_{buff}}^*(t, g_t, g_{t+1}) = \max [p_{buff}(t, g_t), p_{max_{buff}}(t + 1, g_{t+1})], \quad (4.54)$$

and

$$p_{min_{buff}}^*(t, g_t, g_{t+1}) = \min [p_{buff}(t, g_t), p_{min_{buff}}(t + 1, g_{t+1})]. \quad (4.55)$$

Since time has now passed, the current energy states are updated. Once again, these values are temporary and will be denoted with an asterisk. The baseload energy is

$$e_{base}^*(t, g_t, g_{t+1}) = e_{base}(t + 1, g_{t+1}) + \frac{(p_{base}(t, g_t) + p_{base}(t, g_{t+1}))}{2}, \quad (4.56)$$

assuming the time step is 1 second.

Similarly, the buffer energy is

$$e_{buff}^*(t, g_t, g_{t+1}) = e_{buff}(t + 1, g_{t+1}) + \frac{(p_{buff}(t, g_t) + p_{buff}(t, g_{t+1}))}{2}, \quad (4.57)$$

The temporary maximum ($e_{max_{base}}^*$) and minimum ($e_{min_{base}}^*$) baseload energies are

$$e_{max_{base}}^*(t, g_t, g_{t+1}) = \max [e_{base}^*(t, g_t, g_{t+1}), e_{max_{base}}(t + 1, g_{t+1})], \quad (4.58)$$

and

$$e_{min_{base}}^*(t, g_t, g_{t+1}) = \min [e_{base}^*(t, g_t, g_{t+1}), e_{min_{base}}(t + 1, g_{t+1})]. \quad (4.59)$$

The maximum ($e_{max_{buffer}}^*$) and minimum ($e_{min_{buffer}}^*$) buffer energies are of the same form

$$e_{max_{buffer}}^*(t, g_t, g_{t+1}) = \max [e_{buffer}^*(t, g_t, g_{t+1}), e_{max_{buffer}}^*(t+1, g_{t+1})], \quad (4.60)$$

and

$$e_{min_{buffer}}^*(t, g_t, g_{t+1}) = \min [e_{buffer}^*(t, g_t, g_{t+1}), e_{min_{buffer}}^*(t+1, g_{t+1})]. \quad (4.61)$$

With $p_{max_{base}}^*(t, g_t, g_{t+1})$, $p_{min_{base}}^*(t, g_t, g_{t+1})$, $p_{max_{buffer}}^*(t, g_t, g_{t+1})$, $p_{min_{buffer}}^*(t, g_t, g_{t+1})$, $e_{max_{base}}^*(t, g_t, g_{t+1})$, $e_{min_{base}}^*(t, g_t, g_{t+1})$, $e_{max_{buffer}}^*(t, g_t, g_{t+1})$, and $e_{min_{base}}^*(t, g_t, g_{t+1})$ known for all grid positions (g_{t+1}), the uniform pack sizing methodology described in Section 4.5 can be applied to both the baseload and the buffer individually. Using Relations 4.14 and 4.18, the number of parallel cells ($P_{base}^*(t, g_t, g_{t+1})$ and $P_{buffer}^*(t, g_t, g_{t+1})$) are incremented until all constraints are satisfied for each temporary path.

For evaluation of the objective function, the energy lost due to the internal resistance of the baseload and buffer is recorded for each temporary path. First, the instantaneous current is found for both the baseload and buffer using the same method as Equation 4.12. The baseload current is

$$i_{base}^*(t, g_t, g_{t+1}) = -SP_{base}^*(t, g_t, g_{t+1})V_{oc_{base}} \pm \frac{\sqrt{SP_{base}^*(t, g_t, g_{t+1})(V_{oc_{base}})^2 - 4R_{i_{base}}p_{base}(t, g_t)}}{-2R_{i_{base}}}, \quad (4.62)$$

and the buffer current is

$$i_{buffer}^*(t, g_t, g_{t+1}) = -SP_{buffer}^*(t, g_t, g_{t+1})V_{oc_{buffer}} \pm \frac{\sqrt{SP_{buffer}^*(t, g_t, g_{t+1})(V_{oc_{buffer}})^2 - 4R_{i_{buffer}}p_{buffer}(t, g_t)}}{-2R_{i_{buffer}}}. \quad (4.63)$$

With the currents known for each path, the temporary cumulative energy loss can




Figure 4-19: DDP Algorithm Step N-1 Evaluation for UDDS Power Cycle Example

be calculated for the baseload as

$$e_{R_{base}}^*(t, g_t, g_{t+1}) = SP_{base}^*(t, g_t, g_{t+1})(i_{base}^*(t, g_t, g_{t+1}))^2 R_{i_{base}} + e_{R_{base}}(t+1, g_{t+1}), \quad (4.64)$$

and for the buffer as

$$e_{R_{buff}}^*(t, g_t, g_{t+1}) = SP_{buff}^*(t, g_t, g_{t+1})(i_{buff}^*(t, g_t, g_{t+1}))^2 R_{i_{buff}} + e_{R_{buff}}(t+1, g_{t+1}), \quad (4.65)$$

assuming the time step is 1 second.

With $P_{base}^*(t, g_t, g_{t+1})$, $P_{buff}^*(t, g_t, g_{t+1})$, $e_{base}^*(t, g_t, g_{t+1})$, $e_{buff}^*(t, g_t, g_{t+1})$, $e_{R_{base}}^*(t, g_t, g_{t+1})$, and $e_{R_{buff}}^*(t, g_t, g_{t+1})$ known for all all grid positions (g_{t+1}), the objective function is used to evaluate the best path. In Figure 4-18 this path is depicted in red and the yellow paths are discarded. The index at $t+1$ from which the minimum mass pack is generated for g_t is denoted as $g_{min_{t+1}}$.

The parameters resulting in the minimum mass pack are then stored at t for g_t for reference at $t-1$

$$p_{max_{base}}(t, g_t) = p_{max_{base}}^*(t, g_t, g_{min_{t+1}}) \quad (4.66)$$

$$p_{min_{base}}(t, g_t) = p_{min_{base}}^*(t, g_t, g_{min_{t+1}}) \quad (4.67)$$

$$p_{max_{buff}}(t, g_t) = p_{max_{buff}}^*(t, g_t, g_{min_{t+1}}) \quad (4.68)$$




Figure 4-20: DDP Algorithm Step N-2 Evaluation for UDDS Power Cycle Example

$$p_{min_{buffer}}(t, g_t) = p_{min_{buffer}}^*(t, g_t, g_{min_{t+1}}) \quad (4.69)$$

$$e_{base}(t, g_t) = e_{base}^*(t, g_t, g_{min_{t+1}}) \quad (4.70)$$

$$e_{buffer}(t, g_t) = e_{buffer}^*(t, g_t, g_{min_{t+1}}) \quad (4.71)$$

$$e_{max_{base}}(t, g_t) = e_{max_{base}}^*(t, g_t, g_{min_{t+1}}) \quad (4.72)$$

$$e_{min_{base}}(t, g_t) = e_{min_{base}}^*(t, g_t, g_{min_{t+1}}) \quad (4.73)$$

$$e_{max_{buffer}}(t, g_t) = e_{max_{buffer}}^*(t, g_t, g_{min_{t+1}}) \quad (4.74)$$

$$e_{min_{buffer}}(t, g_t) = e_{min_{buffer}}^*(t, g_t, g_{min_{t+1}}) \quad (4.75)$$

$$e_{R_{base}}(t, g_t) = e_{R_{base}}^*(t, g_t, g_{min_{t+1}}) \quad (4.76)$$

$$e_{R_{buffer}}(t, g_t) = e_{R_{buffer}}^*(t, g_t, g_{min_{t+1}}) \quad (4.77)$$

$$P_{base}(t, g_t) = P_{base}^*(t, g_t, g_{min_{t+1}}) \quad (4.78)$$

$$P_{buffer}(t, g_t) = P_{buffer}^*(t, g_t, g_{min_{t+1}}) \quad (4.79)$$

The process of calculating all temporary paths and selecting the best is repeated for each point $(g_t]_1^G)$. Figure 4-19 illustrates the best path for each grid point at t .

This process of finding the best path from $t + 1$ to each point g_t is repeated until $t = 1$ is reached. Figures 4-20 and 4-21 show an example of the repeated process for $t = N - 2$ and $t = N - 3$.

The best of all the paths according to the objective function at $t = 1$ is selected




Figure 4-21: DDP Algorithm Step N-3 Evaluation for UDDS Power Cycle Example




Figure 4-22: DDP Algorithm Best Path for UDDS Power Cycle Example

as the overall minimum mass path. The selected path is shown as a black dashed line in Figure 4-22.

DDP Algorithm Step 4 The mass of the battery pack resulting from the best path at $t = 1$ is stored as m and compared to the minimum pack mass M . If m is not equal to M , the algorithm has not converged, and so M is set equal to m , t is reset to N , and the buffer parameters are reset. The baseload parameters are also reset except for the number of parallel cells, which remains at the value from the end of the current iteration.

For the first iteration, M was set to infinity, so the cycle must repeat at least once. The reason for resetting the buffer is to allow the baseload to grow until more baseload does not improve the solution. Without this buffer resetting technique, the higher power density of the buffer tends to force the algorithm to overutilize the buffer at the beginning of each iteration when the required energy is small in comparison to




Figure 4-23: UDDS Power Cycle Example with Power Baseload Power Thresholds

the minimum and maximum power requirements.

DDP Algorithm Step 5 Once the same algorithm returns the same pack twice in a row, we know that the baseload did not grow over the course of the cycle, and the solution has been reached.

4.6.3 Peak Shaving Strategy as a Benchmark

While the dynamic programming based solution can find a minimum achievable pack size, it doesn't imply that the control strategy is practically realizable. This is due to the omniscient nature of the dynamic programming algorithm. If the future is completely unknown, a less efficient control strategy must be utilized [72]. To see which applications lend well to a more practical control strategy, we will use a peak shaving benchmark heuristic.

Peak Shaving Method

The peak shaving method is very simple in principle. A power threshold (p_{thresh}) is selected for the given application. Any power below that threshold is assigned to the baseload, and any power over that threshold is assigned to the buffer [73]-[75]. Figure 4-23 shows an example UDDS Power cycle with positive and negative baseload power thresholds overlaid on top. The magnitude of the threshold is a design choice, and could be changed dynamically, but in this thesis it is assumed to be constant.




Figure 4-24: Peak Shaving Baseload Power Curve for UDDS Power Cycle Example




Figure 4-25: Peak Shaving Buffer Power Curve for UDDS Power Cycle Example

The baseload power ($p_{base}(t)$) can be represented as a piecewise function

$$p_{base}(t) = \begin{cases} p(t) & \text{if } |p(t)| \leq p_{thresh} \\ p_{thresh} & \text{if } p(t) > p_{thresh} \\ -p_{thresh} & \text{if } p(t) < -p_{thresh} \end{cases} \quad (4.80)$$

and then the buffer power ($p_{buff}(t)$) can be found by subtracting the baseload power from the total power

$$p_{buff}(t) = p(t) - p_{base}(t). \quad (4.81)$$

The baseload power curve is shown for the example UDDS power cycle in Figure 4-24 and the buffer power curve is shown for the same cycle in Figure 4-25. With the power cycles defined for both the baseload and buffer, the blended battery pack sizing problem can be treated as two uniform battery packs. The same uniform battery pack




Figure 4-26: Minimum Pack Mass for DDP, Integer Programming, Peak Shaving, and Uniform Solutions for the Toyota Prius on a UDDS Cycle

sizing methodology described in Section 4.5 is used to identify the minimum necessary number of parallel baseload cells ($P_{baseload}$) and buffer cells (P_{buffer}).

This peak shaving method can be repeated for power thresholds over the entire range of possible power levels with some granularity. It is important to include both zero and a threshold greater than the peak total power so that the all baseload and all buffer cases can be considered. The minimum mass pack out of all the power threshold levels is taken to be the peak shaving solution.

4.7 Study Results

By utilizing all four of the battery pack sizing methodologies listed in Sections 4.5 and 4.6, we can examine particular use cases and understand how the minimum mass uniform battery pack compares to the blended solutions that are sized with the three different methods. As Section 4.6 describes, the integer programming (IP) method provides a lower bound, the dynamic programming (DDP) based method yields a




Figure 4-27: Minimum Pack Mass for DDP, Integer Programming, Peak Shaving, and Uniform Solutions for the Hybrid Bus on a UDDS Cycle

minimum feasible pack, and the peak shaving (PS) method gives a good baseline.

4.7.1 Design of Experiment

To understand which applications lend well to tribridization, we designed and carried out a large experiment that sweeps over a range of different applications and potential energy storage units. For each of the usage cases listed in Table 4.3, we found solutions using all four methodologies (DDP, IP, PS, and uniform) for 1 through 10 repeated drive cycles (7.45 -74.5 miles of driving) while considering all 18 of the energy storage units listed in Table 4.1. We will discuss the results for HEVs, PHEVs, and EVs separately in the following three sections. The full results tables are listed in Appendix A.




Figure 4-28: Minimum Pack Mass for DDP, Integer Programming, Peak Shaving, and Uniform Solutions for the Hybrid SUV on a UDDS Cycle

4.7.2 Hybrid Electric Vehicles

Figures 4-26, 4-27, and 4-28 show the results for the Toyota Prius, Hybrid Bus, and Hybrid SUV respectively. Each figure depicts the minimum mass pack for all four methods across 1 through 10 repeated drive cycles. The actual energy requirement of the application will determine how many repeated drive cycles are appropriate, but for hybrid vehicles in general, the energy requirement corresponds to about 4 repeated cycles.

Taking a closer look at Figure 4-26 shows the minimum mass packs found using the DDP, IP, PS, and uniform solutions for the Toyota Prius. Up until 5 repeated cycles, the uniform and PS solutions converge to the same result of a 1P pack of Ford Fusion SE cells (ID number 12) at 45.4 kg. The DDP solution offers a lower mass pack, but only saves about 6 kg by blending between a 4P pack of Tesla Model S cells and a 1P pack of Maxwell Ultracaps.

