Performance Modeling and Benchmark Analysis of an Advanced 4WD Series-Parallel PHEV Using Dynamic Programming

by

Stefan Kaban
B.Eng., University of Victoria, 2010

A Thesis Submitted in Partial Fulfillment of the
Requirements for the Degree of

Master of Applied Science

in the Department of Mechanical Engineering

Performance Modeling and Benchmark Analysis of an Advanced 4WD Series-Parallel PHEV Using Dynamic Programming

by

Stefan Kaban
B.Eng., University of Victoria, 2010

Supervisory Committee

Dr. Zuomin Dong, Supervisor
(Department of Mechanical Engineering)

Dr. Curran Crawford, Co-Supervisor
(Department of Mechanical Engineering)

Dr. Ashoka K.S. Bhat, Member
(Department of Electrical and Computer Engineering)

Supervisory Committee

Dr. Zuomin Dong, Supervisor
(Department of Mechanical Engineering)

Dr. Curran Crawford, Co-Supervisor
(Department of Mechanical Engineering)

Dr. Ashoka K.S. Bhat, Member
(Department of Electrical and Computer Engineering)

ABSTRACT

Advanced hybrid vehicle architectures can exploit multiple power sources and optimal control to achieve high efficiency operation. In this work, a method for generating the best-possible energy efficiency benchmark for a hybrid architecture is introduced. The benchmark program uses Dynamic Programming to analyse a reduced-fidelity MATLAB model over standard driving cycles, and bypasses vehicle controls to identify the optimal control actions and resulting fuel consumption of the Series-Parallel Multiple-Regime retrofitted PHEV of the UVic EcoCAR2 program.

The simulation results indicate an optimal fuel consumption value of 4.74L/100km, in the parallel regime, compared to the stock Malibu's 8.83L/100km. The results are found to be sensitive to the allowed level of regenerative braking, with an optimal consumption value of 6.56L/100km obtained with restricted regen power limits. The parallel regime provided more efficient operation overall, especially during more aggressive driving conditions. However, the series regime provided more desirable operation during gentle driving conditions, where opportunities for regenerative braking are limited.

The generated powertrain control profiles were then used to drive a higher-fidelity Simulink model. Due to the significant difference between the model structures of the MATLAB and Simulink models, comparison of results were not conclusive. A

different simulation approach is required to make this proof-of-concept more useful for controls development. This research forms the foundation for further studies.

Contents

List of Tables

List of Figures

List of Abbreviations

ANL	Argonne National Laboratories
ASM	dSPACE Automotive Simulation Model
AWD	All Wheel Drive
BAS	Belted Alternator Starter
CD	Charge-Depleting
CS	Charge-Sustaining
CVT	Continuously Variable Transmission
DP	Dynamic Programming
E-REV	Extended-Range Electric Vehicle
E85	Fuel mix of 85% ethanol, 15% gasoline
ECMS	Equivalent Consumption Minimization Strategy
EPA	US Environmental Protection Agency
ESS	Energy Storage System
EV	Electric Vehicle
GHG	Greenhouse Gas
GM	General Motors Corporation
HEV	Hybrid Electric Vehicle
ICE	Internal Combustion Engine

IGBT	Insulated-Gate Bipolar Junction Transistor
Li-ion	Lithium Ion
MBD	Model-Based Design
MG	Motor-Generator
NiMH	Nickel-Metal Hydride
PHEV	Plug-In Hybrid Electric Vehicle
PM	Permanent Magnet
PSD	Power Split Device
RTM	Rear Traction Motor
SOC	State-of-Charge
UF	Utility Factor
UVic	University of Victoria

ACKNOWLEDGEMENTS

The completion of this thesis was no easy feat, and I would like to recognize several individuals for their particular assistance. I would like to thank:

Teghan Godkin for your love, support, encouragement, and patience with me.

Dr. Zuomin Dong, for providing a graduate school opportunity and financial support, but most importantly tolerating my jokes.

Dr. Curran Crawford, for your guidance and advice on this work, and for introducing us to the wonders of Dog Joring.

David Killy, Jian Dong, and Daniel Prescott for your particular contributions to EcoCAR modelling and simulation activities, without which this thesis would not exist.

Dorothy Burrows for always brightening my day.

Susan Wignall for helping me stay in school.

To the optimist, the glass is half full. To the pessimist, the glass is half empty. To the engineer, the glass is twice as big as it needs to be.

Anon.

Chapter 1

Introduction

1.1 The Need for Change

From the beginning of large-scale consumer adoption of the automobile to the present day, the internal combustion engine ('ICE') has provided propulsive power to the vast majority of vehicles on the road. ICE's operate most efficiently in a relatively narrow range of torque and speed, though the typical consumer automobile engine frequently operates outside of this region due to the random events and variable conditions of daily commuting. This technology, powered in most cases by a range of liquid petroleum-based fuels, offers high power density but generally low efficiency. ICE's also generate significant amounts of combustion by-products in the form of gaseous emissions, including greenhouse gasses ('GHG'), carbon monoxide, mixed hydrocarbons, and other compounds.

In recent years, concerns have mounted over volatile fuel costs and rising emissions levels. Crude oil prices in Canada peaked at over $133 per barrel in late 2008, leading to historically high fuel costs for consumers [1]. In 2009, the transportation sector consumed 25% of Canadian energy resources, with automotive gasoline and diesel fuels accounting for 87% of sector energy use [2]. From 1990 to 2007, emissions due to Canadian motor vehicles rose by 35%, almost twice the growth rate of the population over that period [3]. The combination of volatile oil prices and heavy reliance on fuels for transportation has lead to a regulatory mandate for improving vehicular fuel efficiency. Additionally, an overwhelming body of evidence has mounted linking the rising fuel use and GHG emissions to global climate change, more extreme weather events, and other environmental disturbances [4]. Automobile manufacturers have

responded to these concerns, in part, with increased efforts in the development of hybrid, electric, and alternative-fuel vehicles.

1.2 The Hybrid Vehicle Revolution

Hybrid vehicles utilize two or more differing energy sources for propulsion, typically petroleum-based fuel and stored electricity. Hybrid powertrain technology offers a means to reduce fossil fuel consumption by providing a range of propulsion options, selectable on the basis of efficiency, to meet the wide range of operating conditions and varied power demands encountered by consumer automobiles. Additionally, kinetic energy recovery methods such as regenerative braking can be implemented to enhance efficiency, and direct displacement of petroleum fuel use is possible with increased vehicle electrification. The core component of a hybrid powertrain is the Energy Storage System ('ESS') which typically takes the form of a large high-voltage battery, isolated from the standard vehicle electrical system.

Hybrid vehicles vary widely in the degree to which fossil fuel usage is reduced or displaced. Different hybrid powertrains are typically defined according to their degree of electrification. At the low end of the electrification spectrum are 'mild hybrids, which replace the traditional ICE starter motor with a small energy-storage system and electric motor/generator ('MG'), allowing the ICE to be turned off during idling, coasting, or braking, and quickly restarted when needed. Full hybrids ('HEV') typically use an MG for assisting the engine, and often can provide short-duration electric-only operation. HEV's employ variants of the parallel and power-split architectures, which are discussed in the next chapter. The Plug-In Hybrid ('PHEV') enhances an HEV with a much larger ESS, capable of storing electrical energy from the utility grid for later propulsive use. The Extended-Range Electric Vehicle ('E-REV') typically operates as a full-function electric vehicle ('EV'), using an ICE (and a plug-in charger) to replenish their ESS upon its depletion or to assist the traction motor.

A table outlining the various hybrid powertrain types is given in Table 1.1.

1.3 EcoCAR 2: Plugging in to the Future

EcoCAR 2: Plugging in to the Future ('EcoCAR2') is a collegiate student design competition, running from 2011 to 2014, that challenges teams from 15 universities

Table 1.1: Hybrid Classifications with Commercial Examples

Classification	**EV Capability**	**Commercial Example**
Conventional (ICE)	None	Many
Mild Hybrid	ICE Start/Stop	Chevrolet Malibu
Hybrid Electric Vehicle	Limited	Toyota Prius
Plug-In Hybrid Electric Vehicle	Partial	Toyota Prius PHEV
Extended-Range Electric Vehicle	Full	Chevrolet Volt

across North America to re-engineer a donated mid-size sedan into a hybrid, with the goals of improving fuel economy and minimizing vehicle emissions while retaining performance and consumer appeal. Sponsored primarily by the US Department of Energy ('DOE') and General Motors ('GM'), and managed by the US DOE's Argonne National Laboratory ('ANL'), one of the primary goals of the program is to train a new generation of automotive engineers in the skills required for careers in the hybrid vehicle industry.

The University of Victoria ('UVic') was awarded participation in the program in 2011. For its competition vehicle, UVic is developing an all-wheel drive ('AWD') PHEV with an advanced series-parallel architecture, designed into the chassis of a 2013 Chevrolet Malibu. A GM 2.4L Ecotec engine, capable of running either gasoline or E85 ethanol fuel, is mated to a GM 6T40 6-speed automatic transmission and a TM4 80kW peak / 37kW continuous permanent magnet electric motor/generator, providing power to the front wheels. The rear axle of the car is powered by Magna E-Car Systems' E-Drive, which consists of a 90kW peak / 45kW continous permanent magnet motor/generator mated to a 7.82:1 ratio differential, combined into the same compact housing. The vehicle ESS is based around advanced lithium-ion-nanophosphate battery modules from A123 Systems, providing a total of 16.2kWh of energy storage capacity. This architecture was designed to allow significant electric-only propulsion capability, as well as flexibility in power delivery to ensure fuel economy improvement can be achieved under a wide range of operating conditions.

As a corollary to EcoCAR2 deliverables, many graduate-student research projects have been, and are being, pursued at UVic, in the field of hybrid vehicle powertrain and control systems research. In particular, the UVic team emphasis on control system functionality and flexibility has resulted in the implementation of leading-edge optimal control systems research.

1.4 Hybrid Vehicle Development Challenges

The addition of a hybrid powertrain to the vehicle design process adds significant complexity. The ICE, ESS, and electric machines must be selected and sized based on desired vehicle operating characteristics, such as fuel economy and performance. The addition of the ESS, electric machines, and the cooling and electrical systems that accompany them, present a challenge with respect to packaging and mechanical integration. The UVic EcoCAR2 vehicle is no exception; combining a large ESS with an ICE, BAS system, and a large electric motor coupled to the rear axle, it is an advanced hybrid vehicle with a flexible powertrain that offers significant opportunities for reduced fuel consumption using advanced control and optimization methods.

The selection of output power in a conventional or EV vehicle is a straightforward 1-dimensional problem, set by the driver power request. In contrast, a more complex hybrid powertrain presents a multi-dimensional problem, depending on the number and configuration of powertrain components. In this case, many possible solutions exist for producing motive power to meet the driver's request. Advanced control techniques are required to identify solutions that enable the performance potential of such an architecture, and the performance obtained is strongly dependent on the control method selected. In addition, obtaining the true maximum performance requires full knowledge of future driving conditions, which is not possible for a control system operating in real time. The identification of the upper bound on architecture performance is therefore a fairly complex problem.

To quantify the upper performance bound and derive operational benchmarks, the architecture must be analysed independently of any particular control methodology. This analysis must also be performed in a backward-looking format, where the desired driving cycle information is fully known in its entirety, to allow the inclusion of more globally optimal solutions. These benchmarks form a target performance metric to which the performance provided by different control strategies can be compared. Information such as the optimal control policies (component operating profiles), resulting fuel consumption in city and highway driving, and the ideal operating regime for given driving conditions also allows the quantitative assessment of the hybrid vehicle architecture and can be used in future deployment of control strategies.

1.5 Research Goals

Obtaining operational benchmarks requires significant effort in the areas of powertrain modelling and global optimization, and is the focus of this research. UVic's EcoCAR2 competition entry is used as the subject of the benchmarking analysis.

The goals of this work are to:

1. Summarize the UVic EcoCAR2 vehicle powertrain architecture and corresponding simulation model development.
2. Investigate the use of Dynamic Programming ('DP') in establishing benchmarks to aid the development of hybrid vehicle control strategies.
3. Employ DP to determine benchmark values for the EcoCAR2 vehicle's novel Series-Parallel powertrain architecture, including optimal fuel consumption and control policies, over a set of standard driving cycles.
4. Investigate the viability of using DP-generated control inputs in a forward-looking dynamic high-fidelity powertrain model.

1.6 Organization of the Thesis

The remainder of this thesis is organized as follows:

- Chapter 2 discusses the primary hybrid vehicle powertrain configurations, provides an overview of the UVic EcoCAR2 powertrain architecture selection process and the team's chosen architecture, and provides a summary literature review of hybrid powertrain technologies and controls methods, with a focus on the use of optimization in the development of advanced hybrid vehicle control strategies.
- Chapter 3 summarizes hybrid powertrain modelling and simulation techniques, including a discussion of model fidelity concepts and simulation environments. It then provides an overview of specific powertrain component models as used by the EcoCAR2 team.
- Chapter 4 discusses the context of this work, explains the fundamental techniques used, and summarizes the development and implementation of the dynamic programming algorithm used for the benchmark analysis.

- Chapter 5 examines the validation of the DP algorithm, the resulting vehicle operation over standard driving cycles, and computational issues regarding the algorithm. The implications of optimal operation are examined and discussed. The results of applying DP-derived control policies as control inputs in a higher-fidelity, forward-looking model are examined.
- Finally, Chapter 6 provides a summary of this work and offers recommendations for future research and development.

Chapter 2

Background

2.1 Hybrid Vehicle Technology Review

Many types of hybrid vehicles have been built and demonstrated, utilizing a wide range of energy sources and propulsive components. The most common variants employ an ICE in tandem with one or more MG and some type of on-board ESS [5]. The ESS can be electrical, mechanical, or chemical in nature. However, the use of a battery-based ESS is currently the most common variant, due to the falling cost and rising energy density of battery technology. Hybrid vehicles take advantage of the following technologies to improve overall fuel efficiency:

- Engine idle-stop - ICEs are typically very inefficient at idle speeds. By employing larger MG's (compared to standard automotive starter motors) to handle engine starting, very rapid and smooth ICE starts can be realized, thus allowing the ICE to be stopped during deceleration or idling events without affecting consumer acceptability or drive quality. During this time, the ESS can also supply auxiliary electrical loads more readily than the standard automotive 12-volt battery. Engine idle-stop systems alone can provide 4-10% improvement in fuel economy over a combined EPA test cycle, along with an approximately comparable reduction in CO_2 emissions [6]. Often, the larger MG also provides the low-voltage charging function of the alternator; such a configuration is called a Belt-Alternator-Starter ('BAS') system.

- ICE operating point control - In some hybrid architectures, the speed and/or torque of the ICE is de-coupled from that of the driving axle, allowing the

ICE to operate in more efficient torque/speed regions and thereby reduce fuel consumption. This arrangement can make use of advanced optimal control techniques to maximize ICE efficiency, and can allow more full utilization of a physically smaller ICE.

- Regenerative braking - The MG's can recapture kinetic energy from the vehicle during deceleration events, or elevation changes, storing it as electrical energy in the ESS for later use. This is the fuel-displacing mechanism of a standard HEV.

- Plug-In Charging - A PHEV can store energy from off-board sources, most notably the local electrical utility grid, in the ESS. Given enough ESS capacity, this allows the vehicle to operate as an EV over short driving distances, dramatically reducing the use of fossil fuels. This can have a large impact, as the majority of North American drivers live in urban areas, with relatively short commutes. The 2009 National Household Transportation Survey indicates that the distances the majority of US residents travel per day are relatively short: 12.09 miles (19.46km) for the average commute trip length of a private vehicle and approximately 35 miles (56.3 km) total distance travelled per person per day [7]. The ESS capacity required to achieve these distances is well within the range provided by current battery technologies.

In general, PHEV's and E-REV's are better equipped to take advantage of the latter two benefits as compared to standard HEV's; their larger ESS are typically associated with not only a greater energy storage capacity, but also higher power transfer capabilities. This means they are more capable of absorbing larger influxes of power during regenerative braking, and may also source more electric power, improving flexibility in terms of ICE operation.

During operation, the amount of energy stored in the ESS, referred to as the state of charge ('SOC'), may fluctuate, based on the energy source selected to power the vehicle. The way this fluctuation is allowed to occur over a drive is used to define the operating mode of the vehicle, as shown in Fig. 2.1.

Given a large ESS capacity and available electric drive power, a vehicle can operate in EV-only mode, where only electric power is used, and the vehicle operates essentially as an electric vehicle. Alternatively, a blend of fuel and electric energy can be used, biased primarily towards electric, that results in a net discharge of the ESS.

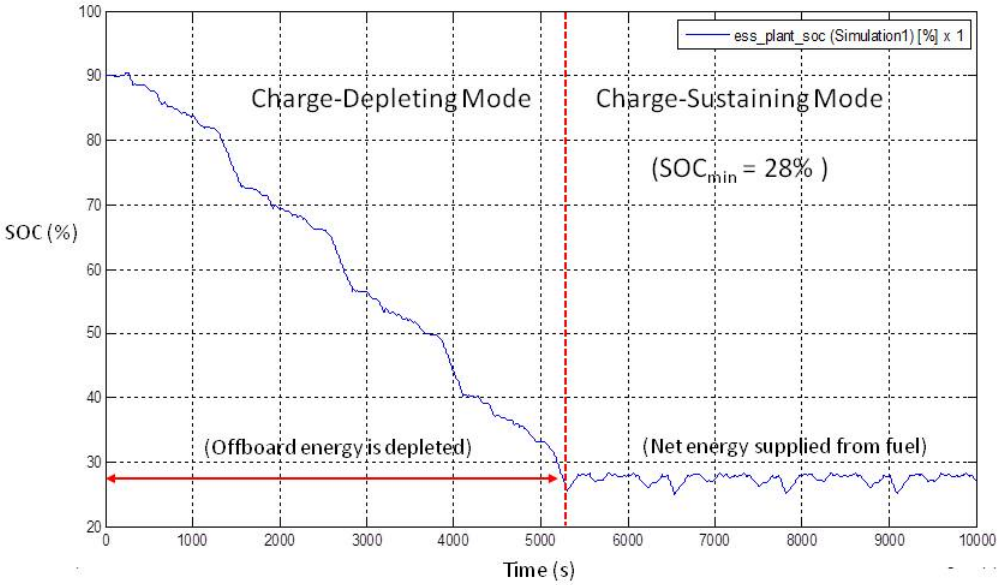

Figure 2.1: Modes of PHEV Operation - Simulink simulation (Image courtesy of UVic EcoCAR2 Team)

This is known as Charge-Depleting ('CD') mode. EV and CD modes can operate until certain conditions are encountered (for example, the ICE is required to meet overall power demand) or, typically, a low limit on ESS SOC is reached. At this point, all propulsive energy must be derived from fuel; any local-in-time fluctuations in ESS SOC, such as from a regenerative braking event, will be compensated for later in the drive. This is referred to as Charge-Sustaining ('CS') mode.

CD and EV modes are used in PHEV's and E-REV's to discharge stored electrical energy from an off-board source, reducing overall fuel consumption of the vehicle over a drive. All mild hybrid and HEV operation is in CS mode; although certain events during a drive cause an increase or decrease in SOC, no off-board energy is used. Note that when comparing energy consumption between different vehicles, ESS SOC in CS operation must be equal at the beginning and end of a drive, to ensure all fuel use is included in the comparison. This SOC restriction is not a requirement for real-world

operation.

2.1.1 Hybrid Vehicle Powertrain Architectures

The primary hybrid vehicle powertrain component configurations, referred to as 'architectures', are series, parallel, and power-split [5], each of which are discussed in turn below.

2.1.2 Series

Series powertrains are composed of an ICE and MG operating together as a generator-set, usually decoupled from the wheels. An additional MG is used to propel the vehicle. All ICE power is converted to electrical energy via the connected MG and, combined with energy from the ESS, is converted back to mechanical power with the traction MG and gearbox. This means that the propulsive power in the series hybrid is always 100% electric. A schematic of a series powertrain is shown below in Figure 2.2.

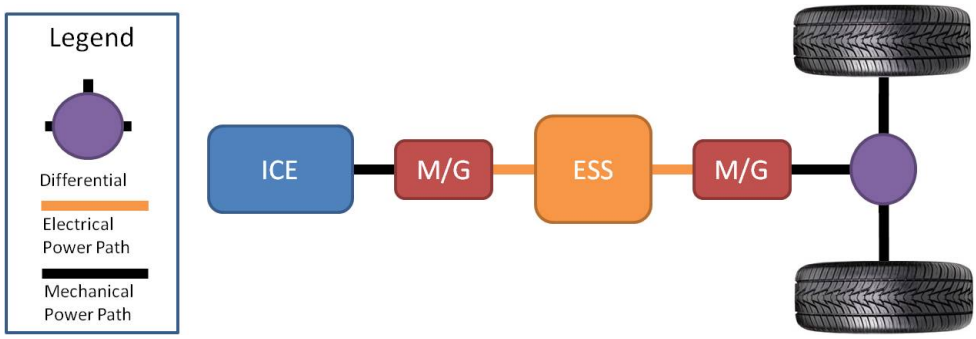

Figure 2.2: Schematic of Series Powertrain Architecture

Series hybrid powertrains have been utilized for many years in industrial and military applications, such as mining haul trucks and submarines. In the typical series configuration, the ICE is completely decoupled from the wheels. This allows the ICE operating point to be selected independently of wheel speed and torque requirements, ensuring that the ICE can always be run at optimal combinations of torque and speed. Additionally, the ICE in a series vehicle can be sized smaller than in a conventional equivalent, and the ICE power more fully utilized. The lack of mechanical connection between ICE and wheels eliminates the need for a transmission, and somewhat reduces the packaging complexity of the vehicle.

The two-stage energy conversion process in the series architecture imposes some additional electro-mechanical conversion losses. This means that the series is most useful in urban driving conditions, with lower average speeds and frequent starts and stops. For highway driving or on a steep grade, the average efficiency of the series is low compared to other hybrid architectures. The propulsive MG must also be powerful enough to meet all vehicle performance specifications, adding significant weight and cost.

The Chevrolet Volt is, to date, the most notable widely-produced series-like consumer hybrid vehicle. Not strictly a series, the Volt is built as an E-REV, with a large lithium-ion battery pack providing 9.4kWh of usable energy to a 100kW traction motor. A small 1.4L ICE drives a generator to provide range-extender capabilities. Using a series of clutches, the ICE can also be coupled to the wheels directly, to eliminate high-speed conversion losses [8]. More exotic ICE types have also been employed in series hybrids; an E-REV bus utilizing a micro-turbine as a range-extender has been developed and trialled [9].

2.1.3 Parallel

In a parallel powertrain, both the ICE and MG are mechanically coupled to the wheels of the vehicle. The vehicle can be propelled by either the ICE or the MG, or a combination of both. A schematic of a parallel powertrain is shown below in Figure 2.3.

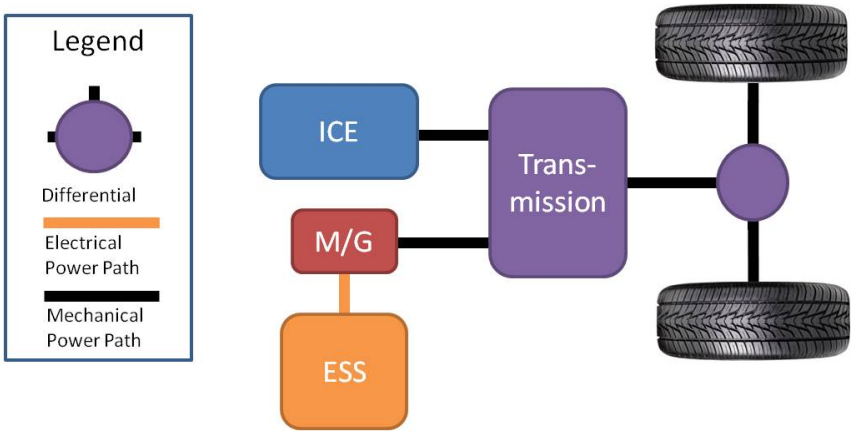

Figure 2.3: Schematic of a Parallel Powertrain Architecture

The advantages of a parallel hybrid are its operational flexibility and reduced

energy conversion losses. Since propulsive power can be produced from a combination of electrical and mechanical pathways, the system can be designed to efficiently meet different operating conditions. The direct mechanical coupling between the ICE and the wheels eliminates the conversion losses associated with the series architecture, which can become significant at higher power levels. A common strategy is to utilize the MG to propel the vehicle at lower speeds, taking advantage of the higher low-speed efficiency of electric drive systems over ICE's, and switch to ICE propulsion at higher speeds or power levels. While the ICE is running, the MG can be used as a generator, to replenish the ESS. When the ICE is not running, it must be decoupled from the powertrain or operated in such a way as to minimize pumping losses. As the ICE is connected to the wheels through a transmission, its operating point cannot be selected optimally, resulting in higher fuel consumption. The use of a continuously-variable transmission ('CVT') can mitigate this somewhat.

Several automotive OEM's have been manufacturing commercially successful parallel hybrids. Starting in 2011, Hyundai began producing a hybrid version of the Sonata mid-size sedan, employing a 30kW electric motor in a parallel architecture. This system, termed Hyundai Blue Drive, allowed the vehicle to achieve a combined 36mpg, a significant improvement over the conventional Sonata's 28 mpg [10]. Honda also markets a parallel hybrid, the Honda Insight, which utilizes their Integrated Motor Assist technology. The MG, a thin brushless permanent magnet DC motor, is directly coupled to the 1.0L ICE's crankshaft in place of a flywheel, and performs ICE assist and idle-stop functions [11]. The 2013 Insight offers a combined fuel economy of 42 mpg [10].