Figure 4-27 shows the minimum mass packs found using the DDP, IP, PS, and




Figure 4-29: Gravimetric Power Density vs. Gravimetric Energy Density for Surveyed Cells (black) and DDP Solutions (red) for 4 Consecutive Cycles

uniform solutions for the Hybrid Bus on a UDSS drive cycle. For this use case, the blended solutions offer an improvement over the uniform solution up until 8 repeated cycles. All of the blended solutions are relatively similar across the board. At 4 repeated cycles, the DDP and PS solutions have converged to a 442.2 kg pack composed of a 30P pack of Tesla Model S cells and a 2P pack of Kia Soul cells. The best uniform pack is a 5P pack of Kia Soul cells that weighs 528 kg. This translates to a potential savings of almost 86 kg over the uniform pack through tribridization.

For the Hybrid SUV on a UDSS drive cycle, Figure 4-28 shows the minimum mass packs found using the DDP, IP, PS, and uniform solutions. While there are areas where blended packs offer significant improvement over the uniform pack, the two feasible blended solutions converge to the uniform pack when the number of repeated drive cycles is 4. At this point, the DDP, PS, and uniform solutions are all 1P packs of Ford Fusion SE cells. The IP solution offers substantial gains, but it is not likely to be achievable in practice.

Figure 4-29 shows all 18 of the uniform storage options (from Figure 4-1) with the




Figure 4-30: Least Massive Pack for DDP, Peak Shaving, and Uniform Solution Methods for 4 Consecutive Cycles

blended solutions for each of the seven use cases overlaid on top in red for 4 consecutive drive cycles. It is clear that each of the red (blended) solutions is interpolating between the energy density and power density of two of the uniform solutions. The hybrid bus solutions, for instance, interpolate between the Tesla Model S cell (number 1) and the Kia Soul cell (number 2). In this way, the blended packs can linearly fill in the gaps between the various uniform options.

To show all of the solutions for each use case with 4 repeated drive cycles, Figure 4-30 depicts a summary bar plot. From this plot, it is clear that the hybrid bus applications (3 and 4) offer the largest room for improvement through blended battery packs. The Ford Focus EV (7) also shows significant room for improvement, but 4 repeated cycles (approximately 30 miles) for this application is not an acceptable EV range. This energy requirement would be more representative of a PHEV.




Figure 4-31: Minimum Pack Mass for DDP, Integer Programming, Peak Shaving, and Uniform Solutions for the PHEV on a UDDS Cycle

4.7.3 Plug-in Hybrid Electric Vehicles

Figure 4-31 shows the minimum mass packs found using the DDP, IP, PS, and uniform solutions for the PHEV on a UDDS drive cycle. With the exception of one repeated cycle, the blended battery packs show potential for improvement over the uniform packs across the board. A realistic energy requirement for the PHEV occurs at approximately 7 repeated cycles (or roughly 8 kWh). At this location on the plot, we can see the best uniform pack is a 105.6 kg pack of Kia Soul cells (1P). The DDP and IP methods converged to a blended solution composed of a 7P pack of Tesla Model S cells and a 1P pack of Maxwell Ultracaps for a total mass of 62.6 kg. The PS benchmark yields a 91.6 kg blended solution made up of a 3P pack of Tesla Model S cells and a 1P pack of Ford Fusion SE cells. This means that even the benchmark blended solution can cut the weight of the pack by 14 kg while the DDP solution can reduce the mass of the pack by 43 kg.

In the same way Figure 4-29 presents the blended solutions for 4 consecutive




Figure 4-32: Gravimetric Power Density vs. Gravimetric Energy Density for Surveyed Cells (black) and DDP Solutions (red) for 7 Consecutive Cycles

cycles, Figure 4-32 shows the 18 uniform storage options (from Figure 4-1) alongside the blended solutions for each of the seven use cases for 7 consecutive drive cycles. We can see that the blended solutions have shifted to the right with respect to the 4 consecutive cycle case (in order to achieve higher energy density). Once again, we can see that the red solutions are interpolating between two different uniform solutions.

Figure 4-33 summarizes the results for all the use cases with 7 repeated drive cycles. It is clear that there is large room for improvement for the PHEV (C6) and also the Toyota Prius and Hybrid SUV (C1 and C5). The PHEV can achieve small gains using the PS solution, but much larger savings can be realized by utilizing the DDP solution (which requires complete knowledge of the drive cycle). The hybrid vehicle use cases offer larger gains for 7 cycles than for 4 cycles because the energy requirements have increased substantially to a point that is not realistic for full hybrid electric vehicles.




Figure 4-33: Least Massive Pack for DDP, Peak Shaving, and Uniform Solution Methods for 7 Consecutive Cycles

4.7.4 Electric Vehicles

Pure electric vehicle applications demand the highest amount of energy density and the lowest amount of power density. This is due to the fact that while the energy requirement becomes multiplied with each consecutive drive cycle, the power requirement remains constant. As the necessary energy grows, the resulting C rate diminishes, and this opens the door for higher energy density cells. Figure 4-34 illustrates this effect and shows the minimum mass packs found using the DDP, IP, PS, and uniform solutions for the Focus EV on a UDDS drive cycle.

Just from the linear integer programming lower bound in Figure 4-34, we can see that EV battery packs can only be downsized with blended battery packs when the overall vehicle range is relatively short. After 9 consecutive cycles, all of the solutions converge to the same uniform pack (composed completely of Tesla Model S cells). This occurs when the power output of the highest energy density cells is sufficient to supply all of the power for the given application. There are regions on the plot




Figure 4-34: Minimum Pack Mass for DDP, Integer Programming, Peak Shaving, and Uniform Solutions for the Focus EV on a UDDS Cycle

that show room for improvement with the blended solutions, but with fewer than 10 repeated cycles, the requirement is more indicative of a PHEV than an EV. This illustrates that when power and energy are the only constraints under consideration, EVs are not susceptible to weight reduction through blending of the battery pack.

Figure 4-35 shows the power density and energy density of the blended solutions for 10 consecutive cycles in red overlaid on top of the original 18 uniform storage options. We can see that the energy densities have increased with respect to Figure 4-32 due to the increased energy requirement of the 10 cycles. Each solution is approaching the Tesla Model S cell by interpolating between it and the Maxwell Ultracap. The Honda Insight pack remains the same, but this is because the energy requirement for each cycle is very low (highly charge sustaining), and so the minimal pack remains unchanged.

To summarize the results for 10 consecutive cycles, Figure 4-36 shows a bar plot with the minimal mass packs for the DDP, IP, PS, and uniform solutions for each of the 7 use cases. We can see that there is no room for improvement on the Focus EV




Figure 4-35: Gravimetric Power Density vs. Gravimetric Energy Density for Surveyed Cells (black) and DDP Solutions (red) for 10 Consecutive Cycles

use case (7), and very little room for improvement on the others with the exception of the Toyota Prius (1). With 10 repeated cycles, the 3.38 kWh energy requirement is unrealistically large for a full hybrid electric vehicle and would be more representative of a limited range PHEV.

4.8 Conclusion

In this chapter, I illustrated that tribridization through the use of blended battery packs has significant potential to reduce the overall weight of hybrid vehicles. This potential depends heavily on the specific parameters of the application in question. To examine which applications can benefit most from tribridization, I introduced three different methods for sizing blended battery packs early on in the design process. An integer programming method provides a lower bound, a dynamic programming based method yields the minimum feasible pack, and a peak shaving strategy gives a benchmark. I compared the results of all three of these methods to the best uniform




Figure 4-36: Least Massive Pack for DDP, Peak Shaving, and Uniform Solution Methods for 10 Consecutive Cycles

pack for 7 different use cases over a range of energy requirements.

Using this experiment, I showed that for the UDDS cycle, PHEVs have the largest potential for improvement. Additionally, some HEVs have potential for gains when their energy constraints are large enough. EVs, on the other hand, are not well suited for blended battery packs since the highest energy density cell can already supply the full power requirement without excess energy capacity. Regardless of the application, I found that a good rule of thumb is to plot the energy and power densities of the available storage devices and see if the application can be served better by linearly interpolating between two sources.

Perhaps the most useful insight of this analysis is that future knowledge of the drive cycle is necessary to realize the full benefits of tribridization. A simple strategy such as peak shaving can yield improvements over a uniform battery pack, but a control strategy that accounts for upcoming events (such as DDP) can offer much larger gains. For example, if a tribrid vehicle can anticipate a regenerative braking event before it occurs, it can ensure that there is enough room in the buffer to absorb

that pulse by increasing the draw from the buffer ahead of time. These types of gains will only be fully realized once vehicles can look outside the bounds of their conventional sensors and communicate with other vehicles on the road.

Chapter 5

Collaborative Learning for Improved Hybrid Energy Management

In this chapter, I discuss the idea of collaborative learning for the purpose of improving energy management in hybrid vehicles. To start, I introduce the high level ideas of connected vehicles and collaborative learning. Since connective vehicle technology is a new and fast-growing area, automotive manufacturers are still trying to discern the potential benefits and costs thereof. The costs and benefits are very application specific, so to understand the overall utility of connected vehicles, we must examine the applications on a case-by-case basis.

I will take a closer look at how these technologies can have an impact on the automatic start/stop system in micro hybrid vehicles. After introducing micro hybrid vehicles and giving an overview of the associated features and issues, I will quantify the potential to resolve these issues while maintaining the fuel economy benefits through the use of three different data sources and two machine learning techniques.

5.1 Connected Vehicles and Collaborative Learning

As vehicles become more and more connected to each other, infrastructure, and the cloud, and the costs of transmitting data continue to fall, applications of this connected vehicle data are becoming increasingly feasible [76]. Automotive manufac-




Figure 5-1: Connected Vehicle Technologies: (a) V2V, (b) V2I, (c) V2C

turers have begun to invest more in connected vehicle technology and are becoming increasingly interested in the implications of such large scale vehicle data collection. Yet, gathering massive amounts of data is not a trivial task, and the question still remains of which applications will benefit from connected vehicle data and how much value will result from such applications.

5.1.1 Connected Vehicle Technology

There are several main categories of connected vehicle technology: vehicle-to-vehicle (V2V), vehicle-to-infrastructure (V2I), and vehicle-to-cloud (V2C) [77]. Each of these methods allows data to disseminate with varying degrees of bandwidth, latency, and cost from one vehicle to other connected vehicles and surrounding connected infrastructure. In this thesis, we will not delve deeply into the consequences of one technology over another. Instead, we will assume the availability of current and historic vehicle data and analyze the potential benefits for certain applications.

5.1.2 Collaborative Learning

Collaborative learning traditionally refers to a situation in which two or more people attempt to learn something through a combined effort [78]. By working together, the objective can be reached much faster, and each individual benefits from the learning.

Current vehicle control systems rely solely on information obtained by the sensors in that vehicle. These control loops offer very precise information with extremely low latency about the state of the systems within that vehicle. Over time, learning algorithms can tune control parameters to compensate for gear lash or sloppy linkages; however, they have no insight into what is happening a few miles down the road or what was sensed in that same spot a few minutes ago by another vehicle.

In a connected automotive system, the meaning of "collaborative learning" is not so different from its pedagogical roots. Instead of each vehicle only having myopic access to the sensors within that same vehicle, every connected vehicle can share sensor data and store it to a common cloud for later reference. Then, learning algorithms can benefit from a much richer data set that extends well beyond the bounds of a single vehicle.

The applications for collaborative learning in automotive systems are practically limitless. For example, a vehicle that hits a pothole can log the position and magnitude of suspension deflection so that vehicles traveling down that same stretch of road can adjust their ride height and damping coefficients accordingly and notify municipalities that road work is necessary. A hybrid vehicle traveling down the interstate can foresee an upcoming accident and lower the state of charge of the battery pack to prepare for the regenerative braking event.

Ultimately, vehicles can use all of this contextual information to predict the future with some level of certainty and use this foresight to improve efficiency, comfort, performance, and the overall driving experience. In this chapter, we will focus on a few select applications with the goal of improving energy management through collaborative learning.

5.2 Micro Hybrids and Automatic Engine Start/Stop

Charge-sustaining hybrid vehicles can boost fuel economy in three main ways: supplementing engine power with an electric motor when the engine is least efficient, recovering energy through regenerative braking, and eliminating engine idling [16].




Figure 5-2: Micro Hybrid Topology

Micro hybrids exclusively target the potential gains of reducing engine idle time. In this way, micro hybrids are essentially the minimum viable hybrid vehicle.

At its core, a micro hybrid revolves around the automatic engine start/stop (AESS) system. As the name implies, this system turns the engine off when the engine would otherwise be idling, and then turns it back on again once the driver is ready to accelerate or some other system (e.g. power steering) requires the engine to start [79]. All the other micro hybrid components make this functionality possible while minimizing any negative impact on reliability, drivability, and overall user experience.

5.2.1 Micro Hybrid Structure

The goal of the micro hybrid topology is to eliminate all engine idling with minimal changes to the conventional internal combustion vehicle. While the necessary changes are not cost-prohibitive, there are certainly changes that must be made. A typical micro hybrid topology is shown in Figure 5-2.

Starter System

First of all, the starter motor system must be designed to handle the significantly increased cycle count. In addition to the typical 35,000 cold starts, an AESS system




Figure 5-3: Types of Starter Motors: (a) Ring Gear Starters and (b) Belt-Integrated Starter Generator

demands an estimated extra 300,000 warm start cycles [10]. Among other things, this includes a high durability starter motor and deep cycle lead acid battery.

Types of Starter Motors

There are two main types of starter motors used in AESS systems: ring gear starters and belt-integrated starter generators (BISG) [80]. Both of these are shown in Figure 5-3. Ring gear starters can be further subdivided into conventional and always engaged. Each one has strengths and weaknesses.

Conventional Ring Gear Starters Ring gear starters transfer torque to the main flywheel through a pinion gear on the starter motor shaft. This pinion gear is pushed into place by the starter solenoid, and so the starter motor pinion is only engaged with the ring gear on the main flywheel during startup. After startup, the pinion gear is retracted and can no longer transfer torque to the crankshaft.

Because of this, the engine can only be started after it has slowed down enough for the pinion gear to engage the ring gear. Despite the challenge that this limitation poses for an AESS system (discussed further in Section 5.2.3) ring gear starters are still commonly used because of their relatively low cost [80].

Belt-Integrated Starter Generator A belt-integrated starter generator works well for AESS systems since it is always engaged with the engine via belts and pulleys. This allows the starter to apply forward torque at any speed and can also perform a mild form of regenerative braking. Because of these additional features (and the extra cost associated with them) vehicle equipped with a BISG are generally classified as mild-hybrids [81].

Always Engaged Ring Gear Starters Always engaged ring gear starter motors have been proposed to allow the application of torque even before the engine has completely stopped [82]. This is similar to the BISG but uses gears instead of belts and pulleys. This solution has the drawback of additional noise due to the always engaged pinion gear.