2.1.4 Power-Split

A power-split HEV utilizes an ICE and two MG's coupled together through a power-splitting device ('PSD'), most often consisting of a planetary gear set. The term 'split' refers to the PSD's action of dividing the ICE input power into mechanical and electrical power paths within the transmission. The PSD also allows the ICE to operate decoupled from the wheels, in effect functioning as a CVT. A schematic of a power-split powertrain is shown below in Figure 2.3.

Power-split's offer a number of advantages over other architectures. Since the ICE can operate decoupled from the road, its torque and speed setpoints can be manipulated to provide more efficient operation. If not needed, the ICE can be shut

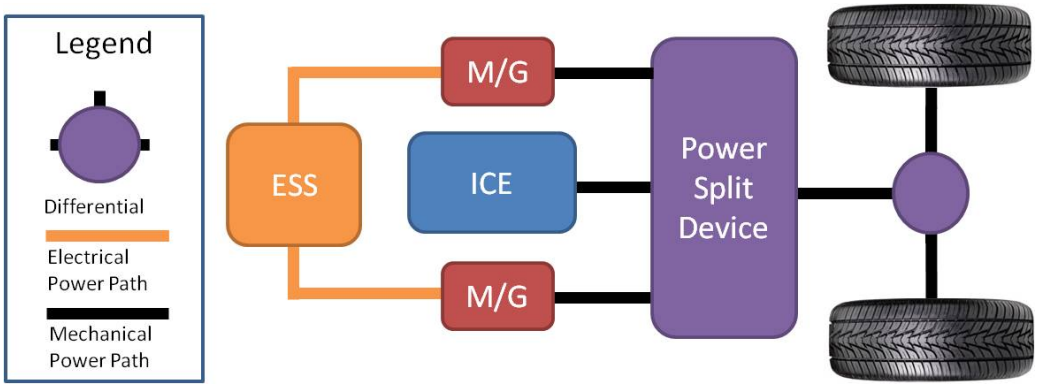

Figure 2.4: Schematic of Power-split Architecture

down, allowing EV-only operation. This flexibility comes at the cost of increased mechanical design and control system complexity. Wishart et al. describe input-split, output-split, and compound split designs, and analyse two commercially successful power-split architectures - Toyota's *Toyota Hybrid System* (now called *Hybrid Synergy Drive*) and GM's *2-Mode* transmission [12]. The HSD has been used in several Toyota models, most notably the well-known Toyota Prius, since 1997, and consists of a single PSD and two M/G's in an input-split configuration. The single power-split mode means that the power path cannot be adjusted for vehicle loading, and HSD-equipped vehicles achieve noticeably reduced fuel economy on steep grades and higher speeds because of this limitation.

The GM Two-Mode was designed to offer significant improvement over the one-mode HSD, by adding an additional PSD and clutches to enable more power paths through the transmission. The Two-Mode features four fixed gears and two power-split modes to be utilized, allowing for a wide range of high-efficiency operation. The GM Two-Mode is currently in service in hybrid variants of the GMC Yukon, Cadillac Escalade, and Chevrolet Tahoe. A donated Two-Mode transmission was also utilized by the University of Victoria to develop a powerful, advanced power-split E-REV architecture for the inaugural EcoCAR competition [13].

2.1.5 Energy Storage Technology

The ESS is an expensive and very critical component of a hybrid powertrain, and many different technologies have been implemented. Especially in the case of PHEVs, the ESS must provide high energy density, high charging capacity (particularly for the

power surges during regenerative braking), long battery life to reduce maintenance costs, and be robust and reliable.

The primary energy storage technology currently being utilized are different chemistries of rechargeable batteries. Chemical batteries typically have high energy densities and store the majority of on-board electric energy. Nickel metal hydride (NiMH) and lithium ion (Li-ion) and their derivatives are the most adopted batteries for HEV's today [14].

The energy density of a NiMH battery is twice that of a typical automotive lead-acid battery and is safe to operate at high voltages and across wide temperature ranges. NiMH batteries have long life cycles when operated in shallow cycles (20% to 50% of rated SOC) and are the most common battery found in HEV's today. In fact, the $600 million NiMH HEV pack business accounted for half of the total Ni-MH business in 2006 [14].

Nickel-Zinc (Ni-Zn) and Nickel-Cadmium (Ni-Cd) batteries are other Nickel-based battery technologies under development. Ni-Zn batteries utilize low-cost, environmentally friendly materials, and can safely be deep cycled, but suffer from poor operating life cycles, which is unfavourable in vehicular applications. Ni-Cd batteries are recyclable, can be fully discharged without damage, and have a long usable life, but currently come with a large price tag making them uneconomical for HEVs [14].

Li-ion batteries have demonstrated excellent performance in ESS applications, providing high energy density and good high temperature performance compared to other batteries for many HEV prototypes. Typically built with a cobalt-oxide cathode material [15], Li-ion batteries have twice the energy density of NiMH and show much less output voltage variation across a wide range of SOC, and are replacing NiMH in hybrid applications.

Variations on the basic Li-ion battery involve subtle chemistry changes, and different cathode materials. A123 Systems manufactures a proprietary Lithium Iron Nanophosphate battery, with a special nanoscale engineered material for the cathode. The nanophosphate material provides greater cathode surface area resulting in higher charge/dis more chemically stable discharge power capabilities. The material is also chemically stable, reducing the chances of energetic thermal runaway under extreme conditions [16]. Another variant, used in Nissan Leaf EV's, is Lithium Manganese Oxide, which exhibits high power capability, but lower energy capacity, and higher capacity loss on storage or cycling due to manganese dissolution in the electrolyte at high temperatures [15].

Flywheel-based ESS systems have recently come under investigation as an alternative to batteries. The rotating flywheel stores energy mechanically, and is coupled to the vehicle's electrical bus via a motor/generator that accelerates or decelerates the high-inertia rotating mass of the flywheel [17]. To reduce losses, flywheels are commonly run in a vacuum, with magnetic bearings and very high efficiency electric machines. The high cost of these components make present-day flywheel ESS more expensive than battery alternatives. However, flywheel ESS have been designed for use in large vehicles (i.e. transit buses) where battery costs are already inherently high [18], and have been used in Formula 1 racing for several years. The typical Formula 1 application, known under the race rules as a Kinetic Energy Recovery System, stores energy during braking, and makes it available as a 60kW 'power boost' for up to 6.6s per lap [19].

The use of ultracapacitors for on-board energy storage is a new area of research and development. Although ultracapacitors still provide significantly inferior energy storage density compared to batteries, the lack of chemical variations on the electrodes means they are able to withstand several orders of magnitude more charge/discharge cycles over their lifetime, and offer much higher power density. These advantages have made ultracapacitors useful as ESS buffers, improving the dynamic response of the ESS and protecting it from harmful operation by allowing rapid and high-magnitude charge/discharge transients to bypass the batteries (e.g. during quick regenerative braking) [14].

2.1.6 Electric Drive Technology

A electric drive system of a hybrid, consisting of a combined power electronics converter and controller, and a motor, is responsible for transmitting the power stored in the ESS to the road, and returning energy to the ESS via regenerative braking. Electric drives for hybrid vehicle applications must exhibit fast response, ease of control, low electrical and acoustic noise, and relatively high efficiency over a wide set of operating conditions. [20]. With the space and weight constraints imposed by HEV's, developing smaller and higher power electric drive systems within an acceptable cost limit is essential for large scale production and uptake of these vehicles.

In recent years, power electronic converters have matured considerably in robustness, cost, and packaging, driven largely by advances in gate-controlled power switches [21]. New and improved semiconductor devices have been developed, most notably

the Insulated Gate Bipolar Transistor ('IGBT') which combines the ruggedness of the well-established bipolar junction transistor with the simple gate drive requirements, high switching speed, and gate terminal electrical isolation of a MOSFET [21]. The increased switching speed and ease of fine control has, in turn, allowed for more complex control techniques (e.g. Vector, D-Q, and Direct Torque control) to be employed, optimizing torque production through more precise magnetic field manipulation. Converter systems for road vehicle applications commonly take the form of a boost converter maintaining a DC link, in series with a three-phase voltage-source inverter in a 6-switch bridge topology [21]. On vehicles with multiple drives sharing an electrical bus, EMI filters are included to reduce crossover electrical noise and prevent catastrophic resonant effects.

Electric machines are differentiated primarily by the means of electrical excitation. DC motors, historically the primary choice in traction applications, can run directly off a DC bus, but require mechanical commutator brushes to switch field polarity, providing a potential maintenance issue. Induction machines are common and inexpensive, but are physically larger at high powers than competing technologies, and require an inverter to drive them. Permanent magnet synchronous motors ('PMSM') replace the induction motor's 'squirrel-cage' rotor with powerful rare-earth permanent magnets, leading to much greater power density and size reduction of the motor, but at much higher cost. All EV's and HEV's currently produced by major automotive OEM's utilize some form of PM synchronous motor.

2.2 UVic EcoCAR 2 Architecture Selection

The University of Victoria's entry to the EcoCAR 2 competition has been termed a Series-Parallel Multiple-Regime PHEV. Multiple-regime refers to the ability and intent to operate the powertrain in series, parallel, ICE-only, or EV-only configurations, depending on driving conditions and availability of fuel and electrical energy. This multi-regime capability is central to the team's design strategy of powertrain flexibility. The vehicle is designed using the 2013 Chevrolet Malibu as a base platform. A CAD representation of the vehicle and its major powertrain components is shown in Fig. 2.5, an architecture power flow diagram is given in 2.6, and component specifications are given in Table 2.1.

Architecture selection was driven by the breakdown of EcoCAR2 competition scoring, as well as the past experience of the team. The scoring methodology em-

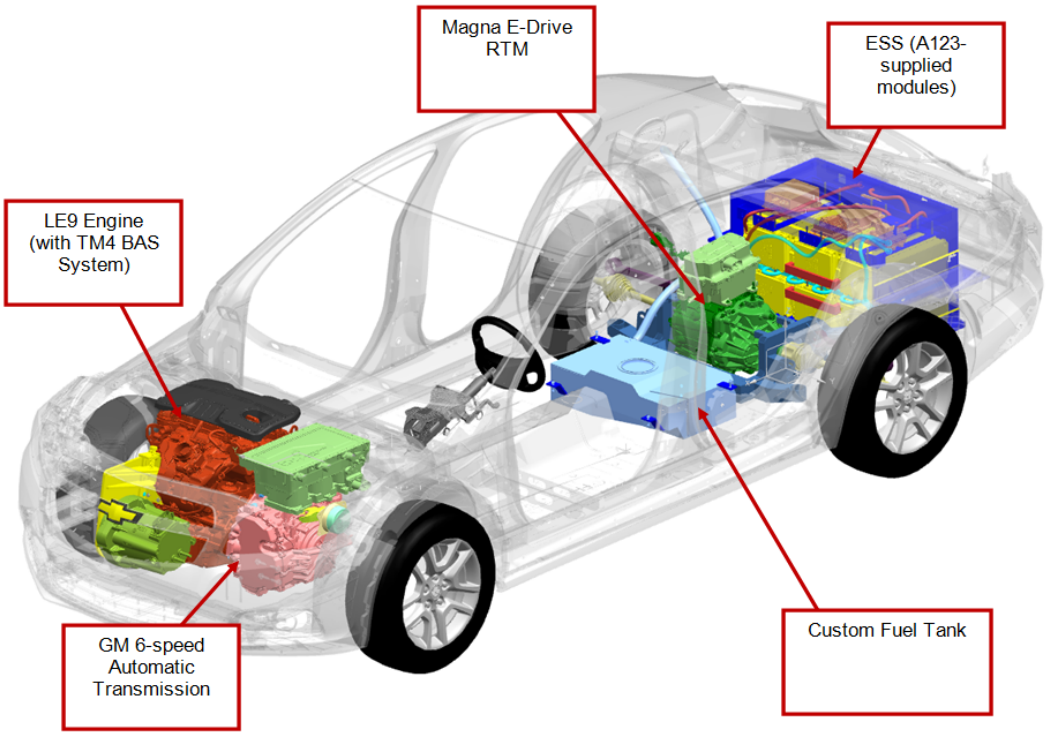

Figure 2.5: UVic EcoCAR2 Powertrain Components. Retrieved from "An Innovative 4WD PHEV Utilizing a Series-Parallel Multiple-Regime Architecture", Kaban, S., Nelford, J., Dong, Z., Dong, J. et al., 2012, *SAE Int. J. Alt. Power*

phasizes safe on-road operation, maximum fuel efficiency and dynamic performance, along with minimum emissions. A simplified scoring breakdown is given in Table 2.2.

Events relating to emissions and energy consumption make up 40% of the total points from the competition. It is clear that selecting a hybrid architecture and components from the perspective of maximizing scoring in these areas would yield the greatest benefit in terms of competition performance. This approach was a strong influence in fuel selection and ESS design.

2.2.1 Fuel Selection

Fuel selection was performed based on an analytical comparison of the available competition fuels: gasoline (E10), biodiesel blend (B20), ethanol blend (E85), and hydrogen. Emissions and energy values were derived using GREET (for US energy mix) and GHGenius (for Canadian energy mix) with 2018 as the target year (theoretically, the year the team's vehicle would enter the market if it were being produced), and

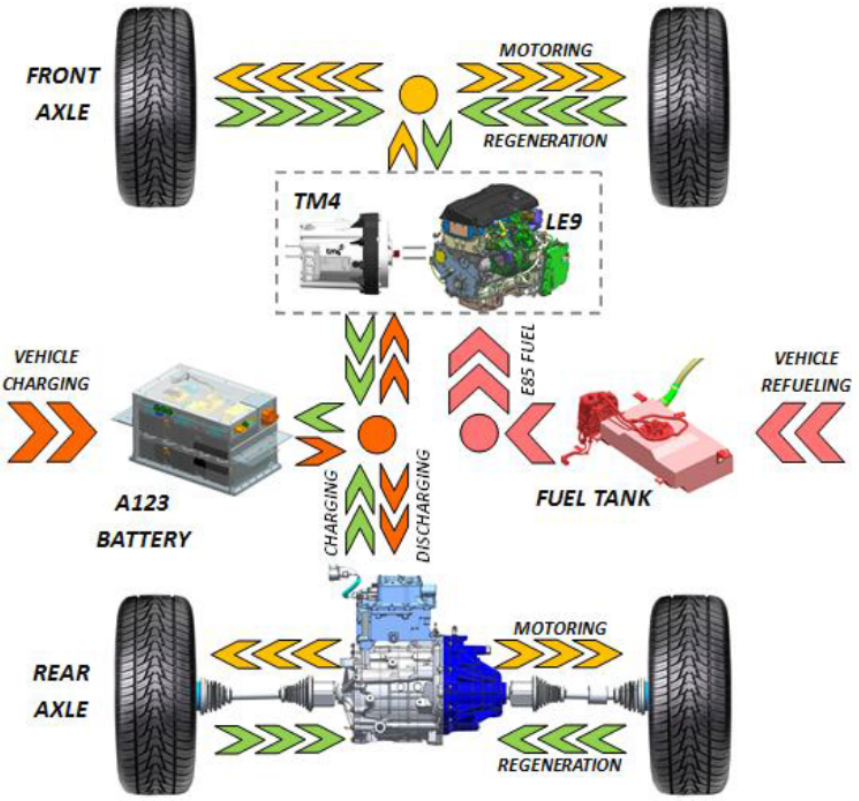

Figure 2.6: UVic EcoCAR2 Power Flow Diagram. Retrieved from "An Innovative 4WD PHEV Utilizing a Series-Parallel Multiple-Regime Architecture", Kaban, S., Nelford, J., Dong, Z., Dong, J. et al., 2012, *SAE Int. J. Alt. Power*

are given below in Table 2.3. Electricity is also included as a fuel, for comparison purposes, but was not considered as a primary fuel choice, as no currently available battery chemistry can compare with liquid fuels on the basis of energy density.

Though B20 offers the highest energy density, the GHG emissions from using it are higher compared to E85. Also, E85 offers the lowest Well-to-Wheels petroleum energy use, due to its low (15%) petroleum content. Electricity offers an even lower petroleum energy content, which suggests that designing a vehicle to have a long electric-only range is a viable strategy for further reducing petroleum energy use, with the trade-off of increased GHG emissions. In the category of criteria (CAC) emissions (hydrocarbon content, CO, and NO_x), E85 and B20 offer the lowest upstream emissions values. However, minimizing the downstream pollutant emissions from B20 would require

Major Component	Specifications
Vehicle	2013 Chevrolet Malibu
GM LE9 2.4L Ecotec Engine	130kW @ 5800rpm
	230Nm @ 5000rpm
GM 6T40 Transmission	6-Speed Automatic
	3.71:1 final drive
TM4 Motive Drive (BAS)	80kW peak, 37kW continuous
	170Nm peak, 65Nm continuous
Magna E-Drive (RTM)	90kW peak, 45kW continuous
	245Nm peak, 150Nm continuous
	7.82:1 final drive
A123 Systems ESS	330V peak, 292V nominal,
	16.2kWh max capacity, 14.5kWh useable
	152kW peak, 51kW continuous

Table 2.1: UVic EcoCAR Powertrain Component Specifications

Event	Points Percentage (Year 2)
0-60mph Acceleration	3.25
50-70mph Acceleration	3.25
Autocross	4
Energy Consumption	10
Petroleum Energy Use	10
Criteria Emissions	10
WTW and GHG Emissions	10
Driving Events Total	**63.8**
Non-Driving Events Total	36.2

Table 2.2: UVic EcoCAR Powertrain Component Specifications

the design and integration of a complex exhaust after-treatment and particulate filter system. E85 was the final choice of fuel, due to its high petroleum displacement, low GHG and CAC emissions, and the UVic EcoCAR team's prior experience with E85 engine operation.

2.2.2 ESS Design using SAE Utility Factor

The ESS selected for the vehicle is comprised of modules and electronics units donated by A123 Systems, one of the major EC2 competition sponsors. The pack consists of 6 series-connected modules, each containing 3 parallel strings of 16 cells utilizing a proprietary Lithium-Iron-Phosphate chemistry, providing a nominal voltage of 292V

	E10	**E85**	**B20**	**Elec**
Fuel-specific energy by mass (kWh/kg)	11.44	7.96	11.55	N/A
Fuel density (kg/L)	0.746	0.7871	0.8552	N/A
Fuel energy density by volume (kWh/L)	8.534	6.265	9.878	N/A
% Petroleum	90%	15%	80%	0%
Downstream g CO_2 / kWh	261	260	277	N/A
GHG_{WTW} (g/kWh)	322	261	288	648
Vehicle PEU_{WTW} (kWh PE/kWh)	0.984	0.316	0.859	0.034
Upstream THC (g/kWh)	0.0612	0.0475	0.0101	0.0031
Upstream CO (g/kWh)	0.0119	0.005	0.0091	0.0326
Upstream NO_x (g/kWh)	0.0279	0.0141	0.0214	0.1012

Table 2.3: EcoCAR 2 Competition Fuel Properties

and an energy storage capacity of 16.2kWh. These modules were selected primarily to obtain a large on-board energy storage capability, in order to maximize vehicle Utility Factor ('UF').

The UF essentially describes the fraction of commuters that would have their daily driving requirements met by a given vehicular CD range. As the CD range is increased, more fuel is displaced by electric energy, and the average fuel consumption rating of that vehicle decreases. The UF weighting used in EcoCAR 2, shown in Fig. 2.7 is derived from SAE standards J2841 and J1711 [22], and specifically accounts for a PHEVs electric-only range capability while operating in CD mode. UF-weighted fuel consumption is a weighted average value, calculated according to the function

$$FC_{UF} = (UF) * FC_{CD} + (1 - UF) * FC_{CS}$$

where FC_{CD} and FC_{CS} are the charge-depleting and charge-sustaining fuel consumption, respectively. The drive cycle used to test fuel consumption in the EcoCAR 2 competition is a longer cycle specifically calibrated to this UF curve, to more clearly separate periods of CD and CS operation. This would suggest that in order to improve the vehicle's scoring, the pack should be designed to be as large (in terms of energy storage) as possible, in order to offset fuel usage. However, a constraint is imposed by space in the vehicle available for pack integration, as well as the mass increase resulting from carrying large sets of batteries and associated equipment.

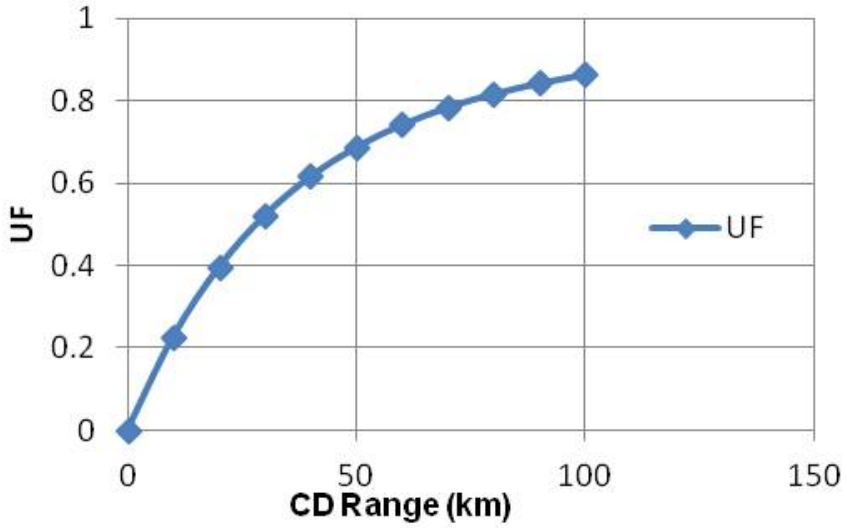

Figure 2.7: EcoCAR2 Utility Factor Curve

2.2.3 Component Selection and Sizing

Engine

Although a wide range of engine sizes and types are available on the market, it was decided to utilize one of the three engines provided for the competition by GM, to benefit from the accompanying technical support, available performance data, and CAD models. The 2.4L Ecotec LE9, the 1.4L LUJ, and the 1.8L LUD were compared in terms of power output, mechanical integration difficulty, and fuel compatibility.

Specification	**LE9**	**LUJ**	**LUD**
Power (kW)	130	104	120
Torque (Nm)	230	167	200
Fuel	E85/Gas	Gas	B20/Diesel
Displacement (L)	2.4	1.4	1.8
Volume (relative to LE9)	1.0	0.75	0.85

Table 2.4: Specifications of Different ICE Options

The GM LE9 was selected for the vehicles engine. The donated base vehicle, a 2013 Chevrolet Malibu, already ships with the LE9, which will reduce the complexity of engine mount modification. The LE9 also natively mates with the selected 6-speed automatic transmission, also from GM. The LE9 is Flexfuel-capable, meaning it can run on gasoline or E85, the fuel of choice, without modification. The higher power

and torque outputs of the LE9 will greatly benefit the vehicle in performance-oriented events such as acceleration or autocross.

Transmission

The GM 6T40 6-speed automatic transmission was selected for the vehicle. The 6T40 is compact and robust, and natively mates with the LE9 engine, which simplifies mechanical integration. The accompanying torque converter also allows lock-up between input and output, eliminating power loss during steady-state operation. The transmission shifting schedules are programmed directly into the transmission controller, and cause a gear shift based on vehicle speed, engine speed, and driver torque demand. It is likely that these schedules will not be flexible, in the sense that the vehicle supervisory controller cannot directly request a shift, meaning that ICE operating torque and speed in parallel regime cannot be directly modified through transmission control. The gear ratios of the 6T40 are given in Table 2.5.

Gear	**Ratio**
1	4.584
2	2.964
3	1.912
4	1.446
5	1.000
6	0.746

Table 2.5: 6T40 Transmission Gear Ratios

Electric Machines

The choice of electric drive components was made based on vehicle power requirements, and a trade-off of physical size and efficient operation. While most electric machines have a very high peak efficiency rating (>90%), this occurs at higher values of torque production. When cruising at a steady speed, the torque demand is low, leading to less efficient operation. With this in mind, motors were selected based on their expected efficiency while operating at key vehicle speeds, such as 50km/hr (urban cruising) and 90km/hr (highway cruising).

The Magna E-Drive, shown in Fig. 2.8, was selected as the rear traction motor ('RTM') of the vehicle. This component was originally designed and built to serve as the traction drive system for the Ford Focus EV. Donated by Magna Inc. for the

EcoCAR competition, the E-Drive provides an innovative integrated design, where the motor, inverter, and a single-speed differential are combined into a single package. This component, while not without integration challenge, offers the opportunity to utilize a unique turn-key solution backed by Magna technical support. An analysis of efficiency also showed that the E-Drive performs comparably to other top-end electric drive systems on the market.

Figure 2.8: Magna E-Drive - Integrated inverter, motor, and differential (Image courtesy of UVic EcoCAR2 Team)

To provide the generator function for the series operating regime, and to enable ICE idle-stop to be utilized, a large EM was selected to form the core of a custom BAS system. The TM4 Motive A, shown in Fig. 2.9, was selected as the BAS motor. This unit operates at the voltage levels of the selected ESS, and uses supervisory control via CANbus, making it compatible with other vehicle systems in electrically. As it is mounted in the crowded engine bay, making a physically small motor more desirable, most smaller PM drive systems operate at lower voltage levels (150 Volts). An in-house designed belt system was implemented to couple the BAS motor and the ICE via the ICE crankshaft pulley. This system is shown in Fig. 2.10.