Electrified Accessories

Since traditional accessories such as the A/C compressor and power steering are run off the serpentine belt that stops moving when the engine is stopped, a new solution is required if the operation of these accessories is going to be maintained at all times. The simplest way to always ensure operation of the accessories is to run a separate accessory motor either for each one or some combination using belts and pulleys.

In addition to allowing these accessories to run while the engine is stopped, there are secondary benefits of decoupling them from the engine. When they are not being used, they no longer apply extra frictional loads to the engine. Each accessory can be run at the most efficient RPM for any given time and is not linked to the engine speed through some gear ratio. Accessories can even be loaded opportunistically to help bring the engine into more efficient operating ranges as the electrical load on the BISG or alternator increases [83].

5.2.2 Estimated Fuel Economy Gains

The actual fuel economy benefits from AESS depend heavily on the use case. The more idling the vehicle would experience, the more room there is for improvement.




Figure 5-4: Idle Time for UDDS Cycle

For the EPA cycles such as the UDDS cycle shown in Figure 5-4, the fuel economy benefits range from 4 – 10% [84] - [85]. In Section 5.3.1, we will discuss idle time reduction further based on real world data.

5.2.3 Problems with Automatic Engine Start/Stop

The potential fuel economy gains from AESS are substantial in theory, but when the vehicle operates in the real world, a number of problems arise that have slowed down the adoption of AESS. While regulations require that the vehicles start with AESS enabled by default, they do not prevent the manufacturers from installing buttons to disable AESS, or discourage drivers from turning it off at the beginning of each trip. Unless these problems are addressed, there will always be drivers who disable AESS and circumvent any fuel economy gains that would have otherwise been achieved.




Figure 5-5: Performance Comparison of BMW X1 with AESS enabled and Disabled (Gathered Using OBDII Data)

Performance Comparison

While vehicles with AESS are tuned to minimize starting time, there is still a noticeable lag between when the engine begins to start and when it finally reaches a high enough speed to produce usable torque. Lag times range from vehicle to vehicle and from stop to stop. Typical startup delays range between 0.5 and 1 second [10], [86]. Figure 5-5 illustrates some acceleration times we gathered for a BMW X1 with AESS enabled and disabled. We observed delays in excess of 1 second.

Change-of-Mind Events

The AESS hesitation is exacerbated during a "change-of-mind" event, in which the driver desires to accelerate as soon as the engine has slowed beyond the point of no return in a micro hybrid equipped with a ring gear starter (around 400 RPM) [80]. This is considered the worst-case-scenario in terms of startup delay and causes the feeling of a distinct disconnect between the pedal and the motion of the vehicle.




Figure 5-6: Critical Idle Duration Before AESS Benefits Fuel Economy

Despite the label, a change-of-mind event does not only occur when a driver changes his or her mind but often occurs in the natural flow of traffic, as we will show in the sections that follow. These events would otherwise be relatively short engine idles in a conventional vehicle but are none-the-less referred to this way.

Effect on Fuel Economy

In terms of fuel economy gains, a change-of-mind event offers the least potential savings and can even hurt fuel economy with respect to idling the engine depending on the duration of the would-be idle. This is due to the fixed fuel cost of starting the engine versus the variable fuel cost of idling the engine. Figure 5-6 illustrates this relationship. According to [10], even under the best conditions for a micro hybrid with a ring gear starter, the engine would have to idle for more than 2 seconds before the fuel consumed during the engine start process is surpassed by the fuel that is saved due to idle time reduction. Other studies have reported this critical threshold as low as 0.4 seconds [87] or as high as 7 seconds for some engines [88].

Stopping the engine for any idle time shorter than this critical threshold will cost as much fuel as idling for the amount of time less than the threshold. For example, when the critical threshold is 2 seconds as in Figure 5-6, a 0.5 second AESS event actually increases the effective idling time (with respect to fuel consumption) by 1.5 seconds.

User Experience

The most prominent annoyances to customers are noise, vibration, and harshness, collectively referred to as NVH. As the NVH in vehicles has improved over the years, drivers have become even more sensitive to new sources of noise and vibration. In a conventional vehicle where the engine is started and stopped once per trip, the NVH experienced by the driver during these events has little consequence on the overall comfort and experience of the rest of the journey.

In a micro hybrid, where the engine is started and stopped frequently throughout the trip, the harshness of the start and stop can disrupt an otherwise smooth experience. Because of the importance of NVH, entire studies have been performed to characterize and optimize the NVH while maintaining drivability and launch-readiness [89] - [90].

Shorter and more frequent stops subject the driver to the most additional NVH with the smallest potential fuel economy gains (and the possibility of increasing fuel consumption). Unless NVH can be reduced to a level below which drivers are not disturbed by the constant start and restart of the engine, they will most likely continue to disable the system instead of benefit from the fuel economy gains.

Disabling Factors

Although AESS systems perform well during EPA testing, there are a large number of real-world factors that cause the system to be disabled in practice. For example, AESS does not function if the ambient temperature is close to or below freezing or above some high threshold (usually about $95^{\circ}F$), or if the engine temperature is too high or too low (still warming up), or if the cabin temperature is not at the target temperature, or if the steering wheel is angled or moved, or if the battery voltage is too low, or if the road grade is too steep, or if the windshield is fogged, or if the car is in reverse, or if the car is in stop and go traffic (vehicle must exceed a certain speed between engine stops).

All of these requirements have good reasons behind them. They are attempting to

ensure that the vehicle functions smoothly, and that any driver annoyance is kept to a minimum while still capturing as much of the fuel economy gains as possible during the EPA tests. Interestingly enough, none of the factors listed above are triggered during fuel economy testing. Even the ambient temperature requirement alone can eliminate the utility of an AESS system for months out of the year depending on the local climate.

Short Idles

Short idles are at the root of most of the problems associated with AESS. They introduce change-of-mind events, which produce the largest lag between driver-demanded acceleration and actual acceleration. They have the lowest potential for fuel economy gains and can even be detrimental to fuel economy when the stop is below the critical time threshold (approximately 2 seconds). They produce the greatest amount of NVH for a given amount of idle reduction. Ultimately, if short idles can be predicted, they should be avoided entirely by not stopping the engine.

5.3 JLR Data Set

We conducted research with Jaguar Land Rover to start answering the question of how valuable large amounts of data can become with applied to the vehicle powertrain and energy management. As their part, JLR has provided access to a data set that they have been building up over the past few years.

5.3.1 JLR Data Statistics

To filter out any extraneous trips, we used a subset of the JLR data set based on a number of criteria. In order to be classified as a valid trip, we required a minimum duration of 6 minutes, minimum distance traveled of 5 kilometers, maximum GPS time discontinuity of 6 minutes, maximum GPS spacial discontinuity of 5 kilometers, a maximum vehicle speed of 150 kilometers per hour.




Figure 5-7: Histogram of Trip Durations from the JLR Data Set

In addition, we added the requirement to the filter that all of the following parameters be reported on the vehicle networks at least once during a given trip: engine speed, AESS available, AESS button status, AESS status, AESS stop-in-progress, and AESS switch status. This ruled out any vehicles without AESS and ensured the validity of any data processing assumptions.

After filtering, the JLR data set contains 79 vehicles, 6144 trips, for a total of 3,176 hours of driving data. The minimum, average, and maximum trip duration after filtering were 6 minutes 23 seconds, 31 minutes 1 second, and 2 hours 4 minutes respectively. Figure 5-7 shows a histogram of the trip durations in minutes after filtering.

JLR Idling Statistics

Figure 5-8 illustrates the engine speed and stop-in-progress data from a typical trip in the JLR data set. When the stop-in-progress parameter is 1, an AESS event is in progress, and when it is 0, there is no stop in progress. Using this parameter, we




Figure 5-8: Sample Trip from JLR Data Set

can evaluate how often and for how long AESS events are occurring throughout the entire JLR data set.

For vehicles reporting the AESS stop-in-progress parameter, 20,146 unique idle events were recorded. The minimum, average, and maximum idles times were 0.1 seconds, 13.7 seconds, 99.8 seconds respectively when filtering out any idles over 100 seconds long. Any idles longer than 100 seconds are considered to be parking events and not a part of traffic flow.

Figure 5-9 shows a histogram of the JLR data idle times. It is clear that the idle times follow an asymmetrical distribution with a bias towards shorter idles. This plot appears to resemble an Inverse Gaussian distribution.

The empirical CDF plot in Figure 5-10 illustrates what fraction of idle events are below a given idle time. This illustrates that the majority of idle events are less than 8 seconds, and 18 % of idles are less than 2 seconds in duration.

Over all the filtered trips, the AESS system indicated a stop-in-progress 6.27% of the time, the engine was idling 16.82% of the time, and the AESS switch status was reported to be disabled 0.89% of the time. The total duration of idles (real-world idle fraction) was 23.1% of the time. Overall system effectiveness is therefore 44.7%.




Figure 5-9: Histogram of Idle Durations from the JLR Data




Figure 5-10: Empirical Cumulative Distribution Function (CDF) of Idle Durations from the JLR Data

Change-of-Mind Detriment to Fuel Economy

As discussed in Section 5.2.3, there is a fixed fuel cost of starting the engine and a variable cost of idling the engine. Due to this relationship, there is a critical time at which the fuel consumed through idling is equal to the fuel used to start the engine. Turning the engine off for an idle that is less than this critical threshold results in more fuel being consumed than if the engine had idled for that same duration. Because of this, any of the potential fuel savings from idle time below the threshold cannot be recovered through AESS.

Since this critical threshold takes on a range of values depending on the application, we have used the JLR data set to show the cumulative fraction of the total idle time that is below a given critical threshold in Figure 5-11. From this figure, we can see that approximately 7% of total idle time is accounted for in the first second of idling. With a critical threshold of two seconds, the fuel savings from about 13.5% of all idle time is not recoverable. At three seconds, over 19% of the fuel consumed through idling cannot be saved using AESS. Understanding this threshold dependent efficiency limit is important when evaluating different AESS strategies.

Comparison to Mercedes Benz Off-cycle Credit Idling Study

For model year 2012 to 2016 vehicles, the EPA established an option for manufacturers to generate additional carbon credits through the use of technologies that cut more carbon dioxide than is captured in the 2-cycle test procedure employed by the EPA. These additional credits are known as “off-cycle” credits [91].

Real-world Idle Fraction By default, the EPA uses a real-world idle fraction of 13.76% when crediting all AESS systems. Mercedes Benz was awarded additional off-cycle credits by proving that the actual real-world idle fraction was closer to 23%. In a Mercedes specific study of 29 vehicles, the idle fraction was found to be 23.83%. They also used data collected from 17,484 Mercedes vehicles by Progressive Insurance to find an idle fraction of 23.9% [91]. This agrees generally with the 23.1% idle fraction found for the JLR data.




Figure 5-11: CDF of % Idle Time Due to Idles Less than the Threshold

AESS System Effectiveness The real-world idle fraction accounts for how much time the engine is idling but not how often it can actually be shut off by the AESS system. Normally, the EPA assumes a system effectiveness of 87.75%. After considering all the factors listed in Section 5.2.3 aside from driver disablement, Mercedes estimated the weighted AESS effectiveness of 60-65% for their vehicles. The EPA reviewed this claim and concluded that 52% system effectiveness was more realistic [91].

Driver Disablement As discussed in Section 5.2.3, drivers can easily disable the AESS system with the push of a button. An example of a trip where the driver disables the AESS system is shown in Figure 5-12. When the button status parameter is at 2, the AESS system is enabled, and when it goes to 5, the driver has disabled it. This trip shows four AESS events before the driver disables the system for the rest of the journey.

In the Mercedes Benz off-cycle credit application, they estimated driver disablement at less than 1% of the time based on European survey data [91]. This agrees




Figure 5-12: Sample Trip with Driver Disablement from JLR Data Set

well with the 0.89% driver disablement found in the JLR data. The EPA used a more conservative estimate of 11% for driver disablement yielding an overall system effectiveness of 46%, which agrees well with the 44.7% found in the JLR data.

5.3.2 JLR Data Set Limitations

The JLR data set has a large quantity of data, but it only has the data that was available on the vehicle networks at the time it was gathered. In other words, a parameter that is important in retrospect, may not have been broadcast at the time, and is therefore unavailable.

Of particular consequence to this study (as will be shown in the rest of this chapter) is the relative position and speed of the vehicle in front of the vehicle in question. Although the vehicles in the JLR fleet are equipped with parking sensors and radar modules for active cruise control, and the position and speed of the vehicle in front could be obtained using these sensors, this data was not logged due to the nature of the vehicle network architecture.

5.4 Traffic Simulation Data Set

In order to understand traffic flow at a more fundamental level and determine the best-case feasibility of predicting idle time, we decided to start with traffic simulation. If it is not possible to make reliable predictions in a controlled simulation, it is not likely to work in the real world. This also allowed for the relatively quick exploration of the usefulness of various data sources without having to go through the process of real-world data collection for each one.

5.4.1 Traffic Modeling with the Intelligent Driver Model

To model the behavior of individual cars on the road, we created a traffic model in MATLAB based on the Intelligent Driver Model (IDM) [92], which is a microscopic traffic flow model. In this time-continuous car-following model, the acceleration of each vehicle is determined based on the current speed of that particular vehicle, the bumper-to-bumper distance and relative velocity to the car in front, and driver parameters such as desired speed and acceleration.

Intelligent Drive Model Definition and Parameters

For a given vehicle (i), the 1-dimensional position of that vehicle is denoted as x_i , and so the velocity of vehicle i is represented by the ordinary differential equation

$$v_i = \dot{x}_i = \frac{dx_i}{dt}. \quad (5.1)$$

Taking another derivative yields acceleration

$$a_i = \dot{v}_i = \frac{dv_i}{dt}, \quad (5.2)$$

which is modeled as

$$\frac{dv_i}{dt} = A_i \left[1 - \left(\frac{v_i}{V_{i0}} \right)^\delta - \left(\frac{s^*(v_i, \Delta v_i)}{s_i} \right)^2 \right], \quad (5.3)$$

Table 5.1: IDM Parameter Values

Parameter	Value
A_i	2 [ms^{-2}]
V_{i_0}	20 [ms^{-1}]
δ	4
s_{i_0}	2 [m]
L_i	4 [m]
T	1 [s]
B_i	2 [ms^{-2}]
$B_{i_{max}}$	6 [ms^{-2}]

where A_i is desired acceleration in traffic, V_{i_0} is desired speed on a free road, δ is the acceleration exponent, Δv_i is the relative speed to the car in front ($v_{i-1} - v_i$), s_i is defined as

$$s_i = x_{i-1} - x_i - L_{i-1} \quad (5.4)$$

where x_{i-1} is the position of the car in front of car i , L_{i-1} is the length of the car in front, and s^* is defined as

$$s^*(v_i, \Delta v_i) = s_{i_0} + \max \left[0, \left(v_i T + \frac{v_i \Delta v_i}{2\sqrt{A_i B_i}} \right) \right] \quad (5.5)$$

where s_{i_0} is the minimum desired bumper-to-bumper distance for car i , T is the desired safety time headway to the car in front, B_i is the comfortable desired braking deceleration, and $B_{i_{max}}$ is the maximum allowed breaking deceleration.

For this simulation, the vehicle parameters were chosen to be representative of a typical driver in city traffic. The IDM parameter values utilized in this study are listed in Table 5.1.

Simulation Algorithm

The Intelligent Drive Model explained in Section 5.4.1 governs the acceleration of each vehicle in the simulation, but the overall simulation requires further explanation.

For simplicity, we chose to create a one-dimensional simulation composed of both vehicles and traffic lights as shown in Figure 5-13. Whenever a vehicle reaches the



Figure 5-13: 1-D Traffic Simulation with Traffic Light

Table 5.2: Traffic Simulation Parameter Values

Parameter	Value
N_c	40 <i>cars</i>]
L_r	800 [<i>m</i>]
N_l	3 <i>lights</i>
T_{rg}	40 [<i>s</i>]
T_y	3 [<i>s</i>]
T_{sim}	48 [<i>hours</i>]
τ	0.1 [<i>s</i>]

end of the road, it returns to the beginning of the road ($x = 0$) forming a virtual ring road. The number of vehicles (N_c), length of the road (L_r), number of traffic lights N_l , red and green light duration T_{rg} , yellow light duration T_y , time step τ , duration of simulation T_{sim} , and IDM parameters are all inputs to the simulation. The values of these overall simulation parameters are listed in Table 5.2.

Initialization, $t = 0$ The beginning of the simulation is seeded with some random components to ensure each result is unique. The last car (N_c) is placed at position $x_{N_c} = 0$ with a speed of $v_{N_c} = 15 [ms^{-1}]$. Each subsequent car is placed a random distance between 5 and 10 meters in front of the car before it with a random speed between 15 and 17 [ms^{-1}]. So the initial position of each car is

$$x_{i_0} = x_{(i+1)_0} + X \quad (5.6)$$

where X is a discrete random variable with a uniform distribution between 5 and 10. The initial speed is

$$v_{i_0} = V \quad (5.7)$$

where V is a discrete random variable with a uniform distribution between 15 and 17.

The lights are also initialized semi-randomly to attempt to simulate a more realistic environment without perfect synchronization. Lights are spaced evenly across the total road length at position x_j . The red and green light durations for each light (j) are set to the base duration plus a random component,

$$T_{rg_j} = T_{rg} + Q_{rg}, \quad (5.8)$$

where Q_{rg} is a discrete random variable with a uniform distribution between 0 and 15, and the same is done for the yellow light duration,

$$T_{y_j} = T_y + Q_y, \quad (5.9)$$

where Q_y is a discrete random variable with a uniform distribution between 0 and 2.

Each light is also started at a random point in the cycle,

$$T_{start_j} = Q_{start}, \quad (5.10)$$

where Q_{start} is a discrete random variable with a uniform distribution between 0 and $2T_{rg_j} + T_{y_j}$.

Simulation Main Loop, $t > 0$ After initialization, the simulation proceeds at increments of τ . At the start of each step, the first step is to update the colors of the lights based on the current time, the start time of the light, and the red, green, and yellow light durations.

A red light acts like a stationary car at the position of the light, and a green light has no effect on the flow of traffic. Treating the red light like a stopped car assumes that the yellow light duration (T_{y_j}) gives the car enough time to stop at its maximum desired speed V_{i_0} .

When a car encounters a yellow light, it must decide if it should stop or go through

the yellow light. Just treating the yellow light like a red light would result in cars attempting to stop when it is physically impossible for them to do so.

For each vehicle (starting with $i = 1$), the desired acceleration is calculated using the Intelligent Driver Model (discussed in Section 5.4.1) based on the parameters of the vehicle in front ($a_{i_c}(t)$). Note that for the first vehicle in line ($i = 1$), the vehicle in front is the last vehicle in line ($i = N_c$) since this is a ring road, but otherwise it is $i - 1$.

If all the lights are green, then the desired acceleration is just set to the acceleration due to the car in front,

$$a_i(t) = a_{i_c}(t). \quad (5.11)$$

The desired acceleration for a red light ($a_{i_r}(t)$) is calculated using the IDM as though the red light were a stationary vehicle at the position of the light. The acceleration due to the light is then set equal to the red light acceleration

$$a_{i_l}(t) = a_{i_r}(t). \quad (5.12)$$

In the case of a yellow light, the minimum stopping distance ($d_{stop}(t)$) is calculated base on the maximum braking deceleration and the current speed,

$$d_{stop}(t) = \frac{v_i(t)^2}{2B_{i_{max}}}. \quad (5.13)$$

Then the acceleration due to the light is defined by the following piecewise function,

$$a_{i_l}(t) = \begin{cases} a_{i_r}(t) & \text{if } d_{stop}(t) \leq x_j - x_i \\ a_{i_c}(t) & \text{if } d_{stop}(t) > x_j - x_i \end{cases} \quad (5.14)$$

which ensures that the car will stop for the yellow light as though it were a red light if it can physically do so.

With the desired acceleration calculated for the car in front and the next red or yellow light, we take the minimum of these two accelerations to find the acceleration

of car i .

$$a_i(t) = \begin{cases} a_{i_c}(t) & \text{if } a_{i_c}(t) \leq a_{i_l}(t) \\ a_{i_l}(t) & \text{if } a_{i_c}(t) > a_{i_l}(t) \end{cases} \quad (5.15)$$

If the desired acceleration $a_i(t)$ exceeds the maximum allowed acceleration or deceleration, then the value is saturated to the maximum.

With the acceleration known for each vehicle at the current time step, the resulting speeds and positions are calculated assuming constant acceleration for the duration of the time step. The velocity is

$$v_i(t) = v_i(t - \tau) + \tau a_i(t), \quad (5.16)$$

and the resulting position is,

$$x_i(t) = x_i(t - \tau) + \tau v_i(t) + \frac{a_i(t)\tau^2}{2}. \quad (5.17)$$

Whenever the velocity ($v_i(t)$) of a vehicle is 0, that vehicle is considered to be idling. The durations and positions of each idle are logged for later analysis.

5.4.2 Traffic Simulation Data Statistics

For a 48 hour simulation, the forty simulated vehicles logged 59,323 idle events. The minimum, average, and maximum idles times were 0.1 seconds, 26.67 seconds, and 59.1 seconds respectively.

A histogram of idle times is shown in Figure 5-14. The distribution here looks substantially different from the real data collected in the JLR data set. From 0 to 35 seconds, the distribution appears more or less uniform until there is a spike around 38 and 47 seconds, which corresponds to the red and green light durations for the simulation.

Figure 5-15 shows the empirical CDF for the simulation and illustrates what fraction of idle events are below a given idle time. This distribution highlights how much




Figure 5-14: Histogram of Idle Durations from the Simulated Data

more slowly short idles accumulate with respect to the real-world JLR data. There are many factors that the simulation does not capture which could cause the discrepancy between the simulation and reality. Stop signs, pedestrian crossings, turning right on red (causing all the cars to move up), creeping forward in heavy traffic, and other driver behaviors can all cause more short idles. While the simulated distribution is different from the real world distributions, it still serves as a good test case.

5.5 Cambridge Data Set

Because the JLR data set is limited to the parameters that were broadcast on the vehicle networks and does not necessarily include all the pertinent information, we also gathered a local idle data set in Cambridge that includes the parameters we required. The goal of our local data collection was to mimic the data available in the simulation for the purposes of validating any hypotheses that came out of the simulation.




Figure 5-15: Empirical CDF of Idle Durations from Simulated Data

As the Intelligent Driver Model shows, the behavior of a driver and the acceleration of a given vehicle can be modeled by observing the relative speed and position of the vehicle directly in front of that given vehicle. This fact compelled us to focus on gathering location and speed data of a single vehicle as well as the speed and position of the vehicle in front of it. This allows us to compare the model more directly with real-world data.

5.5.1 Experimental Method

To gather idle time data in Cambridge, we used two vehicles equipped with custom data acquisition rigs. Each rig consisted of a differential GPS unit for position, OBD-II logger for vehicle speed and engine RPM, two single-point LIDAR and sonar units for the distance to the car in front (see Figures 5-16 and 5-17), and a laptop for processing and saving the data.

In order to parallel the simulation and limit the number of driving hours necessary to acquire sufficient data density, a small circuit was selected around MIT campus with



Figure 5-16: Windshield Mounted LIDAR and Sonar Rig (1 of 2)

four traffic lights. Unlike the simulation, this circuit also included four crosswalks. A map view of the circuit with markers for idle locations is shown in Figure 5-18.

The chase car pictured in Figure 5-17 acts as the vehicle i while the lead car acts as vehicle $i - 1$ from Section 5.4.1. When the vehicle speed parameter of vehicle i read from the OBD-II logger reaches zero, the position and time from the GPS are logged, and position and speed of vehicle $i - 1$ are logged by both the LIDAR and sonar sensors on vehicle i as well as the OBD-II logger in vehicle $i - 1$. When the speed of vehicle i goes above zero again, the time of the idle is logged.

5.5.2 Cambridge Data Statistics

Over the course of approximately 6 hours, our two-vehicle data convoy logged 375 unique idle events around the Cambridge circuit (pictured in Figure 5-18). The minimum, average, and maximum idles times were 0.2 seconds, 7.35 seconds, 29.8 seconds respectively.



Figure 5-17: Chase Car with LIDAR and Sonar Mounted on Windshield Near Side Mirrors



Figure 5-18: Map View of Cambridge Idle Data, Greener Markers Correspond to Shorter Idles, Bluer Markers Correspond to Longer Idles

A histogram of idle times is shown in Figure 5-19. It is clear that the idle times follow an asymmetrical distribution with a bias towards shorter idles in the same way as the JLR data.

The empirical CDF plot in Figure 5-20 shows what fraction of idle events are below a given idle time. This illustrates that the majority of idle events are less than




Figure 5-19: Histogram of Idle Durations from the Cambridge Data




Figure 5-20: Empirical CDF of Idle Durations from the Cambridge Data

4 seconds, and over 35% are shorter than 2 seconds.

5.6 Smart/Stop: Making AESS Seamless

A smarter solution is required to help alleviate the issues that have plagued automatic engine start/stop systems (Section 5.2.3). In this section, we introduce the idea of a Smart/Stop system that can learn how to predict a change-of-mind event before it occurs, and then temporarily disable the AESS system, thus eliminating what could otherwise be an unpleasant experience for the driver and a waste of fuel.

5.6.1 Short Idle Classification for Smart/Stop

The fundamental requirement of the Smart/Stop system is to be able to classify short and long idle events **before** they happen. In order to do this, we will draw on current and historical contextual data.

5.6.2 Confusion Matrix

A confusion matrix is a useful tool in understanding the performance of a particular classification method. An example confusion matrix is shown in Table 5.3. This matrix will show the number of times the algorithm classifies each event correctly or incorrectly.

It is important to separate it out in terms of correct short classifications (true positive), incorrect short classifications (false positive), correct long classifications (true negative) and incorrect long classifications (false negatives). Ideally, every prediction is either a true positive or a true negative, but in reality, there will be misclassifications.

It is useful to understand the consequence of a false positive versus a false negative classification for the specific application. In the case of a Smart/Stop system, a false positive will result in the engine idling for a long idle, when the AESS system would have otherwise shut the engine off. This results in lower fuel economy with respect to the baseline (standard AESS). A false negative, on the other hand, would lead to the engine stopping for a short idle, which could lead to a change-of-mind event. This is equivalent to the baseline case.

Table 5.3: Example Confusion Matrix

Actual Class	Classified Short	Classified Long
Short	# True Positive	# False Negative
Long	# False Positive	# True Negative

5.6.3 Learning Feature Selection

Before we can attempt to classify short and long idles, we must identify a list of parameters to include in the learning vector. We decided to start by focusing in on the parameters that are most fundamentally related to idle time.

Using the IDM to Understand Important Features

By manipulating the Intelligent Driver Model, we will find that the speed and acceleration of the vehicle in front of the vehicle in question are the most critical factors. When the vehicle in question stops (and is just starting an idle event), we know that the speed of the vehicle is zero by definition. Plugging $v_i = 0$ into Equation 5.3 yields the acceleration of a stationary vehicle as follows

$$a_i(v_i = 0) = A_i \left[1 - \left(\frac{s^*(v_i, \Delta v_i)}{s_i} \right)^2 \right], \quad (5.18)$$

where A_i is the desired acceleration and s_i is the bumper-to-bumper distance to the car in front (Equation 5.4). Plugging in $v_i = 0$ into Equation 5.5 simplifies the expression further to

$$a_{i_0} = A_i \left[1 - \left(\frac{s_{i_0}}{s_i} \right)^2 \right], \quad (5.19)$$

where s_{i_0} is the minimum desired bumper-to-bumper distance.

It is important to note that after the vehicle stops, any negative acceleration is saturated to zero. The vehicle inevitably overshoots the desired bumper-to-bumper distance, and this prevents it from moving backward to correct. With this in mind, in order to predict how long the vehicle will be idling, we must find the critical point at which the stationary acceleration (a_{i_0}) exceeds zero. In other words, the time at which the following is true

$$\left(\frac{s_{i_0}}{s_i}\right)^2 = 1. \quad (5.20)$$

Of course, the desired bumper-to-bumper distance is constant with respect to time, but the actual distance will increase as the vehicle in front moves forward according to Equation 5.4. So the condition is met when the actual bumper-to-bumper distance exceeds the desired distance. This means that the idle time for a stopped vehicle with a moving vehicle in front should be the time that it takes the vehicle in front to move a fixed distance $(s_{i_0} - s_{i_I})$, where s_{i_I} is the bumper-to-bumper distance at the start of the idle.