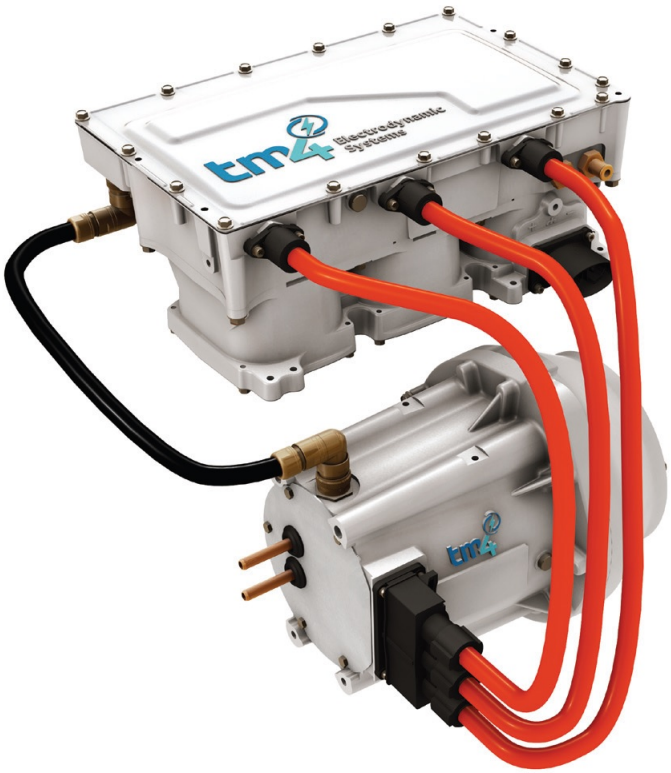

Figure 2.9: TM4 Motive-A Drive System (Inverter and Motor)(Image courtesy of UVic EcoCAR2 Team)

2.3 Hybrid Vehicle Control Strategies

A significant amount of research and development has been conducted in the area of hybrid powertrain control, especially as more flexible and complex architectures and components are developed. Managing energy transfer in a vehicle powertrain in a safe and efficient way is a complex topic that requires significant analysis and design work. The most commonly used hybrid control strategies will be briefly reviewed here.

Control strategies can be broadly divided into rule-based methods and optimization-based methods [23], with each of these groups having two distinct sub-groups, as shown in Fig. 2.11.

2.3.1 Deterministic Rule-Based Control

Rule-based control has been the traditional control methodology used in the automotive industry. A rule-based strategy typically uses 'if-then' and 'switch' logic, based on

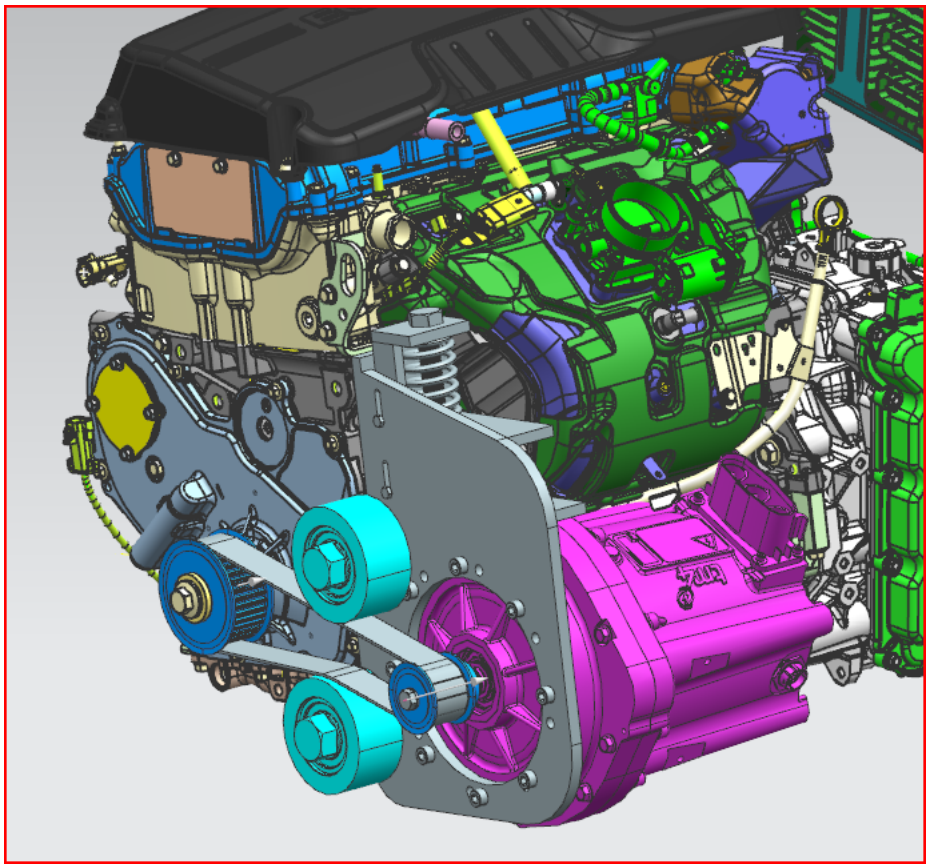

Figure 2.10: BAS belt system, NX CAD rendering (Image courtesy of UVic EcoCAR2 Team)

mathematical models, sets of heuristics, or some other set of hard-coded behaviours defined by current operating conditions and constraints.

The simplest rule-based controller is the hysteresis (or 'bang-bang') controller, which turns a component or process on or off based on a feedback signal crossing some predetermined threshold. State machines can be used to give structure to the operating rules, making the supervisory control system more resilient against faults. Phillips, Jankovic, and Bailey [24] implemented such a system for a parallel architecture, noting the advantages of ready understanding by engineering personnel and easy adaptation to different architectures and component configurations. Early Toyota Prius and Honda Insight HEVs made use of the Power Follower method, to sustain the charge in the ESS by supplementing ICE power with electric power as needed [25]. Although these methods are practical and have been successfully implemented, they are limited by their inflexibility. Careful tuning to particular

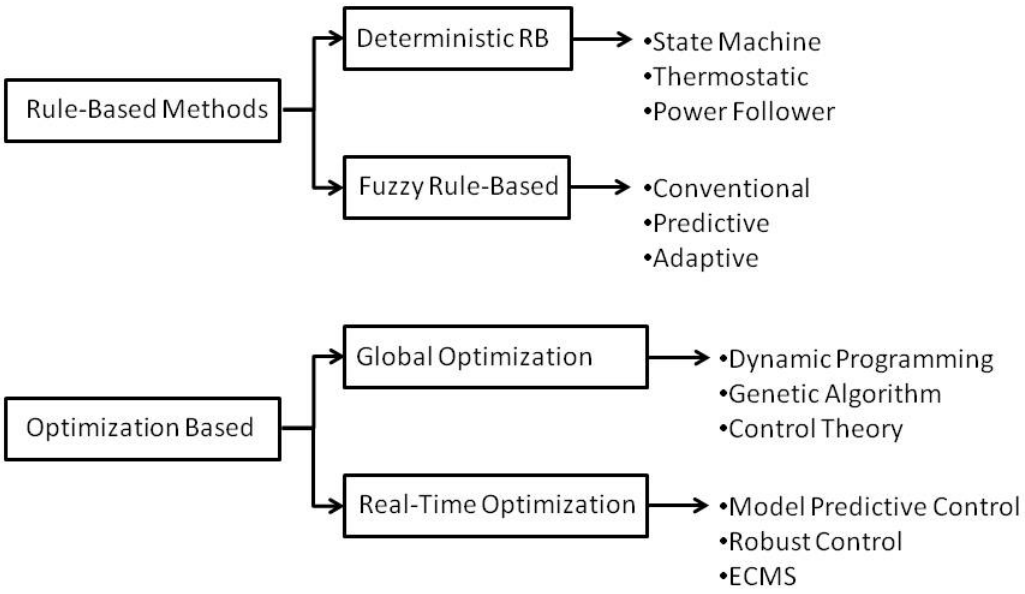

Figure 2.11: Control Strategy Classifications

operating conditions is required to ensure anything close to optimal operation is achieved. To accommodate a wider range of operating conditions, this tuning must be relaxed, resulting in poorer average performance.

2.3.2 Fuzzy Logic Control

Fuzzy logic was developed to address the performance limitations of classical logic when applied to probabilistic problem sets. Rather than the 'in' or 'out' membership, items in a 'fuzzy set' can exhibit a varying degree of membership. Many real-world problems can be addressed with this method, such as temperature control in a room; rather than being only 'hot' or 'cold', the controller input could be 'somewhat hot' or 'very cold', and the desired output (heater control signal) would be 'low cooling' or 'high heat', respectively.

Applied to hybrid powertrain control, fuzzy logic functions as an improved ex-

tension of a rule-based controller. Fuzzy controllers are more robust, as the use of fuzzy sets adds adaptability to imprecise measurements or component variations [23]. Schouten *et al* used a fuzzy logic controller to determine electrical-mechanical power split in a post-transmission parallel vehicle, noting that the controller was able to keep both ICE and EM in ideal operating areas [26]. Many other examples of fuzzy logic use in powertrain control are available in the literature [27] [28]. Although a significant improvement over hysteresis control, fuzzy logic controllers still lack the ability to explicitly determine optimal or pareto-optimal control parameters.

2.3.3 Optimal Control

Optimization techniques can be incorporated into a control strategy by employing an objective function representing fuel economy or emissions levels to calculate the optimum values of control outputs. In contrast with rule-based strategies, optimization-based control strategies directly determine the desired component operating point(s), and can maximize performance of the vehicle as a whole, rather than focusing on individual subsystems. The primary difference between the global- and real-time subgroups shown in Fig.2.11 is that global optimization strategies are determined over a complete driving cycle, necessitating *a priori* knowledge of vehicle operating conditions and restricting them to off-line optimization. Real-time strategies, as their name implies, are suitable for on-board use where this information is not available.

2.3.4 Real-Time Optimization

Real-time optimization strategies are a major topic of research, as they offer a means to utilize a complex hybrid architecture to its full potential. Under real driving conditions, future operating points are not known in advance, and so a true optimal solution cannot be found. Instead, real-time optimization strategies obtain sub-optimal solutions at discrete time intervals to manage power flows in the vehicle. These methods are computationally intensive, requiring high processing power and careful and efficient algorithm design to be implemented in practice. Real-time optimization strategies can also adversely affect vehicle driveability, from a consumer acceptability perspective. The sequence of instantaneous optimization results at subsequent time intervals can result in rapid changes of component operating points, resulting in an unusual driving experience for the end user - in testing a real-time optimization strategy for a hybrid vehicle, Waldner noted 364 gear shifts over a 30 minute period while

driving a concatenated UDDS/HWFET cycle [29].

The Equivalent Consumption Minimization Strategy ('ECMS') is a real-time optimization strategy developed by Paginelli *et al.* This method takes into account that to maintain SOC, any EM use at a given time will require some future ICE use, resulting in an 'equivalent fuel consumption' value [30]. This value is combined with the typical ICE fuel consumption value, and the value is optimized at some time interval to yield desired component operating points. With careful tuning, an early ECMS implementation showed a 17.5% reduction in fuel consumption over an ICE-only version of the same vehicle [30]. However, careful tuning is required to achieve optimal results, as the equivalence parameter relating fuel and electric use is very sensitive to drive cycle variations [31]. Waldner further noted that an ECMS algorithm constantly 'chasing' optimal operating points frequently passes components through non-optimal operating areas, resulting in worse fuel economy than with a much simpler rule-based system [29]. Musardo *et al* implemented an improved version, Adaptive ECMS, that incorporates a driving condition predictor, such that a mission window is built using past and present driving conditions, and equivalence parameters are optimally determined over that window [31]. Musardo reports the A-ECMS algorithm as attaining fuel economies very close to those provided when the optimal equivalence factor is calculated with *a priori* knowledge of the driving cycle.

2.3.5 Global Optimization

Determining a truly global optimal solution requires knowledge of the vehicle's operating conditions over a complete driving cycle. As it is not possible to obtain this information using real time control, global optimization is not directly applicable as a control strategy. Instead, global optimization techniques are typically used during vehicle design stages, to establish performance benchmarks with which to compare the performance of other control strategies [23]. Several global optimization methods have been successfully applied to this type of analysis.

Tate and Boyd [32] pioneered this approach by using linear programming to optimally determine a series hybrid powertrain's performance characteristics independent of any specific control laws. This method requires the formulation of a piecewise-linear model, involving many approximations, to describe vehicle operation, making it not suitable for more complex powertrain architectures. Piccolo *et al* [33] determined an optimal control strategy for a parallel architecture, using the genetic algorithm, to

minimize CO emissions and fuel consumption. Gielniak and Shen [34] performed a similar analysis of a fuel cell hybrid vehicle using game theory, and found that the algorithm needed to be very carefully tailored to a specific architecture and powertrain component set to return an optimal solution. Particular swarm optimization [35], simulated annealing [36], optimal control theory [37], and dynamic programming have also been successfully used for this type of analysis.

Dynamic Programming is a method for finding the globally optimal solution to a multi-stage problem, by breaking it down into simpler sub-problems and iteratively solving them. When used with an applicable problem, DP offers dramatically reduced computation time compared to brute force, and many other graph search methods, as it avoids repetitious evaluations of similar sub-problems and excludes infeasible solution paths, only searching over admissible state or control values [38]. That being said, DP still requires storing all valid state transition costs, resulting in relatively high memory requirements. Although similar search methods exist (ex. A* ('A-star')) that also guarantee an optimal solution, DP is well-suited to deal with overlapping sub-problems, such as a performance analysis where vehicle state at consecutive time steps is dependent on state at previous steps.