If we assume that vehicle in front is accelerating a constant rate with an initial velocity, $v_{(i-1)_I}$, we can model the change in bumper-to-bumper distance as

$$\Delta s_i = \frac{1}{2}a_{i-1}t^2 + v_{(i-1)_I}t, \quad (5.21)$$

which allows us to estimate the idle time by plugging in $\Delta s_i = s_{i_0} - s_{i_I}$ and solving for t using the quadratic formula

$$t_I = \frac{-v_{(i-1)_I} + \sqrt{v_{(i-1)_I}^2 + 2a_{i-1}(s_{i_0} - s_{i_I})}}{a_{i-1}}. \quad (5.22)$$

From Equation 5.22, we can see that the idle time of a vehicle with a moving car in front depends on the initial velocity of the car in front, the acceleration of that car, and the amount of distance the vehicle needs to cover to reach the desired bumper-to-bumper distance. This equation is only valid to estimate the time from when the vehicle in front starts moving until the vehicle in question starts to move (not valid when the vehicle in front is stationary). However, it can be applied recursively for a queue of stationary vehicles.

Potential Features vs. Idle Time from Simulation

To confirm if the velocity and acceleration of the vehicles in front of the current vehicle are correlated with idle time, we will take a closer look at the simulated traffic data.




Figure 5-21: (a) Speed of Vehicle $i - 1$ at Start of Idle and (b) Acceleration of Vehicle $i - 1$ at Start of Idle vs. Actual Idle Time

Figure 5-21a shows the speed of vehicle in front of vehicle i at the start of the idle versus the duration of the idle. The vast majority of cases show a speed of zero for all idles over two seconds. For idles less than two seconds, on the other hand, we see the speed approach 5 kph as the idle duration moves toward zero.

This trend agrees well with the projected idle times based on $v_{(i-1)_I}$ from Equation 5.22 (overlaid on top of the simulated data with a dashed line in Figure 5-21a). This projection assumes a constant acceleration of $0.2 [ms^{-2}]$ and $s_{i_0} - s_{i_I} = \frac{1}{3} [m]$. Of course, the acceleration is not going to be completely constant, but using an average acceleration illustrates the dependence of idle time on the initial speed of the vehicle.

The acceleration of the vehicle $i - 1$ shows a similar trend in Figure 5-21b. For idles longer than two seconds, the acceleration of vehicle $i - 1$ is generally zero. For idles less than two seconds, the acceleration of vehicle $i - 1$ approaches a value slightly greater than $1 [ms^{-2}]$. Outliers are due to situations in which vehicle $i - 1$ has already passed the light at which vehicle i is stopped.

Figure 5-22 shows the speed and acceleration of the second vehicle in front of the vehicle in question ($i - 2$). Similar trends to Figure 5-21 can be observed for both speed and acceleration. The main difference is that idle time at which the speed and acceleration converge to zero has approximately doubled to about four seconds.




Figure 5-22: (a) Speed of Vehicle $i - 2$ at Start of Idle and (b) Acceleration of Vehicle $i - 2$ at Start of Idle vs. Actual Idle Time

As mentioned in Section 5.6.3, this is due to the “recursive” nature of the traffic queue. Vehicle i must wait for vehicle $i - 1$ to move which must wait for vehicle $i - 2$ to move and so forth down the line. The further away the first moving vehicle is in the queue, the longer vehicle i must wait.

Potential Features vs. Idle Time, Comparison with Cambridge Data

To see if the trends experienced in the simulation and discussed in Section 5.6.3 are just artifacts of the simulation or have some bearing on reality, we will compare the simulated data to the real-world data gathered in Cambridge.

Figure 5-23 shows the speed of vehicle $i - 1$ at the start of the idle versus actual idle time for both the simulated and real-world data. In both data sets, we see that after a certain idle time threshold, the speed of vehicle $i - 1$ converges to zero. For the Cambridge data set, this point is somewhere between 2 and 3 seconds.

The very tight nature of the simulated data with respect to the real-world data is most likely due to the fact that real drivers don’t always accelerate at the same rate for a given situation (unlike the IDM simulated drivers). This large spread of speeds for a given idle time with the real-world data will make prediction of actual idle time more difficult based only on speed. However, if the goal is to classify the idle as short or long, this variance will have less of an effect.




Figure 5-23: Speed of Vehicle $i - 1$ at Start of Idle vs. Actual Idle Time for Simulated and Cambridge Data

Learning Feature Selection Conclusions

By checking the insights gathered from the Intelligent Driver Model with the simulated data and confirming them with real data in the above sections, we can be confident that the speeds and accelerations of the vehicles in front of vehicle i are useful for idle time prediction. We saw that for the real-world data, additional noise in the speed data suggests that more factors should be considered to capture all the variance. This could make other parameters such as location useful for differentiating between different idle durations.

Ideally, we could also examine the acceleration of vehicle $i - 1$ as well as the speed and acceleration of vehicle $i - 2$ for the Cambridge data set. Unfortunately, we did not measure acceleration directly, and taking the derivative of the noisy speed data signal only amplifies the noise, which makes it largely unusable. As for the second vehicle in front, we only had access to two vehicles and two drivers, so we could not include a third vehicle in the data collection convoy.

5.6.4 Classification Methods

In this study, we will explore the utility of two different classification methods. We will utilize the Naive Bayes Classifier as a baseline and kNN for enhanced performance.

Naive Bayes Classifier

As a baseline method, we used a Naive Bayes Classifier to predict if an idle will be short or long using only the features in the learning vector. The Naive Bayes Classifier is based on Baye's theorem which defines the probability of event A given event B in terms of the probability of B given A , and the independent probabilities of each [93]

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}. \quad (5.23)$$

This theorem can be expanded into a classifier for a short idle given a learning vector \mathbf{x} of length n as follows

$$p(short|\mathbf{x}) = \frac{p(\mathbf{x}|short)p(short)}{p(\mathbf{x})}. \quad (5.24)$$

where $p(short|\mathbf{x})$ is the posterior, $p(\mathbf{x}|short)$ is the prior, $p(short)$ is the likelihood, and $p(\mathbf{x})$ is the evidence. Since the evidence does not depend on if the idle is long or short, it can be considered constant and effectively ignored. By using the chain rule, Equation 5.24 can be expanded to

$$p(short|\mathbf{x}) = \frac{p(short)}{p(\mathbf{x})} \prod_{k=1}^n p(x_k|short). \quad (5.25)$$

Similarly, the posterior probability for a long idle can be defined as

$$p(long|\mathbf{x}) = \frac{p(long)}{p(\mathbf{x})} \prod_{k=1}^n p(x_k|long). \quad (5.26)$$

To predict if an idle will be short or long, the posterior probabilities for both long (Equation 5.26) and short (Equation 5.25) are calculated for the given learning vector, and the greater of the two probabilities is taken as the prediction.

***k*th Nearest Neighbor Classification**

For improved performance, we used a *k*th Nearest Neighbor Classifier using a Euclidean distance metric. As the name suggests, this classifier finds the *k* nearest neighbors with the closest Euclidean distance in the training set to the test vector for the idle in question [94]. If more of the neighbors are short idles, the classifier predicts a short idle. If more of the neighbors are long idles, the classifier predicts a long idle.

5.6.5 Cross-Validation

We tested each classifier using an exhaustive leave-one-out cross-validation. One by one, each idle is removed from the set and a prediction is made based on the remaining idles. That idle is then replaced, and the process is repeated. While this method is computationally intensive, it ensures no overfitting and minimizes any wasted data. This is particularly important for the Cambridge data set where there are only 375 idles total.

5.6.6 Idle Time Prediction

Predicting Idle Time with Simulated Data

Figures 5-24 - 5-27 show the results of using a kNN classifier on the simulated traffic data set for various learning vectors. Each plot shows the classification accuracy for short idles (a) and for long idles (b) at three different short idle thresholds ($c = 2$ [s], $c = 3$ [s], $c = 4$ [s]) using both kNN and the Naive Bayes classifiers for reference. These accuracies are calculated using the confusion matrix described in 5.6.2.

Short idle classification accuracy is found by taking the number of short idles properly classified over the total number of short idles. In other words, short idle accuracy is the true positive rate (*TPR*),

$$TPR = \frac{\sum True\ Positive}{\sum True\ Positive + \sum False\ Negative}. \quad (5.27)$$




Figure 5-24: Classification Accuracy for (a) Short and (b) Long Idles Using Position Only with Simulated Data

As discussed in Section 5.6.2, a false negative event occurs when the idle is actually short, but it is predicted to be long. In this case, a false negative results in the Smart/Stop system turning the engine off for a short idle. This is equivalent to the normal AESS system turning the engine off for a short idle. This is equivalent to the normal AESS system response, so it is merely a return to the status quo (no different from the baseline). Every true positive prediction will reduce the number of change-of-mind events, improve drivability, and cut down on fuel wasted from excessive starting and stopping.

Similarly, long idle classification accuracy can be found by taking the number of properly classified long idles divided by the total number of long idles. This accuracy is known as the true negative rate (TNR),

$$TNR = \frac{\sum True\ Negative}{\sum True\ Negative + \sum False\ Positive}. \quad (5.28)$$

In terms of the Smart/Stop system, a true negative prediction will result in the engine stopping for a long idle event (which is equivalent to the baseline). This maintains the fuel economy benefits of AESS. A false positive prediction will cause the system to idle the engine for a long event, which circumvents any fuel economy gains for that particular event, but also eliminates any driver frustration due to a change-of-mind hesitation. In this sense, a false positive prediction is comparable to




Figure 5-25: Classification Accuracy for (a) Short and (b) Long Idles Using Speed of First Car in Front Only with Simulated Data

driver disablement.

Figure 5-24 shows the classification accuracies from simulation for short and long idles using position only in the learning vector. This means that at the start of the idle, the only information available is where that idle is taking place along the simulated road. We can see in Figure 5-24a that the best performance for short idle classification occurs with one neighbor ($k = 1$) and drops off quickly as k increases for each of the three short idle thresholds. The peak short idle accuracy is between 65-68% using position only.

The simulated long idle classification accuracy for position only is shown in Figure 5-24b. Long idle classification accuracy is much higher across the board than for short idles in this case, and the accuracy increases as k increases. This follows from the fact that there are far more long idles than short idles (see Figure 5-15). Taking the best k value for short idle accuracy still yields between 97-99% accuracy for long idles.

It is interesting to note that the Naive Bayes classifier always predicts a long idle for the position-only case (Figure 5-24). This results in 0% accuracy for short idles and 100% accuracy for long idles. This indicates that for any given position, there are more long idles than short idles, and makes this method equivalent to driver disablement for Smart/Stop.

Short and long idle classification accuracies based on speed alone are shown in




Figure 5-26: Classification Accuracy for (a) Short and (b) Long Idles Using Position and Speed of First Car in Front with Simulated Data

Figures 5-25a and 5-25b for the simulated data set. Unlike the position-based predictions, the speed-based predictions are relatively constant for all k values. For short idles, the classification accuracy ranges substantially between different short idle thresholds from as low as 22% for $c = 4$ [s] to as high as 84% when $c = 2$ [s]. Figure 5-21a illustrates why speed is a much better parameter for lower short idle thresholds (since the speed is almost always zero for idles longer than 2 seconds). For long idles, using speed as a predictor results in over 99% accuracy for both kNN and Naive Bayes.

Using both position and speed in the learning vector leads to better results than either position or speed alone, as Figure 5-26 illustrates for the simulated data. Like the position only plots, the performance tapers off for short idles as k increases, but increases for long idles. Using a k value of 1 results in classification accuracies between 88-97% for short idles and from 98-99% for long idles. Naive Bayes, on the other had, predicts long idles every time, resulting in perfect long accuracy and zero short accuracy. This would be the same as disabling AESS entirely.

Figure 5-27 shows the classification accuracies when the position of vehicle i , speed of vehicle $i - 1$, and speed of vehicle $i - 2$ are all included as features in the learning vector. The addition of the speed of vehicle $i - 2$ does not appear to appreciably increase the classification accuracy with respect to position of vehicle i and speed of




Figure 5-27: Classification Accuracy for (a) Short and (b) Long Idles Using Position and Speeds of First and Second Car in Front with Simulated Data

vehicle $i - 1$ alone. There is a slight improvement, but there seems to be diminishing returns. The Naive Bayes classifier still does poorly but has significantly improved when the short idle threshold is two seconds.

Predicting Idle Time with Real-World Data

Since kNN performed well on the traffic simulation data in Section 5.6.6, we attempted to validate these results using real-world data. Figures 5-28 through 5-30 show the results of applying two different classifiers on three different learning vectors for the Cambridge data set. Each figure plots the short and long classification accuracies using both kNN and a Naive Bayes baseline over three short idle thresholds ($c = 2 [s]$, $c = 3 [s]$, $c = 4 [s]$). As before, these accuracies were calculated using a confusion matrix and Equations 5.27 and 5.28.

Figure 5-28 shows the short and long classification accuracies using GPS position only on the Cambridge data set. For short idle accuracy shown in Figure 5-28a, performance ranges between 50-70% depending on the short idle threshold as long as k is less than 25, but drops off as k increases further. Long idle classification accuracy increases gradually with k and hovers between 70-80% when k is less than 25. This performance is comparable to simulation for short idles and slightly reduced for long idles. Naive Bayes still performs relatively poorly for short idles and very well for





Figure 5-28: Classification Accuracy for (a) Short and (b) Long Idles Using Position Only with Cambridge Data

long idles (due to the statistical bias toward long idles).


Classification accuracies using the speed of vehicle $i - 1$ alone on the Cambridge data set are shown in Figure 5-29 for short and long idles. For low values of k , accuracy is very high for short idles but very low for long idles. Once k increases to 3 for $c = 2$ [s] and 7 for $c = 3$ [s] and $c = 4$ [s], the situation flips as accuracy is high for long idles and low for short idles. The value of k cannot be varied preemptively, so in practice, a single value must be chosen, which means that the high accuracy cannot simultaneously be maintained for both short and long idles using speed alone. Once again, Naive Bayes performs poorly for short idles and extremely well for long idles.

Figure 5-30 shows the short and long classification accuracies on the Cambridge data set using a combination of GPS position of vehicle i and the speed of vehicle $i - 1$. As in the simulation, combining both position and speed yield accuracies higher than either of the two alone. Short idle classification reaches peak performance when $k = 3$, where accuracies range between 60-73% depending on the short idle threshold. As k increases beyond 3, the accuracies continue to decline.