As mentioned above, dynamic programming has been successfully used to determine the globally-optimal performance of vehicle architectures, with the results often being used to inform the development of rule-based control systems. Wang and Lukic used DP to evaluate an HEV model with Toyota's Hybrid System for optimal performance, and created a lookup table-based controller from the results, realizing a 27% improvement in overall efficiency [39]. Dokuyucu and Cakmakci obtained optimal response characteristics for energy management and vehicle stability for a parallel powertrain [40]. These were considered concurrently, and the authors noted promising benefits to using DP for rule extraction in this type of problem. Others have tried a similar approach for series [41] and power-split [42] powertrains.

A major limitation of this approach is the requirement to develop code for the DP algorithm around a specific powertrain model, as the operation and limits of the physical system define the constraints and penalties required for the optimization. To create a general-purpose analysis tool using DP would thus require significant abstraction of the problem. However, promising results have been shown with the use of metamodels in this application. A metamodel is a surrogate that is used in place of a more computationally expensive function, or to represent 'black box' functions where details may be obscured [43]. Metamodels can use analytical, polynomial, or

statistical methods to create a response function that replicates the behaviour of the original system.

Very recently, Murgovski *et al* [44] developed a method and tool using dynamic programming that allowed any Simulink powertrain model to be analysed for optimal power-split, even if the model details are hidden (i.e. a black box model). First, the Simulink model to be analysed is subjected to a series of tests to determine the characteristic responses of its powertrain components, in order to define metamodel parameters. For example, a varying ESS current would be simulated and the resulting voltage and SOC change would be observed, in order to calculate the equivalent circuit parameters of the ESS metamodel. Next, the metamodel is used as part of the DP analysis to obtain the desired optimal control policies. The authors noted that their method dramatically reduced the time required to perform the analysis, as the slowest component of the process (Simulink interface) was used sparingly, and the much-faster metamodel was used to populate a DP search space and determine results [45].

2.4 Background Summary

Hybrid powertrains offer the potential to reduce fuel consumption through a range of different technologies, in particular a high degree of vehicle electrification. As such, the Series-Parallel PHEV architecture selected for UVic's EcoCAR2 competition entry uses a high-capacity ESS and powerful electric machines to achieve performance increases as well as a fuel consumption reduction. The vehicle's multi-regime capability allows it to take advantage of the strengths of series or parallel operation in the operating conditions where they are most useful, especially when combined with advanced control techniques. A range of control techniques have been developed which can maximize the potential of a complex hybrid architecture, though they require careful tuning and setup. Global optimization can be used to determine optimal performance benchmarks and control trajectories for a vehicle architecture, by taking into account prior knowledge of driving conditions over a complete cycle.

Chapter 3

Hybrid Powertrain Modelling Fundamentals

This chapter discusses the concepts and strategies used in vehicle powertrain modelling, for performance analysis and/or control system development purposes. It also provides examples of specific powertrain component models used in EcoCAR simulation activities. This material provides a background for the modelling and simulation activities conducted for the EcoCAR competition, and for this thesis work.

3.1 Core Modelling Concepts

The use of simulation models in powertrain development is very widespread. With the increasing adoption of Model-Based Design (see Section 3.2), development of simulation models and performing simulations with them now forms the core of a range of related design activities. Different types and structures of powertrain component models can be implemented depending on the needs of the specific application.

A critical input to most powertrain modelling activities is the driving cycle or drive cycle. A drive cycle is a set of data, containing a specific profile of vehicle speed over time. A range of standardized driving cycles are available, mostly developed by the automotive regulatory bodies of different countries or regions. Use of drive cycles in automotive design allows the performance of different vehicle architectures to be assessed over different driving conditions, while use in testing allows for direct fuel economy comparisons between different vehicles. An example of a regulatory drive cycle is given in Fig. 3.1.

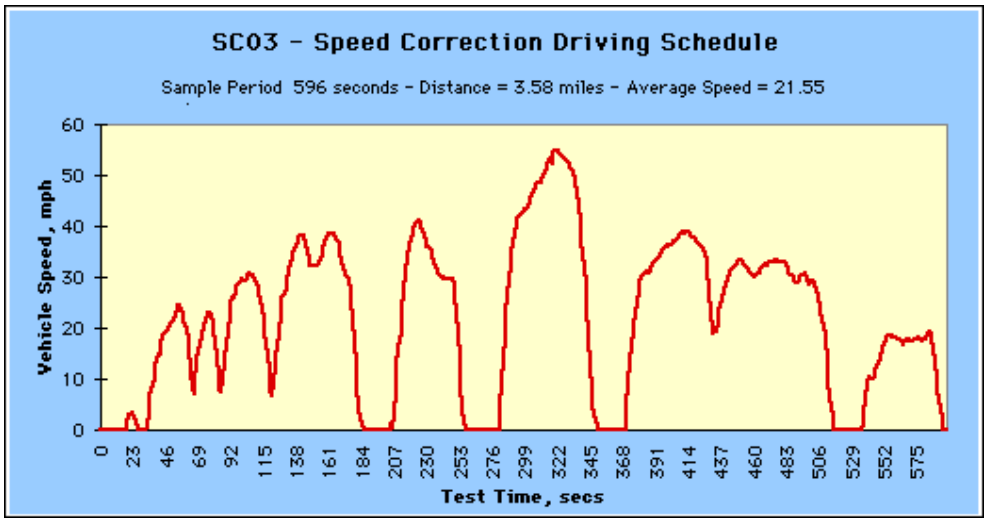

Figure 3.1: An example driving cycle, the US EPA SC03 (United States Environmental Protection Agency)

3.1.1 Model Types

Model types can generally be categorized as empirical or physics-based. This parameter indicates the degree of abstraction with which the physical principles governing the modelled system are represented.

Empirical models allow quick and low-cost simulation of existing systems. In this method, previously measured experimental data are employed to formulate a model that predicts system behaviour without directly considering the physical phenomena that define the system. Empirical models take the form of look-up tables, low-order circuit models, or similar abstractions of the physical system. A drawback of these models is that they are typically parametrized to data gathered under a specific set of operating conditions (e.g. temperature), and as such the model may not accurately reflect system performance under different conditions.

In physics-based models, the state variables of a system are modelled according to the physical laws that represent the underlying principles at work. This can require much greater computational effort, but results in a more accurate representation of the system. Physics-based models, by their nature, can also account for wider variations in operating conditions than empirical models. The Resistive Companion Form ('RCF') method is a widely used example. Borrowed from techniques developed for electronics simulation, RCF involves modelling system components as blocks, within which the equations describing the system dynamics are worked into a standard ma-

trix form that is conducive to numerical integration [50]. When all desired components are modelled, they can be integrated into one set of algebraic equations by applying connectivity constraints between neighbouring modular components, and solved to get system state variables [50].

3.1.2 Model Structures

Model structures, especially in the context of vehicle powertrain simulation, are categorized as forward- or backward-looking. This quality indicates a conceptual 'direction' for data flowing from input(s) to output(s) through the modelled system. The choice of model structure is highly application specific, as forward- and backward-looking models are suited to different tasks.

Forward-looking models are structured to replicate the causal nature of real-world events. In a forward-looking model of an automotive powertrain, a vehicle is desired to follow a driving cycle, and a simulated driver will send accelerator or brake signals to the powertrain in an effort to follow the desired speed as closely as possible. The powertrain produces torque, which acts on the driveline to cause a resulting speed increase or decrease. The torque and speed are used to determine component power inputs and outputs, efficiencies, and resulting energy and/or fuel consumption. Forward-looking models are commonly used for control systems development; controllers typically provide the link between the driver's operating requests and the powertrain components, and the forward-looking structure provides a model setup most conducive to developing that link.

In contrast, backward-looking models assume a system perfectly follows a desired path, and derive the internal states needed to achieve that path. Continuing the automotive example, the vehicle is assumed to perfectly follow the drive cycle, and the simulation derives how the powertrain must operate to achieve that. The driveline speed is set by the drive cycle, and the required component torques and powers, and resulting energy and/or fuel consumption, can then be determined. Since output torques and speeds in a backwards-looking model are constrained by the driving cycle, this method cannot be used to develop realistic control systems, and this lack of simulated control systems also means that some transient operating events, such as ICE starts/stops, cannot be represented. However, the structure is well-suited for defining overall operating trends and performing high-level analyses of a powertrain's operation under different conditions.

3.1.3 Model Fidelity Levels

Model fidelity refers to the level of detail captured by the simulation model. A model can be implemented to varying degrees of fidelity, depending on the particular application. Fidelity is strongly related to model type, as physics-based models will generally produce higher-fidelity results than empirical models. It should be noted that in modelling, 'fidelity' and 'resolution' are not interchangeable terms. Resolution in this case refers to the quantization step of simulated variables, or the fixed time step used in a discrete-time simulation.

A key aspect of modelling is the selection of the appropriate level of model fidelity, depending on the model's purpose and the design task being performed. The operation of a hybrid vehicle powertrain involves a wide range of mechanical, electrical, and chemical processes. Each powertrain component can be broken down to a collection of different subsystems and processes, which can themselves be modelled with different degrees of fidelity. For example, depending on application requirements, a battery may be modelled as a constant voltage source, a higher-order RLC circuit, a set of lookup tables derived from captured test data, or a complex multi-physics model simulating electro-chemical reactions. Overall, powertrain modelling requires the selection or development of component models of a fairly uniform fidelity to ensure that, when combined together, they can effectively simulate the powertrain system.

Selection of model fidelity level is influenced by the methods and tools used. When developing a model, different sets of modelling methods can be used that are based around higher- or lower-level conceptualizations of a system. If using commercial modelling tools, different products are often available offering differing fidelity levels. For example, SimDriveline's native vehicle dynamics model only implements longitudinal vehicle dynamics, with stiff driveshaft models. This type of model is very useful for a higher-level examination of powertrain performance and fuel economy, and for supervisory controls development, but would be useless in the development of traction control systems or suspension geometry. In contrast, dSPACE offers a model set that can simulate full 3D vehicle dynamics with compliant driveshafts and variable tire friction for safety system development work, which provides a much more complete picture of vehicle behaviour at the expense of vastly increased complexity in implementation and tuning.

3.2 Model-Based Design in Powertrain Development

Modern vehicles, whether conventional, electric, or hybrid, are built from complex mechatronic systems which can involve a lengthy and costly development process. To deal with the costs and challenges arising from development of these complex systems, the automotive industry has adopted a development methodology termed 'Model-Based Design' ('MBD') . The concept of MBD involves the use of a progression of mathematical models to simulate the operation of a given physical system instead of implementing a series of costly prototypes. The easily-modifiable simulation models allow development to be performed in a safe, timely, and inexpensive manner. Under MBD, by the time actual prototypes are built, the majority of design challenges have already been encountered and overcome.

Under MBD, plant models (representing the system to be controlled) are created and maintained to facilitate the implementation and integration of corresponding control systems. These models can be of varying levels of complexity. Consider the battery example given above, where models ranged from elementary circuits to complex electro-chemistry representations. In the context of powertrain development, the MBD process is used to develop plant models of engines, transmissions, electric motors, and often the entire vehicle, in order to develop and tune control systems for those components and/or a vehicle-level supervisory controller. The MBD process is used extensively in the EcoCAR competition, to facilitate early modelling and simulation efforts supporting powertrain development, and continued over the course of the competition to aid in supervisory control implementation.

3.2.1 MBD Process

The MBD methodology typically follows a specific process, outlined in Fig. 3.2.

The preliminary stages involve system requirements definition and high-level design activities. At the MIL stage, a simple controller mock-up is used to drive a plant model. The intent is to design and validate the algorithms and model systems that make up the core functionality of the plant - complexity and detail is added gradually, as needed. Neither real-time operation nor a fully realistic plant model is required at this stage, since a high degree of accuracy or refinement is not needed for basic performance validation. Controller development can proceed separately, or

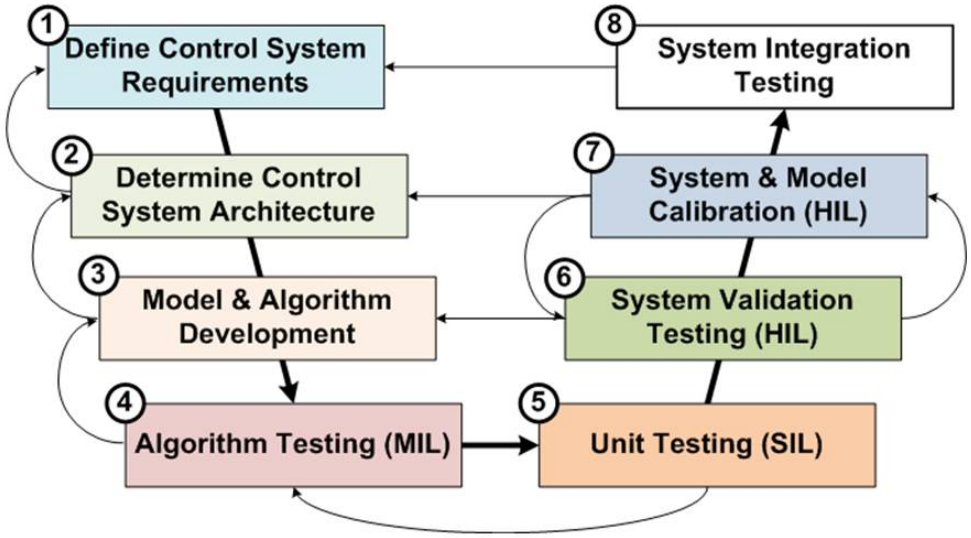

Figure 3.2: The Model-Based Design pathway, or 'V-Diagram' (Image courtesy of UVic EcoCAR2 Team)

incrementally with plant model development.

At the SIL stage, the production controller software is incorporated into the simulation of the plant model. This step allows the validation of most of the controller's functionality through rigorous testing, as well as identification of any major software components for which no corresponding model exists. Models at the SIL stage do not need to run in real time.

At the HIL stage, the controller code is downloaded to the physical hardware and connected to the simulated plant model. In real time, the controller hardware sends commands via discrete electrical connections or digital communications buses, where they are received and converted to system variables in the plant. The plant model then calculates the values of the variables that represent plant outputs, and sends them to the controller [49]. The intent at this stage is to have the model function such that the controller hardware performs as though it were integrated into the real physical plant.

Once the control system has been realized at the HIL stage, development can proceed to system integration, where the controller hardware can be connected to the real plant and finalized. If the plant model has been developed correctly, and HIL integration has been successful, only a small amount of work should remain to complete the integration, reflecting the various simplifying assumptions and abstractions made during development of the plant model.

It should be noted that the process is not rigid; at any stage of the MBD process, the current state can regress to a previous state, or even to the start of the process. This is to account for the effects of unforeseen design changes that occur during development. For example, if a system fails later-stage testing, a different architecture may be required, or the system design requirements may need revising. The process is also not restricted to a single modelling environment; an increase in fidelity requirement may necessitate a change to a more capable modelling tool, or a design change may allow regression to a simpler tool to free up computational resources.

3.3 Powertrain Simulation Environments and Tools

A wide variety of simulation environments exist in industry and academia to support control systems implementation, research and development, and other projects. Many environments function as a turn-key solution for powertrain simulation, while others require or allow in-house model development and customization. The two modelling environments used in EcoCAR are Autonomie and MATLAB Simulink. These provide a wide range of tools and options for powertrain analysis to controls development activities.

Autonomie, like its predecessor PSAT, is an open plug-and-play environment based on MATLAB that is intended to support the rapid evaluation of new powertrain architectures or propulsion technologies for improving fuel economy [51]. Autonomie models utilize predefined architecture layouts that allow simulation of a wide range of vehicle types. Individual component models are empirical and highly customizable; component power capabilities can be scaled up or down to perform parametric studies, and different components can be simulated by directly loading different sets of operating data.

MATLAB Simulink is widely utilized in the automotive industry for performing simulations of vehicle powertrain systems [52]. Simulink code consists of functional blocks which are combined using a graphical interface to form algorithms, making

it more accessible to developers lacking a strong background in conventional programming techniques. A variety of Simulink block sets are available that provide commonly-used modelling functionality. The SimDriveLine block set, for example, contains representations of a selection of automotive powertrain components including simplified ICE's, transmissions, and tires.

The dSPACE Automotive Simulation Model ('ASM') series is a set of high-fidelity automotive component models, built in Simulink, that can be combined to simulate anything from extended automotive subsystems to a complete virtual vehicle. Many ASM blocks can also be parametrized to emulate the performance of a specific vehicle using captured test data. dSPACE ASM blocks are geared towards modelling at the MIL through HIL levels, and the company also produces specialized hardware for HIL testing.

Simulink and ASM models use a torque-forward, forward-looking modelling structure. A driver or supervisory controller model commands the powertrain components to produce torque, which acts on a gearbox or shaft, cascading forward in this manner through the system until it reaches the driving wheels. The applied torque causes a change in wheel speed, and the speed is then cascaded backward through the system until it reaches each powertrain components. Once all component torques and speeds are known, efficiencies or power losses are determined and electric power use calculated. The driver model modifies its output to regulate how closely the vehicle follows a drive cycle. The EcoCAR team's primary powertrain model was developed using Mathworks' Simulink software, utilizing a combination of custom Simulink, SimDriveLine, and dSPACE blocks.

At the time of writing, there are no readily-available off-the-shelf tools for performing design optimization or performance benchmarking of complex hybrid powertrains. This is primarily due to the need to include specific powertrain and component data in the optimization problem formulation, making it impractical to implement a truly general-purpose tool. A Global Optimization Toolbox is available in MATLAB that implements genetic algorithm, simulated annealing, and particle-swarm optimization methods that can be utilized in design optimization activities. The user must formulate their problem to fit a specific format required by the toolbox.

3.4 Powertrain Subsystem Models

This section outlines the formulation of key automotive component models, and describes the corresponding model used by the EcoCAR team. The majority of the EcoCAR team's component models are of the empirical type.

3.4.1 Vehicle Dynamics

A real automobile moves through 3D space on a sprung suspension, with a load-varying mass distribution and level of traction on each wheel. Development of a full 3D dynamics model of a vehicle is a massive and complex undertaking, with factors such as body roll and flex, steering and suspension geometries, and mass distribution (among many others) requiring consideration. Dynamics are typically separated into the longitudinal, or forward-aft component, and lateral, or side-to-side component. Longitudinal dynamics includes the various resistances to forward motion, such as drag, rolling resistance, and inertia. Lateral dynamics include the effects governing cornering and stability, such as body roll, understeer and oversteer, and lateral weight transfer.

For most supervisory control systems development or benchmark analysis, the simulation of lateral dynamics is not required, and even the longitudinal dynamics can be somewhat simplified without compromising the quality of the results. For example, forward-aft mass transfer during acceleration and braking is not considered, and the friction between tires and road is assumed to be constant. This is because we are assuming stable operation for the purpose of benchmark assessment, and that the magnitude of the base road load is much greater than the magnitude of load perturbations during longitudinal driving.

A typical longitudinal vehicle dynamics model implements the equation

$$P_{req}(t) = (m\frac{dv}{dt} + F_{roadload}(t)) * v(t) \tag{3.1}$$

where $P_{req}(t)$ is the propulsive power required to move the vehicle at time t, m is the vehicle mass, v(t) is the vehicle's velocity at time t, and $F_{roadload}(t)$ is the sum of the various forces such as air resistance, rolling resistance, and grade resistance, that act against the propulsive force. An example of these forces is shown in Fig. 3.3.

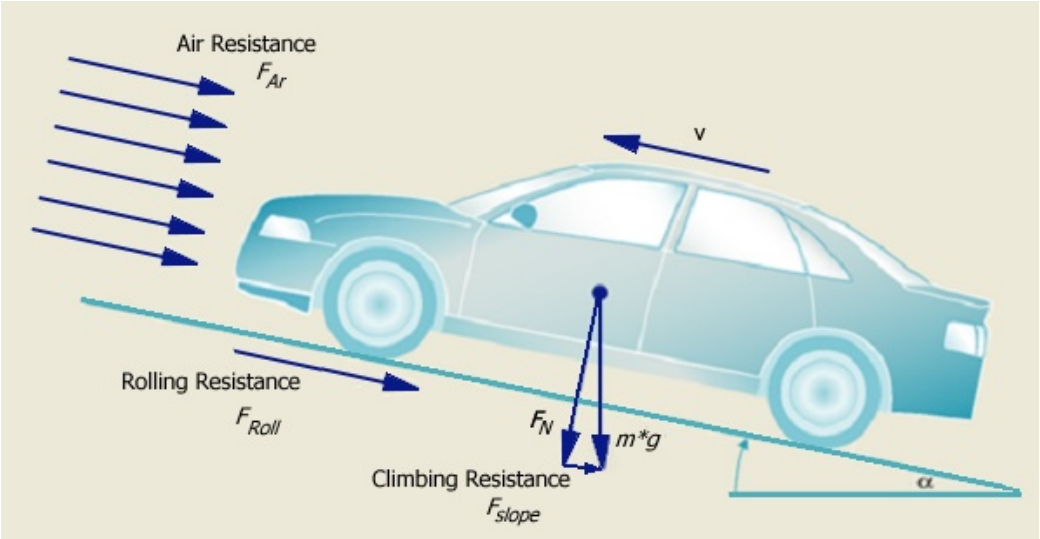

Figure 3.3: Simplified Longitudinal Vehicle Dynamics

Alternatively, $F_{roadload}(t)$ can be approximated by the polynomial

$$F_{roadload}(t) = A + B * v(t) + C * v(t)^2 \tag{3.2}$$

The coefficients A, B, and C form a polynomial that describes the forces on the vehicle resisting its longitudinal motion. Coefficients A and B represent the static and dynamic friction forces due to bearing losses and tire/road rolling resistance, and C represents the drag and aerodynamic resistances. These parameters can been derived from the road load and drag of a real vehicle, measured in a standard procedure known as a 'coastdown test'. In this test, a vehicle is accelerated on a test track to a high speed, then allowed to coast to a stop while the velocity is measured and recorded. With a known mass and the coastdown velocity data, numerical analysis can be performed to derive Eq. 3.2 via curve-fitting. This method is ideal for modelling road load resistances, as the myriad of technical details required to calculate resistive forces from first principles (ex. bearing friction values, or frontal area of a vehicle body) are difficult to gather.

The vehicle dynamics model used in implements the coastdown-derived resistance equation. The coefficient values used were provided by General Motors, and are given in Table 3.1.

Coeff	Value
A	21.77
B	0.3673
C	0.01847

Table 3.1: 2013 Malibu Road Load Coefficients

3.4.2 Engine

An ICE is a complex machine, even without considering fuel injection and electronic control. Controlled combustion reactions convert the chemical energy of the intake air/fuel mixture into thermal energy, which is harnessed by a mechanical linkage to provide a power output. Fundamentally, an ICE can be modelled as a torque actuator, providing an output torque while consuming a proportional amount of fuel. A simplified ICE model contains this functionality, and may even omit the fuel use; for example, the Simulink 'Generic Engine' model only considers torque, speed, and rotational inertia.

A more complex ICE model uses first principles to perform detailed analysis of combustion, heat transfer, and flow rates at the engine valves as a function of crankshaft angle. Engine manufacturers may use such advanced, and likely proprietary, simulation models or modelling programs as part of their design efforts. Companies such as dSPACE also provide complex simulation models of a generic nature that can be parameterized to represent a specific engine using measured data. For example, the dSPACE 'ASM Engine Gasoline' blockset models crank angle-based torque generation, dynamic manifold pressures, engine component temperatures, and fuel injection dynamics [53].

Two critical sets of data are required to model an ICE: the torque/speed or power/speed curve, and a brake-specific fuel consumption ('BSFC') map. The power and torque outputs of an ICE vary non-linearly with speed, following a curve similar to that shown in the example in Fig. 3.4. This type of curve is used to define the maximum torque and mechanical power the ICE can produce. The BSFC map provides the ICE fuel consumption as a function of operating torque and speed.

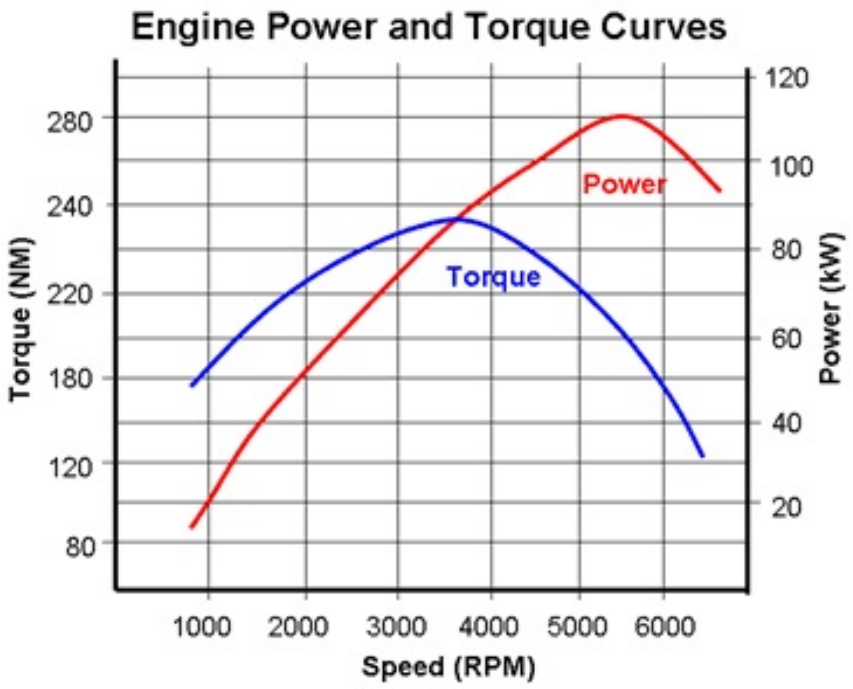

Figure 3.4: Example ICE Power and Torque Curve (www.http://overvoltage.org/wp-content/uploads)

The EcoCAR ICE model is quasi-static, using a BSFC map created from measured data provided by General Motors. The model acts as a torque source, responding to commanded torques up to a maximum limit, which is a function of operating speed. The fuel consumption map is used to determine the instantaneous rate of fuel use at the current operating point. The BSFC map of the LE9 engine is given in Fig.3.5.