Selecting the same value of $k = 3$ for long idles (Figure 5-30a) provides performance levels between 67% and 80% accuracy. Depending on the requirement, k can be increased to provide better long idle prediction at the cost of short idle accuracy. The Naive Bayes classifier reaches levels between 20-26% for short idles and 90-94%




(a)




(b)

Figure 5-29: Classification Accuracy for (a) Short and (b) Long Idles Using Speed of First Car in Front Only with Cambridge Data



(a)



(b)

Figure 5-30: Classification Accuracy for (a) Short and (b) Long Idles Using Position and Speed of First Car in Front with Cambridge Data

for long idles.

AESS and Smart/Stop Fuel Efficiency

Before we can evaluate the utility of Smart/Stop, we need to establish some metrics for measuring the relative fuel efficiency of various strategies. As discussed in Section 5.2.3, one important consideration for AESS is how much fuel is consumed during the startup sequence versus the amount of fuel consumed for a given amount of engine idling. The critical threshold ($T_{I_{crit}}$) at which the fuel saved from the elimination of engine idling exceeds the fuel consumed during engine start is on the order of 2 seconds [10].

When conventional AESS systems are operating within their operational regions, they shut the engine down for every possible idle event. Since the duration of idles are widely distributed (see Figure 5-9), many of these events will be under the critical threshold and actually consume more fuel than idling the engine for that same period. Because of the fixed cost of fueling during startup, all of the fuel savings available through raw idle elimination cannot be achieved. We can represent the efficiency of a conventional AESS system as

$$\eta_0 = \frac{\sum_{z=1}^{Z_0} T_{I_z} - Z_0 T_{I_{crit}}}{\sum_{z=1}^{Z_0} T_{I_z}}, \quad (5.29)$$

where T_{I_z} is the idle time of idle event z and Z_0 is the total number of idles. In order to achieve perfect efficiency, starting the engine would have to use zero fuel, which is not feasible for micro hybrids.

Similarly, the efficiency of the Smart/Stop system can be represented as

$$\eta_{SS} = \frac{\sum_{z=1}^{Z_0} T_{I_z} - \sum_{q=1}^{Z_P} T_{I_q} - Z_N T_{I_{crit}}}{\sum_{z=1}^{Z_0} T_{I_z}}, \quad (5.30)$$

where Z_P is the number of idles predicted short, Z_L is the number of idles predicted

Table 5.4: Smart/Stop Results Summary

c	k	TPR	TNR	$\eta_0(\frac{1}{2})$	$\eta_{SS}(\frac{1}{2})$	$\eta_0(1)$	$\eta_{SS}(1)$	$\eta_0(2)$	$\eta_{SS}(2)$	$\eta_0(3)$	$\eta_{SS}(3)$
2	3	60.1	80.2	93.4	88.1	86.9	82.7	73.7	72.1	60.6	61.4
2	5	58.0	85.5	93.4	89.6	86.9	84.0	73.7	72.7	60.6	61.4
2	7	54.5	84.7	93.4	90.9	86.9	85.1	73.7	73.4	60.6	61.6
2	9	59.4	85.5	93.4	91.0	86.9	85.0	73.7	73.1	60.6	61.2
3	3	68.7	72.2	93.4	82.1	86.9	77.8	73.7	69.1	60.6	60.4
3	5	68.7	79.7	93.4	84.2	86.9	79.5	73.7	70.1	60.6	60.7
3	7	67.0	81.1	93.4	84.6	86.9	80.0	73.7	70.8	60.6	61.5
3	9	66.5	80.7	93.4	85.5	86.9	80.7	73.7	71.2	60.6	61.7
4	3	71.9	66.5	93.4	77.8	86.9	74.0	73.7	66.5	60.6	59.0
4	5	73.9	75.0	93.4	81.6	86.9	77.5	73.7	69.4	60.6	61.3
4	7	70.9	75.0	93.4	82.4	86.9	78.2	73.7	69.9	60.6	61.5
4	9	68.0	77.1	93.4	83.7	86.9	79.4	73.7	70.8	60.6	62.2

long, and T_{I_q} is the idle time of smart idle q .

As we can see from Equations 5.29 and 5.30, the fuel efficiencies of using conventional AESS and Smart/Stop depend heavily on the critical time threshold ($T_{I_{crit}}$) at which idling the engine uses more fuel than stopping and starting it again.

Predicting Idle Time Results Summary

Now that we have measured the performance of the classification algorithms over a wide range of parameters, and defined metrics for evaluating relative fuel efficiency, we can compare the tradeoff between eliminating short idles and decreasing fuel efficiency using Smart/Stop.

As we showed in Section 5.6.6, the total percentage of change-of-mind events that can be eliminated using Smart/Stop (TPR) and total percentage of long idles that can be maintained (TNR) both depend on the threshold (c) below which idles are considered to be short as well as the k value used in the kNN classifier.

Table 5.4 shows TPR and TNR for c thresholds of 2, 3, and 4 seconds and k values of 3, 5, 7, and 9. For each combination, the baseline efficiency of AESS (η_0) and efficiency of Smart/Stop (η_{SS}) are included for critical idling time thresholds ($T_{I_{crit}}$) of 0.5, 1, 2, and 3 seconds.

From examining Table 5.4, a few trends become clear. As k increases, the total




Figure 5-31: Histogram of Idle Durations for Cambridge Data when AESS is Disabled and when Smart/Stop is Enabled

positive rate tends to decrease and the total negative rate tends to increase in general. On the contrary, higher c thresholds yield higher TPR and lower TNR . The lower the critical idle time threshold, the higher the absolute efficiencies become with 0.5 seconds giving the highest efficiencies. As the critical threshold increases, however, the gap between the baseline efficiency and the Smart/Stop efficiency narrows until Smart/Stop actually overtakes the baseline efficiency for most cases when $T_{I_{crit}} = 3$ seconds.

The particular requirements and parameters of a given application will determine which values in Table 5.4 are most valid, but in any case, several general conclusions can be drawn about the potential benefits of Smart/Stop. With sufficient density of data, anywhere between 55 to 74% of change-of-mind events can be eliminated using a kNN classifier. The fuel efficiency that must be sacrificed to eliminate these would-be short idles is small in the worst case, negligible for most cases, and even flipped (efficiency is improved) for the best cases.

It is important to remember the efficiency differences between AESS and Smart/Stop

may be on the order of 1 – 12% in terms of idling time, but idling time makes up between 3 – 10% of fuel consumption for most driving cycles. This means that the effect on overall fuel economy will be on the order of 0.03 – 1.2% of overall fuel economy. These difference are small when compared to the potential to reduce driver disablement of the start/stop system, which the EPA estimates as 11% of the time when granting off-cycle credits [91]. On top of this, there are the additional benefits to drivability, user experience, and safety.

As an example of how Smart/Stop looks in terms of distribution of idle times, Figure 5-31 shows a comparison between AESS disabled and the Smart/Stop when $c = 4$ and $k = 5$. From Table 5.4, we can see that for this case, 73.9% of short idles were properly classified, and therefore, the engine was allowed to idle through them. For long idles, 75.0% were correctly predicted, so the engine was shut down for those. From the plot, it is clear that majority of misclassified long idles are relatively close to the cutoff threshold. When the critical idle time threshold is 2 seconds, the Smart/Stop algorithm only results in 4.3% additional effective idling time over the baseline AESS case. Given a cycle where eliminating all idling increases fuel economy by 5%, Smart/Stop would still capture a 4.8% improvement in fuel economy all while eliminating the vast majority of change-of-mind events.

5.7 Conclusion

In this chapter, I showed how connected vehicle technology and collaborative learning can be applied to the AESS system in micro hybrid vehicles. In particular, I introduced the concept of a Smart/Stop system that uses machine learning to predict and eliminate change-of-mind events before they occur. These change-of-mind events are not only annoying to the driver, but inefficient with respect to fuel economy, and have been shown to cause drivers to disable the AESS system entirely (circumventing all efficiency improvement).

I proved that for our real-world data set the Smart/Stop system can remove as many as 74% of change-of-mind events while still maintaining the fuel economy ben-

efits of AESS. With even richer data sets, this accuracy rate will only increase, and so the potential for a Smart/Stop system to reduce driver disablements is great. This application is only one example of data-based prediction for the purpose of energy management in hybrid vehicles, but it illustrates the potential for massive improvements when vehicles control strategies extend far beyond the bounds of a single vehicle.

Chapter 6

Conclusions

6.1 Insights on Hybrid Vehicle Optimization

In this thesis, I showed that there are still many opportunities to improve upon state-of-the-art hybrid vehicle technology. These opportunities are enabled by the increased volumes of hybrid vehicles and the potential of control strategies to look further into the near future by utilizing the latest sensor technology and the cloud for historical contextual insights. This revolution in data availability will take vehicles from reacting to the environment after the fact to predicting future events and preparing for them ahead of time. Without this type of foresight, automotive manufacturers would have to continue to build in excess hardware capacity to account for the uncertainty in vehicle control and energy management. In the next three sections, I will discuss the major insights from each chapter of this thesis.

6.1.1 Battery Pack Design

Many factors go into cell selection for hybrid and electric vehicles, but as the sales volumes of electrified vehicles continue to rise, manufacturers have more freedom to customize a particular cell for a specific vehicle. In the past, some manufacturers have turned to cells that benefit from high sales volumes in the consumer electronics markets, but this forces them to design around the parameters of that cell.

A single vehicle platform can now have high enough volumes to justify tooling up for entirely customized cell sizes and form factors. With this freedom to choose an arbitrary cell size, shape, and chemistry comes the question of which one is best. In Chapter 3, I showed that due to the tradeoff between the average and delta temperature constraints, this question has a definitive answer with respect to air-cooled battery packs.

Choosing a suboptimal size or shape can multiply the necessary cost of cooling by as much as two to three times over the best case cell. I also showed that some form factors perform better than others regardless of cell size for a given module-level energy density. Specifically, rounded prismatic cells exhibited the lowest cost of cooling across the board.

Finally, after defining a new analytical method to calculate the minimum cost cell size, I showed that it matches the results of the simulation. This provides a design tool for battery pack designers that can be used with a simple spreadsheet instead of requiring time consuming multi-physics simulations. With this tool, designers are better equipped to ensure they are minimizing the cost of air-cooling before they get locked into a cell choice that is inherently less efficient.

6.1.2 Energy Management and Blended Battery Packs

Energy management is often viewed as an optimization problem after the hardware has been defined. In this thesis, I showed that the energy management strategy can have a significant effect on the minimal size (and relatedly cost) of a blended battery pack. If the power split is not defined in such a way as to minimize the hardware size from the start, it is not possible to fairly compare potential energy storage types in both uniform and blended configurations and understand which offers the best solution for an application.

In Chapter 4, I defined a dynamic programming based method that allows blended battery packs to be sized for minimal weight without predefining a power split strategy. Then, I utilized this method along with an integer programming solution to show the lower bound and a peak shaving heuristic strategy as a benchmark to carry

out a large experiment across 18 different energy storage types and 7 different automotive use cases. The results of this study allowed me to draw conclusions about when blended battery packs can be beneficial and how much potential there is for improvement over uniform packs.

If there is a single cell type that can supply the required power and energy with a very small amount of excess power or energy capacity, this source is likely to be the best solution. If there is excess power or energy capacity in the best uniform battery pack, however, this capacity represents additional cost, and there is most likely an opportunity for tribridization. The strength of blended battery packs is derived from the extra degree of flexibility that they unlock. When a high energy and high power source can be sized independent of each other, any excess capacity that would be necessarily present in a uniform battery pack can be eliminated.

From the study, it is clear that blended battery packs can decrease the mass of packs for HEVs and PHEVs but are not well suited for pure EVs. This is due to the fact that while the energy required for a given EV application scales up linearly with the required range, the power requirement stays constant. As the required range is increased, the necessary C rate drops proportionally and allows higher and higher energy density cells to deliver the required power without excess energy capacity. These results suggest that pack designers for HEV and PHEV vehicles should consider tribridization as a means for weight reduction before locking into a uniform battery pack.

6.1.3 Collaborative Learning for Future Insights

While the potential gains of blended battery packs are large, I illustrated that the power split strategy that enables the minimization of pack size requires future knowledge of the drive cycle at hand. Until vehicles become fully autonomous, this requires some pattern recognition and collaborative learning on the part of the vehicle control system in order to predict the future with some level of certainty.

In Chapter 5, I went into detail on how collaborative learning can be used to predict the future for the purpose of improving the drivability of micro hybrids without

sacrificing the fuel economy gains of the automatic engine start/stop (AESS) system. The baseline AESS stops the engine every time the vehicle stops (when able). The main issue with this strategy is that many of these stops are only momentary, so the driver wishes to accelerate again faster than the engine can possibly stop and then start.

These change-of-mind events are not only a detriment to drivability, but capture little to no fuel economy benefit, and can even result in more fuel consumption than idling the engine. To prevent this hesitation, some drivers disable the system entirely. In order to eliminate these change-of-mind events, I introduced the idea of a Smart/Stop system, and illustrated how such a system could be implemented using historical context and machine learning classifiers.

Using the Smart/Stop system, I showed that as many as 74% of change-of-mind events can be eliminated while maintaining over 97% of the fuel economy benefits of AESS. This system has the potential to not only improve drivability, but convince drivers who would otherwise disable AESS to leave it on and increase the overall fleet fuel economy savings. While this is only one application of collaborative learning for hybrid vehicles, it shows that there is great potential for improvement over current systems when vehicles can utilize proactive control strategies.

6.2 Summary of Contributions

To summarize, the main contributions of my thesis are as follows.

- I have advanced the understanding of how cell-level properties affect the overall costs of air-cooling in battery packs and have shown how to minimize those costs. Through both simulation and analytical studies, I illustrated that cell size and geometry can be selected to balance the tradeoff between average and delta temperature constraints. At the intersection of these two constraints, a cell size can be found that results in the minimum cost of air-cooling for the parameters of that application. While the simulation results take a good deal

of time to compute, the analytical method can be used to narrow down a design space with very little effort.

- I have advanced the understanding of how blended battery packs can be sized for a given application and what applications will benefit most from tribridization. Using integer programming for a lower bound, a dynamic programming based method for minimum feasibility, and a peak shaving heuristic as a baseline, I carried out a large experiment to quantify how much room for improvement there is for a variety of hybrid applications using current energy storage technology. Based on the results, I extracted rules of thumb for when to consider blended battery packs.
- I have advanced the understanding of how collaborative learning can benefit energy management for hybrid vehicles. By defining a traffic model and comparing that to real-world datasets, I identified parameters that are most useful in predicting idle time of vehicles in traffic. Using these predictors in combination with machine learning classifiers, I quantified the potential for a Smart/Stop system to improve over current AESS performance. The result is a system that can significantly increase drivability while maintaining or even improving on baseline fuel economy.

6.3 The Death of the Drive Cycle

The drive cycle has served the automotive industry very well in the quest for lower emissions and higher fuel economy, but we have reached a point where it is becoming more of a detriment than a benefit. The findings of this thesis suggest that there is a better way.

In the same way that teachers are now incentivized to focus on skills that are measured on standardized tests instead of those that are applicable in the real world, automotive manufactures are encouraged to optimize their control strategies for the EPA tests instead of targeting better real-world fuel economy. Just as students might

cheat on a standardized test, we have seen certain automotive manufacturers doing the same on the EPA tests.

This thesis outlined a number of ways that the drive cycles can be used to optimize the performance of hybrid vehicles, but it is critical to emphasize the importance of not falling into the trap of overfitting, or real-world fuel economy and emissions will suffer. When control strategies are tuned for these particular cycles, we are not achieving the overall optimal performance and efficiency.

As discussed in Section 5.3.1, the EPA has begun allowing manufacturers to gain off-cycle credits by proving that their vehicles are actually achieving higher fuel economy than the drive cycle tests might show. This acknowledges the problem that a small number of fixed cycles are not representative of the sum total of driving cycles and address the issue of underestimating fuel economy, but it does not solve the issue of overestimation.

Just as the connected vehicle data was shown to allow for further optimization of the AESS system in micro hybrids in Chapter 5, this data can be used to evaluate the fuel economy of vehicles in the real world. By using drive cycles as an initial benchmark, and then measuring real world usage data with some sample set of instrumented vehicles, the EPA can automatically reward off-cycle credits after the fact and penalize any manufacturers that overfit to the EPA cycles. This will incentivize real-world fuel economy and reward strategies that have measurable impact on fuel consumed in practice.

6.4 Future Work

This thesis showed that there is still a great deal of room for improvement when it comes to the efficiency and performance of hybrid vehicles. These opportunities rest at every level of hybrid vehicle design, from the choice of what size and shape of cell is best for air-cooling, to yet another layer of hybridization through blended battery packs, to the optimization of energy management strategies through collaborative learning. Even with all of these areas better understood, there is still much to be

learned.

While I identified that there is a minimum cost cell size and geometry for a particular application with respect to air-cooled battery packs, a similar study is required for water-cooled battery packs. Furthermore, the temperature gradient across the pack was found to be sensitive to the intensity of turbulence at the inlet. If the intensity is high, the turbulence will decrease as it travels through the pack. If it is low, the intensity will increase. As heat transfer increases with higher turbulence, this effect could be used as a design variable to help mitigate issues with excessive temperature gradient. Additionally, these types of generalized analyses should be extended to other aspects of battery pack design to develop more fundamental understanding of how packs should be designed for minimum cost.

I showed that blended battery packs have the potential to reduce the mass of hybrid and plug-in hybrid vehicle battery packs, and that the gains are maximized when an omniscient control strategy is utilized. Since the speed profile of the vehicle is not predefined, the next step in fully minimizing the mass of these packs is to attempt to approach this optimal control strategy with drive cycle prediction algorithms. By utilizing connected vehicle technology this type of speed profile prediction could be feasible. To better understand just how close we can come to the optimal strategy, a study using stochastic programming in combination with speed predictions will be required.

Through the application of collaborative learning to the AESS system in micro hybrids, I illustrated the power of predictive algorithms for the purpose of hybrid powertrain optimization. While the gains appear promising from the initial study, further development and testing will be required to ensure that it can be a viable feature in production vehicles. This application shows great potential, but it is really only the tip of the iceberg when it comes to automotive application areas. More and more areas of improvement will need to be identified and studied in order to justify the costs and bring collaborative learning into the mainstream.

Appendix A

Result Summary Tables from Chapter 4

This appendix provides the summary tables of the results from Chapter 4. These tables show which energy storage types from Table 4.1 were chosen for each of the blended solutions (DDP, PS, IP, and Uniform) for all 7 applications from Table 4.3.

Table A.1: Overall Results for DDP Solution for 4 Consecutive Cycles

ID#	Base	Buff	P_{base}	P_{buff}	Mass [kg]
1	1	16	4	1	39.5
2	1	16	1	1	16.4
3	1	2	7	3	370.7
4	1	2	30	2	442.2
5	12	—	1	—	45.4
6	1	16	5	2	55.9
7	1	16	7	2	71.3

Table A.2: Overall Results for PS Solution for 4 Consecutive Cycles

ID#	Base	Buff	P_{base}	P_{buff}	Mass [kg]
1	12	—	1	—	45.4
2	1	—	3	—	23.1
3	1	2	7	3	370.7
4	1	2	30	2	442.2
5	12	—	1	—	45.4
6	6	—	1	—	68.5
7	1	6	3	1	91.6

Table A.3: Overall Results for IP Lower Bound for 4 Consecutive Cycles

ID#	Base	Buff	P_{base}	P_{buff}	Mass [kg]
1	1	16	2	1	24.1
2	1	16	1	1	16.4
3	1	2	17	2	342.1
4	1	2	28	2	426.8
5	1	16	2	1	24.1
6	1	16	4	2	48.2
7	1	16	8	1	70.3

Table A.4: Overall Results for Uniform Solution for 4 Consecutive Cycles

ID#	ESU	P	Mass [kg]
1	12	1	45.4
2	1	3	23.1
3	2	4	422.4
4	2	5	528
5	12	1	45.4
6	6	1	68.5
7	2	1	105.6

Table A.5: Overall Results for DDP Solution for 7 Consecutive Cycles

ID#	Base	Buff	P_{base}	P_{buff}	Mass [kg]
1	1	16	4	1	39.5
2	1	16	1	1	16.4
3	1	2	47	2	573.1
4	1	2	80	1	721.6
5	1	12	1	1	53.1
6	1	16	7	1	62.6
7	1	16	12	1	101.1

Table A.6: Overall Results for PS Solution for 7 Consecutive Cycles

ID#	Base	Buff	P_{base}	P_{buff}	Mass [kg]
1	1	12	2	1	60.8
2	1	—	3	—	23.1
3	1	2	62	1	583
4	1	2	80	1	721.6
5	1	16	6	1	54.9
6	1	6	3	1	91.6
7	1	16	12	2	109.8

Table A.7: Overall Results for IP Lower Bound for 7 Consecutive Cycles

ID#	Base	Buff	P_{base}	P_{buff}	Mass [kg]
1	1	16	2	1	24.1
2	1	16	1	1	16.4
3	1	2	59	1	559.9
4	1	16	90	2	710.4
5	1	16	3	1	31.8
6	1	16	7	1	62.6
7	1	16	12	1	101.1

Table A.8: Overall Results for Uniform Solution for 7 Consecutive Cycles

ID#	ESU	P	Mass [kg]
1	6	1	68.5
2	1	3	23.1
3	2	6	633.6
4	1	101	777.7
5	6	1	68.5
6	2	1	105.6
7	1	16	123.2

Table A.9: Overall Results for DDP Solution for 10 Consecutive Cycles

ID#	Base	Buff	P_{base}	P_{buff}	Mass [kg]
1	1	16	4	1	39.5
2	1	16	1	1	16.4
3	1	16	101	1	786.4
4	1	—	129	—	993.3
5	1	16	6	1	54.9
6	1	16	10	1	85.7
7	1	—	17	—	130.9

Table A.10: Overall Results for PS Solution for 10 Consecutive Cycles

ID#	Base	Buff	P_{base}	P_{buff}	Mass [kg]
1	1	16	7	1	62.6
2	1	—	3	—	23.1
3	1	2	89	1	790.9
4	1	—	129	—	993.3
5	1	16	7	1	62.6
6	1	16	12	1	101.1
7	1	—	17	—	130.9

Table A.11: Overall Results for IP Lower Bound for 10 Consecutive Cycles

ID#	Base	Buff	P_{base}	P_{buff}	Mass [kg]
1	1	16	3	1	31.8
2	1	16	1	1	16.4
3	1	16	101	1	786.4
4	1	—	129	—	993.3
5	1	16	4	1	39.5
6	1	16	10	1	85.7
7	1	—	17	—	130.9

Table A.12: Overall Results for Uniform Solution for 10 Consecutive Cycles

ID#	ESU	P	Mass [kg]
1	6	1	68.5
2	1	3	23.1
3	1	105	808.5
4	1	129	993.3
5	6	1	68.5
6	2	1	105.6
7	1	17	130.9

Bibliography

- [1] Gaines, L.; Cuenca, R. Costs of Lithium-Ion-Batteries for Vehicles; Argon National Laboratory: Argonne, IL, 2000.
- [2] Chan, C.C., “The State of the Art of Electric, Hybrid, and Fuel Cell Vehicles,” Proceedings of the IEEE , vol.95, no.4, pp.704,718, April 2007 doi: 10.1109/JPROC.2007.892489
- [3] Robin Cowan, Staffan Hultén, Escaping lock-in: The case of the electric vehicle, Technological Forecasting and Social Change, Volume 53, Issue 1, September 1996, Pages 61-79, ISSN 0040-1625, 10.1016/0040-1625(96)00059-5.
- [4] Liu, Jinming, and Huei Peng. “Modeling and control of a power-split hybrid vehicle.” Control Systems Technology, IEEE Transactions on 16.6 (2008): 1242-1251.
- [5] Åhman, Max. “Primary energy efficiency of alternative powertrains in vehicles.” Energy 26.11 (2001): 973-989.
- [6] Otmar Bitsche, Guenter Gutmann, Systems for hybrid cars, Journal of Power Sources, Volume 127, Issues 1-2, 10 March 2004, Pages 8-15, ISSN 0378-7753, 10.1016/j.jpowsour.2003.09.003.
- [7] Anderson, David L. An evaluation of current and future costs for lithium-ion batteries for use in electrified vehicle powertrains. Diss. Duke University, 2009.

- [8] Ye, Fei, Matthew Adams, and Sumit Roy. "V2V wireless communication protocol for rear-end collision avoidance on highways." Communications Workshops, 2008. ICC Workshops' 08. IEEE International Conference on. IEEE, 2008.
- [9] Bergenheim, Carl, Erik Hedin, and Daniel Skarin. "Vehicle-to-vehicle communication for a platooning system." *Procedia-Social and Behavioral Sciences* 48 (2012): 1222-1233.
- [10] Fesefeldt, Thomas, and Soren Müller. Optimization and comparison of quick and hybrid start. No. 2009-01-1340. SAE Technical Paper, 2009.
- [11] Lukic, Srdjan M., et al. "Energy storage systems for automotive applications." *Industrial electronics, IEEE Transactions on* 55.6 (2008): 2258-2267.
- [12] Peter Van den Bossche, Frédéric Vergels, Joeri Van Mierlo, Julien Matheys, Wout Van Autenboer, SUBAT: An assessment of sustainable battery technology, *Journal of Power Sources*, Volume 162, Issue 2, 22 November 2006, Pages 913-919, ISSN 0378-7753, 10.1016/j.jpowsour.2005.07.039.
- [13] Randolph A. Leising, Marcus J. Palazzo, Esther Sans Takeuchi, Kenneth J. Takeuchi, A study of the overcharge reaction of lithium-ion batteries, *Journal of Power Sources*, Volumes 97-98, July 2001, Pages 681-683, ISSN 0378-7753, 10.1016/S0378-7753(01)00598-5.
- [14] Lee, Yuang-Shung, and Ming-Wang Cheng. "Intelligent control battery equalization for series connected lithium-ion battery strings." *Industrial Electronics, IEEE Transactions on* 52.5 (2005): 1297-1307.
- [15] Cao, Jian, Nigel Schofield, and Ali Emadi. "Battery balancing methods: A comprehensive review." *Vehicle Power and Propulsion Conference, 2008. VPPC'08. IEEE. IEEE, 2008.*
- [16] Chau, K. T., and Y. S. Wong. "Overview of power management in hybrid electric vehicles." *Energy conversion and management* 43.15 (2002): 1953-1968.

- [17] Emadi, Ali, Young Joo Lee, and Kaushik Rajashekara. "Power electronics and motor drives in electric, hybrid electric, and plug-in hybrid electric vehicles." *Industrial Electronics, IEEE Transactions on* 55.6 (2008): 2237-2245.
- [18] Erb, Dylan Charles. Optimization of blended battery packs. S.M. Thesis. Massachusetts Institute of Technology, 2013.
- [19] Cao, Jian, and Ali Emadi. "A new battery/ultracapacitor hybrid energy storage system for electric, hybrid, and plug-in hybrid electric vehicles." *Power Electronics, IEEE Transactions on* 27.1 (2012): 122-132.
- [20] Lukic, Srdjan M., et al. "Power management of an ultracapacitor/battery hybrid energy storage system in an HEV." *Vehicle Power and Propulsion Conference, 2006. VPPC'06. IEEE. IEEE, 2006.*
- [21] Kuperman, Alon, et al. "A frequency domain approach to analyzing passive battery-ultracapacitor hybrids supplying periodic pulsed current loads." *Energy Conversion and Management* 52.12 (2011): 3433-3438.
- [22] Haifeng, Dai, and Chang Xueyu. "A Study on Lead Acid Battery and Ultracapacitor Hybrid Energy Storage System for Hybrid City Bus." *Optoelectronics and Image Processing (ICOIP), 2010 International Conference on. Vol. 1. IEEE, 2010.*
- [23] Gao, Lijun, Roger A. Dougal, and Shengyi Liu. "Power enhancement of an actively controlled battery/ultracapacitor hybrid." *Power Electronics, IEEE Transactions on* 20.1 (2005): 236-243.
- [24] Douglas, H., and P. Pillay. "Sizing ultracapacitors for hybrid electric vehicles." *Industrial Electronics Society, 2005. IECON 2005. 31st Annual Conference of IEEE. IEEE, 2005.*
- [25] Lu, Shuai, Keith A. Corzine, and Mehdi Ferdowsi. "A new battery/ultracapacitor energy storage system design and its motor drive integration for hybrid electric vehicles." *Vehicular Technology, IEEE Transactions on* 56.4 (2007): 1516-1523.

- [26] Liu, Shengyi, and Roger A. Dougal. "Design and analysis of a current-mode controlled battery/ultracapacitor hybrid." Industry Applications Conference, 2004. 39th IAS Annual Meeting. Conference Record of the 2004 IEEE. Vol. 2. IEEE, 2004.
- [27] Miller, John M., and Gianni Sartorelli. "Battery and ultracapacitor combinations - Where should the converter go?" Vehicle Power and Propulsion Conference (VPPC), 2010 IEEE. IEEE, 2010.
- [28] Awerbuch, Jonathan J., and Charles R. Sullivan. "Control of ultracapacitor-battery hybrid power source for vehicular applications." Energy 2030 Conference, 2008. ENERGY 2008. IEEE. IEEE, 2008.
- [29] Govindaraj, Arvind. "Design and Characterization of Various Circuit Topologies for Battery-Ultracapacitor Hybrid Energy Storage Systems." (2010).
- [30] Li, Zhihao, et al. "Design and control of a multiple input DC/DC converter for battery-ultra-capacitor based electric vehicle power system." Applied Power Electronics Conference and Exposition, 2009. APEC 2009. Twenty-Fourth Annual IEEE. IEEE, 2009.
- [31] Garcia, F. S., A. A. Ferreira, and J. A. Pomilio. "Control strategy for battery-ultracapacitor hybrid energy storage system." Applied Power Electronics Conference and Exposition, 2009. APEC 2009. Twenty-Fourth Annual IEEE. IEEE, 2009.
- [32] Amjadi, Zahra, and Sheldon S. Williamson. "Prototype design and controller implementation for a battery-ultracapacitor hybrid electric vehicle energy storage system." Smart Grid, IEEE Transactions on 3.1 (2012): 332-340.
- [33] Burke, Andrew, Marshall Miller, and Hengbing Zhao. "Lithium batteries and ultracapacitors alone and in combination in hybrid vehicles: Fuel economy and battery stress reduction advantages." Energy 113.206 (2010): 073.

- [34] Onar, Omer C., and A. Khaligh. "A novel integrated magnetic structure based DC/DC converter for hybrid battery/ultracapacitor energy storage systems." *Smart Grid, IEEE Transactions on* 3.1 (2012): 296-307.
- [35] Anstrom, Joel R., et al. "Simulation and field-testing of hybrid ultracapacitor/battery energy storage systems for electric and hybrid-electric transit vehicles." *Applied Power Electronics Conference and Exposition, 2005. APEC 2005. Twentieth Annual IEEE. Vol. 1. IEEE, 2005.*
- [36] Yalamanchili, Krishna P., and Mehdi Ferdowsi. "Review of multiple input DC-DC converters for electric and hybrid vehicles." *Vehicle Power and Propulsion, 2005 IEEE Conference. IEEE, 2005.*
- [37] Wong, JennHwa, et al. "A parallel energy-sharing control for fuel cell-battery-ultracapacitor hybrid vehicle." *Energy Conversion Congress and Exposition (ECCE), 2011 IEEE. IEEE, 2011.*
- [38] Kuperman, A., et al. "Design of a semiactive battery-ultracapacitor hybrid energy source." *Power Electronics, IEEE Transactions on* 28.2 (2013): 806-815.
- [39] García, Oscar, et al. "Automotive DC-DC bidirectional converter made with many interleaved buck stages." *Power Electronics, IEEE Transactions on* 21.3 (2006): 578-586.
- [40] Berdichevsky, Gene, et al. "The tesla roadster battery system." *Tesla Motors Inc* (2006).
- [41] Kennedy, B., D. Patterson, and S. Camilleri. "Use of lithium-ion batteries in electric vehicles." *Journal of Power Sources* 90.2 (2000): 156-162.
- [42] Thackeray, Michael M., Christopher Wolverton, and Eric D. Isaacs. "Electrical energy storage for transportation – Approaching the limits of, and going beyond, lithium-ion batteries." *Energy & Environmental Science* 5.7 (2012): 7854-7863.

- [43] Park, Chanwoo, and Arun K. Jaura. "Dynamic thermal model of li-ion battery for predictive behavior in hybrid and fuel cell vehicles." SAE transactions 112.3 (2003): 1835-1842.
- [44] Motloch, Chester G., et al. "PNGV battery testing procedures and analytical methodologies for HEV's." Proc. SAE Future Car Congr (2002): 01-1950.
- [45] Stamps, Andrew T., et al. "Analysis of capacity fade in a lithium ion battery." Journal of Power Sources 150 (2005): 229-239.
- [46] Wright, Randy Ben, et al. "Calendar-and cycle-life studies of advanced technology development program generation 1 lithium-ion batteries." Journal of power sources 110.2 (2002): 445-470.
- [47] Mahamud, Rajib, and Chanwoo Park. "Reciprocating air flow for Li-ion battery thermal management to improve temperature uniformity." Journal of Power Sources 196.13 (2011): 5685-5696.
- [48] Rao, Zhonghao, and Shuangfeng Wang. "A review of power battery thermal energy management." Renewable and Sustainable Energy Reviews 15.9 (2011): 4554-4571.
- [49] Nelson, Paul A., Danilo J. Santini, and James Barnes. "Factors determining the manufacturing costs of lithium-ion batteries for PHEVs." Center for Transportation Research, Argonne National Laboratory, presented at EVS24, Stavanger, Norway, May (2009): 13-16.
- [50] Goldberg, Michael. "The packing of equal circles in a square." Mathematics Magazine (1970): 24-30.
- [51] Nurmela, Kari J., and Patric RJ Ostergard. "Packing up to 50 equal circles in a square." Discrete & Computational Geometry 18.1 (1997): 111-120.
- [52] Mollard, Michael, and Charles Payan. "Some progress in the packing of equal circles in a square." Discrete Mathematics 84.3 (1990): 303-307.

- [53] Zaghbi, K., et al. "LiFePO₄/polymer/natural graphite: low cost Li-ion batteries." *Electrochimica acta* 50.2 (2004): 263-270.
- [54] Erb, Dylan C., et al. "Size matters: Why cell size is vital for minimizing cost of air-cooling in battery packs." *Transportation Electrification Conference and Expo (ITEC), 2015 IEEE*. IEEE, 2015.
- [55] Turner, Mike, and Comair Rotron. "All you need to know about fans." *Electronics Cooling* 2 (1996): 10-13.
- [56] Erb, Dylan C., et al. "Effects of cell geometry on thermal management in air-cooled battery packs." *Transportation Electrification Conference and Expo (ITEC), 2015 IEEE*. IEEE, 2015.
- [57] Sabbah, Rami, et al. "Active (air-cooled) vs. passive (phase change material) thermal management of high power lithium-ion packs: Limitation of temperature rise and uniformity of temperature distribution." *Journal of Power Sources* 182.2 (2008): 630-638.
- [58] Zukauskas, A., and Romanas Ulinskas. "Heat transfer in banks of tubes in cross-flow." (1988).
- [59] Mills, A. F. *Heat transfer*. Prentice Hall, 1999.
- [60] Lienhard, John H. *A heat transfer textbook*. Courier Corporation, 2013.
- [61] Kumar, Sumeet, Marco A. Cartas-Ayala, and Todd Thorsen. "Thermal modeling and design analysis of a continuous flow microfluidic chip." *International Journal of Thermal Sciences* 67 (2013): 72-86.
- [62] Burke, Andrew; and Miller, Marshall. (2009). *Performance Characteristics of Lithium-ion Batteries of Various Chemistries for Plug-in Hybrid Vehicles*. UC Davis: Institute of Transportation Studies (UCD).

- [63] Burke, Andrew. (2000). Ultracapacitors: Why, How, and Where is the Technology. UC Davis: Institute of Transportation Studies (UCD). Retrieved from: <http://escholarship.org/uc/item/9n905017>
- [64] Baisden, Andrew C., and Ali Emadi. "ADVISOR-based model of a battery and an ultra-capacitor energy source for hybrid electric vehicles." *Vehicular Technology, IEEE Transactions on* 53.1 (2004): 199-205.
- [65] B.E. Conway, V. Birss, J. Wojtowicz, The role and utilization of pseudocapacitance for energy storage by supercapacitors, *Journal of Power Sources*, Volume 66, Issues 1-2, May-June 1997, Pages 1-14, ISSN 0378-7753, 10.1016/S0378-7753(96)02474-3.
- [66] T. Horiba, T. Maeshima, T. Matsumura, M. Koseki, J. Arai, Y. Muranaka, Applications of high power density lithium ion batteries, *Journal of Power Sources*, Volume 146, Issues 1-2, 26 August 2005, Pages 107-110, ISSN 0378-7753, 10.1016/j.jpowsour.2005.03.205.
- [67] Michael J Lain, Recycling of lithium ion cells and batteries, *Journal of Power Sources*, Volumes 97-98, July 2001, Pages 736-738, ISSN 0378-7753, 10.1016/S0378-7753(01)00600-0.
- [68] Kroeze, Ryan Carter, and Philip T. Krein. "Electrical battery model for use in dynamic electric vehicle simulations." *Power Electronics Specialists Conference, 2008. PESC 2008. IEEE. IEEE, 2008.*
- [69] Sundström, Olle, Lino Guzzella, and Patrik Soltic. "Optimal hybridization in two parallel hybrid electric vehicles using dynamic programming." *Proceedings of the 17th IFAC world congress*. Vol. 17. No. part 1. 2008.
- [70] Bayindir, Kamil Çagatay, Mehmet Ali Gözükcük, and Ahmet Teke. "A comprehensive overview of hybrid electric vehicle: Powertrain configurations, powertrain control techniques and electronic control units." *Energy Conversion and Management* 52.2 (2011): 1305-1313.

- [71] Bertsekas, Dimitri P., et al. Dynamic programming and optimal control. Vol. 1. No. 2. Belmont: Athena Scientific, 1995.
- [72] Vahidi, Ardalan, and Wesley Greenwell. "A decentralized model predictive control approach to power management of a fuel cell-ultracapacitor hybrid." American Control Conference, 2007. ACC'07. IEEE, 2007.
- [73] Dixon, Juan W., and Micah E. Ortuzar. "Ultracapacitors + DC-DC converters in regenerative braking system." Aerospace and Electronic Systems Magazine, IEEE 17.8 (2002): 16-21.
- [74] Simpson, A. G., and G. R. Walker. "Lifecycle costs of ultracapacitors in electric vehicle applications." Power Electronics Specialists Conference, 2002. pesc 02. 2002 IEEE 33rd Annual. Vol. 2. IEEE, 2002.
- [75] Li, Lianbing, et al. "Ultra-capacitor control strategy of EV with energy hybridization." Vehicle Power and Propulsion Conference, 2008. VPPC'08. IEEE. IEEE, 2008.
- [76] Siegel, Joshua. CloudThink and the Avacar. S.M. Thesis. Massachusetts Institute of Technology, 2013. Print.
- [77] Siegel, Joshua. Data Proxies, the Cognitive Layer, and Application Locality: Enablers of Cloud-Connected Vehicles and Next-Generation Internet of Things. Diss. Massachusetts Institute of Technology, 2016. Print.
- [78] Dillenbourg, Pierre. "What do you mean by collaborative learning." Collaborative-learning: Cognitive and computational approaches 1 (1999): 1-15.
- [79] Kramer, Ulrich. Potentialanalyse des Direktstarts für den Einsatz in einem Stopp-Start-System an einem Ottomotor mit strahlgeführter Benzin-Direkteinspritzung unter besonderer Berücksichtigung des Motorauslaufvorgangs. Diss. Universität Duisburg-Essen, Fakultät für Ingenieurwissenschaften, 2005.

- [80] Wang, Xiaoyong, Ryan McGee, and Ming Kuang. Vehicle system control for start-stop powertrains with automatic transmissions. No. 2013-01-0347. SAE Technical Paper, 2013.
- [81] Gao, B., Svancara, K., Walker, A., Kok, D., Conen, M. and Kees, D., "Development of a BISG Micro-Hybrid System," SAE Technical Paper 2009-01-1330, 2009, doi: 10.4271/2009-01-1330.
- [82] Sakai, K., Asada, T., Sugimura, K., and Nagata, R., "Development of Permanently Engaged Gear Mechanism for Idle Stop," JSAE Technical Paper No. 20095042, 2009
- [83] Schmidt, Martin, et al. Potential of regenerative braking using an integrated starter alternator. No. 2000-01-1020. SAE Technical Paper, 2000.
- [84] Robinette, Darrell, and Michael Powell. "Optimizing 12 Volt Start-Stop for Conventional Powertrains." SAE International Journal of Engines 4.1 (2011): 837-849.
- [85] Dhand, Aditya, et al. "Evaluation of Impaired Vehicle Launch for Start-Stop Micro Hybrid Vehicles." Europe in the second century of automobility, Proceedings of the 12th EAEC European Automotive Congress, Bratislava, Slovak Republic. Vol. 29. 2009.
- [86] Dhand, Aditya, et al. "Optimization potential of the vehicle launch performance for start-stop micro-hybrid vehicles." Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering 224.8 (2010): 1059-1070.
- [87] Ueda, Katsunori, et al. Idling stop system coupled with quick start features of gasoline direct injection. No. 2001-01-0545. SAE Technical Paper, 2001.
- [88] Matsuura, Moritaka, Koji Korematsu, and Junya Tanaka. Fuel consumption improvement of vehicles by idling stop. No. 2004-01-1896. SAE Technical Paper, 2004.

- [89] Wellmann, Thomas, et al. "Influence of Automatic Engine Stop/Start Systems on Vehicle NVH and Launch Performance." SAE International Journal of Engines 8.4 (2015): 1938-1946.
- [90] Wellmann, Thomas, Kiran Govindswamy, and Dean Tomazic. "Integration of engine start/stop systems with emphasis on NVH and launch behavior." SAE International Journal of Engines 6.2 (2013): 1368-1378.
- [91] "EPA Decision Document: Mercedes-Benz Off-cycle Credits for MYs 2012-2016" epa.gov. United States Environmental Protection Agency, Sept. 2014. Web. 7 Jan 2016.
- [92] Treiber, Martin, Ansgar Hennecke, and Dirk Helbing. "Congested traffic states in empirical observations and microscopic simulations." Physical Review E 62.2 (2000): 1805.
- [93] Rish, Irina. "An empirical study of the naive Bayes classifier." IJCAI 2001 workshop on empirical methods in artificial intelligence. Vol. 3. No. 22. IBM New York, 2001.
- [94] Horton, Paul, and Kenta Nakai. "Better Prediction of Protein Cellular Localization Sites with the k Nearest Neighbors Classifier." Ismb. Vol. 5. 1997.