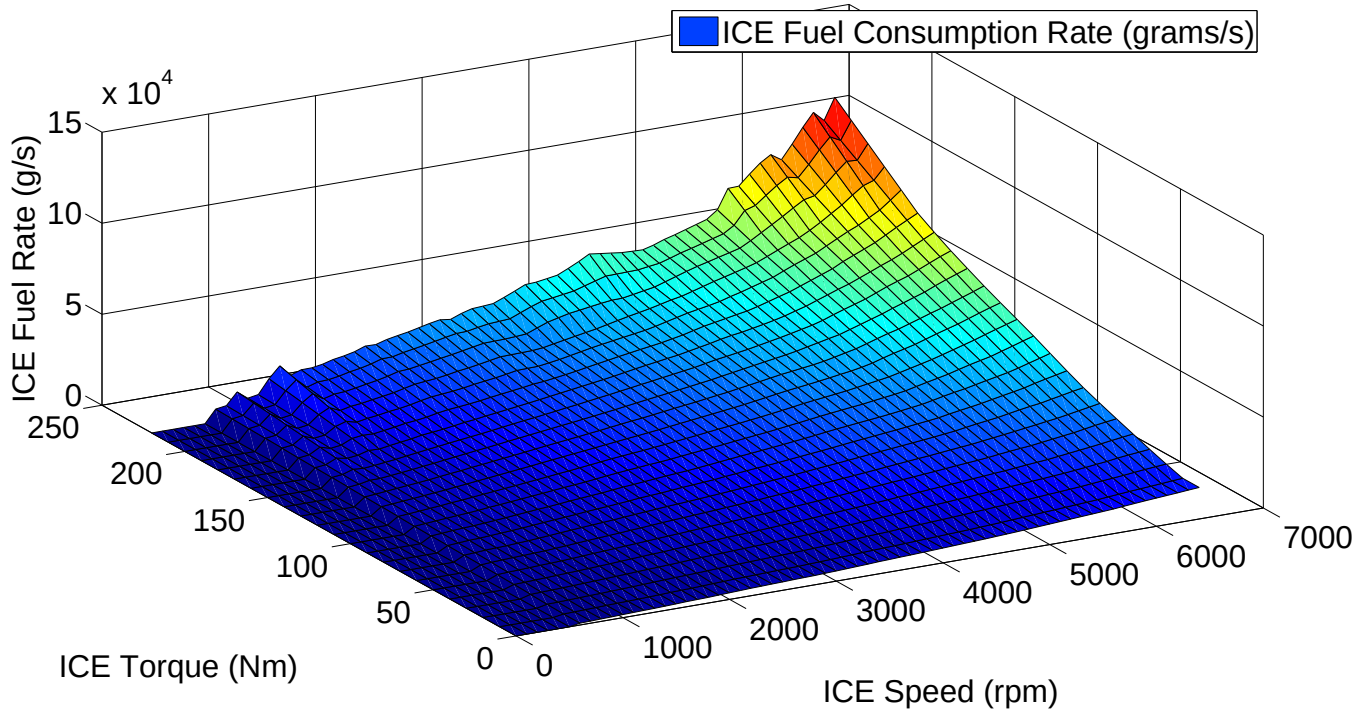

Figure 3.5: LE9 BSFC Map

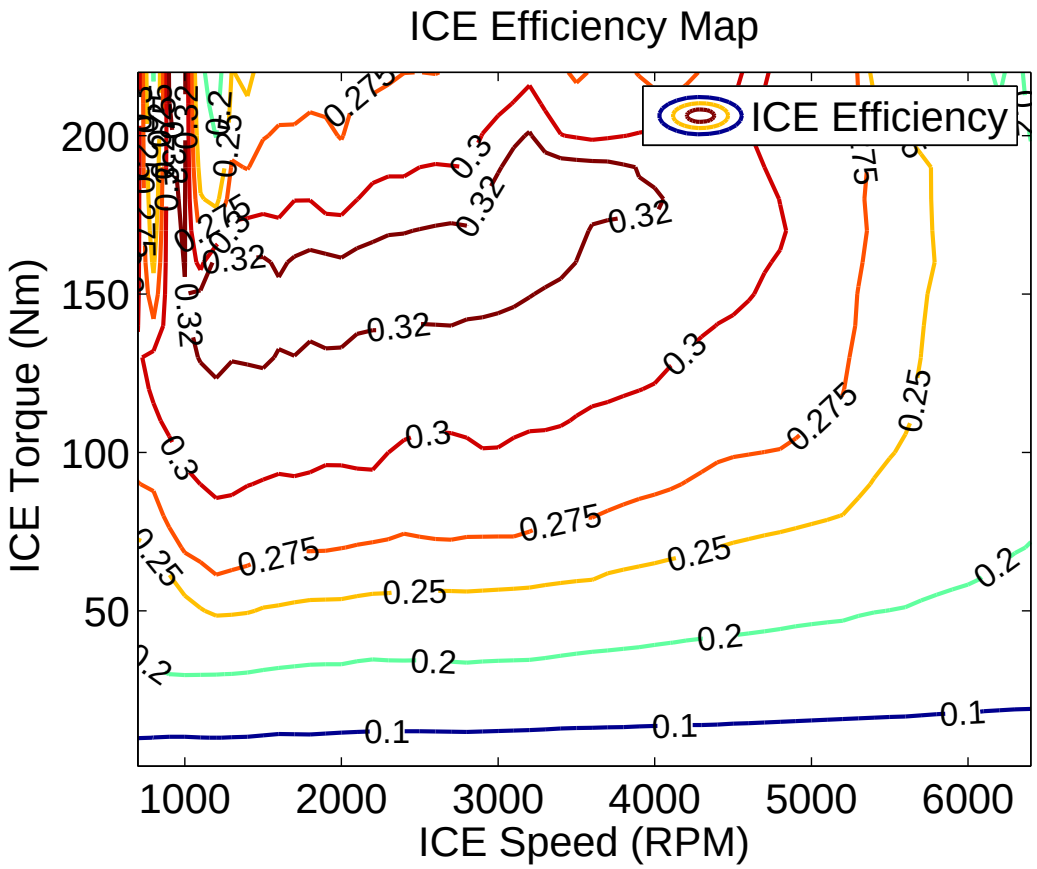

Figure 3.6: LE9 Efficiency Map

It should be noted that due to the static nature of the ICE model, the engine dynamics (transitions between operating points) are not considered separately, and any discontinuities resulting from the BSFC are managed by interpolation.

3.4.3 ESS

The trade-off in modelling battery systems is between the inclusion and depth of physico-chemical effects, which generally improve predictability, and the amount of computation effort required. Thus, battery modelling techniques can be roughly grouped into four major categories: empirical models, electrochemical engineering models, multi-physics models, and molecular/atomistic models [54]. The latter three methods are used when electrochemical details are especially critical, such as during development of new electrode materials or battery chemistries. For the majority of powertrain modelling, the primary concern is the macro-scale electrical behaviour of the battery, thus empirical models offer sufficient fidelity for this application. An

example multi-order equivalent circuit model of a battery is shown in Fig. 3.7.

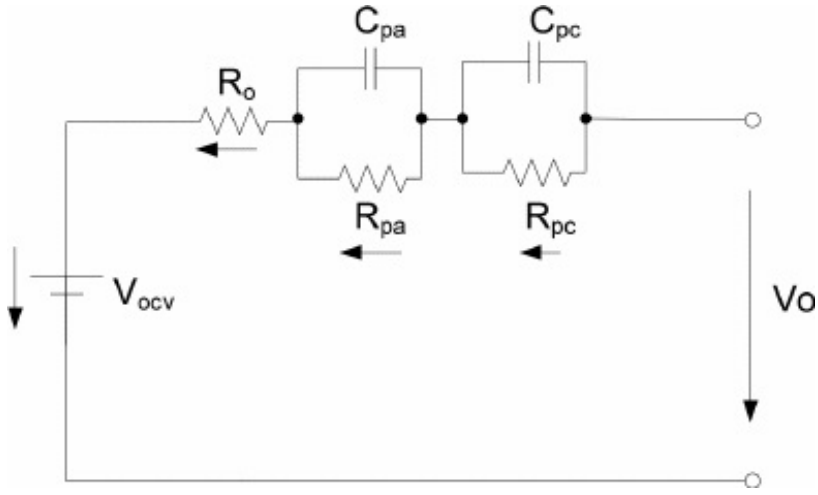

Figure 3.7: ESS Equivalent Circuit Example

The ESS model used for EcoCAR simulation work is empirical, based on performance and internal data supplied by A123 for the battery system described in Section 2.2. This data is confidential, and is thus not given here. To illustrate ESS behaviour, vehicle simulation results from the EcoCAR Team's Simulink model are given for operation in EV mode, with the critical ESS parameters of pack voltage, pack current, and state of charge shown in Fig. 3.8.

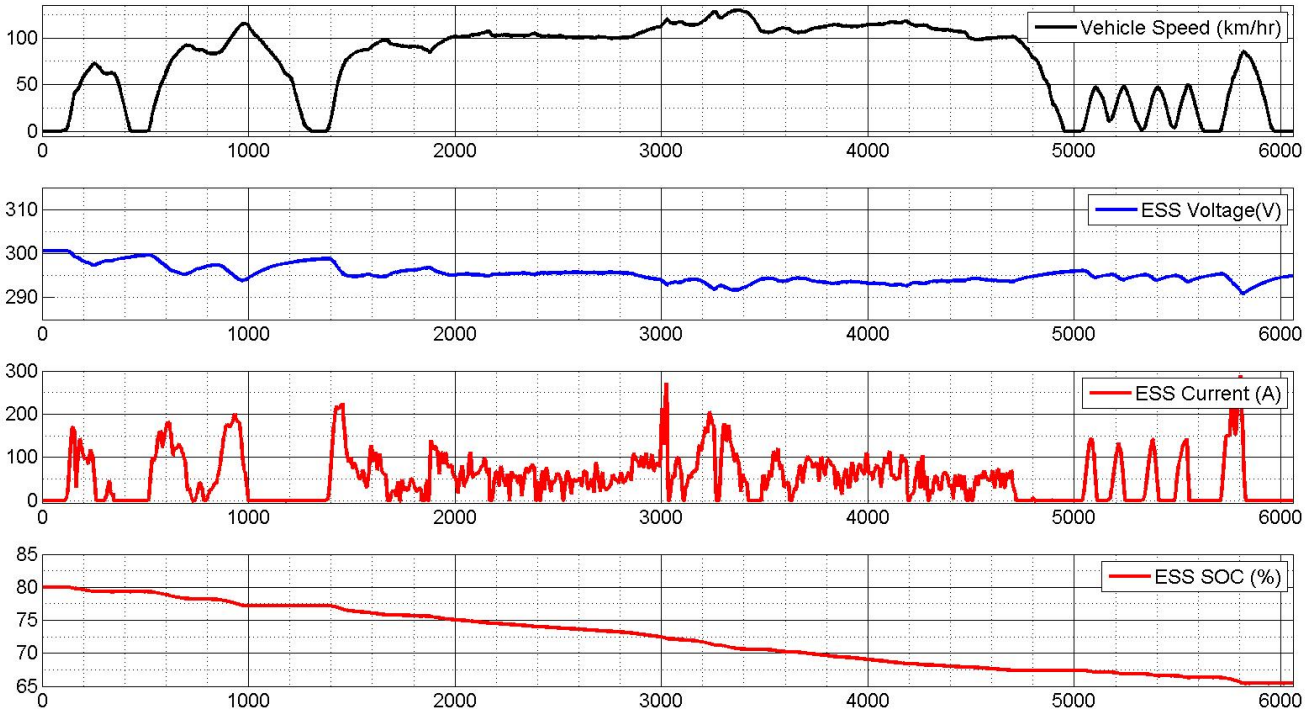

Figure 3.8: ESS Discharge Example

3.4.4 Electric Machines

Modelling techniques for electric motors are divided along empirical and multi-physics lines. Physics-based models of motors are most useful in design optimization applications, or in the development of fault detection and motor control strategies. In these cases, a detailed picture of magnetic field distribution, winding currents and voltages, and rotor dynamics is critical. In most powertrain modelling, an empirical model capturing input (electrical) and output (mechanical) power and efficiency is sufficient. The main information of interest is the interaction of the motor with the rest of the vehicle, such as the power drawn from the ESS and the torque applied to the wheels.

In the EcoCAR team model set, both the BAS and the RTM are modelled empirically as torque actuators, with the output torque modified to be compliant with the EM's physical torque and speed limits. The RTM acts through a reduction gearbox on the rear axle, and the BAS is coupled to the ICE. The models determine the elec-

trical input power required to achieve a requested output using component efficiency data, as provided by the motor's manufacturers. The Magna E-Drive's efficiency map, combined with the maximum torque-speed curve, is given in Fig. 3.9. The BAS efficiency map is similar.

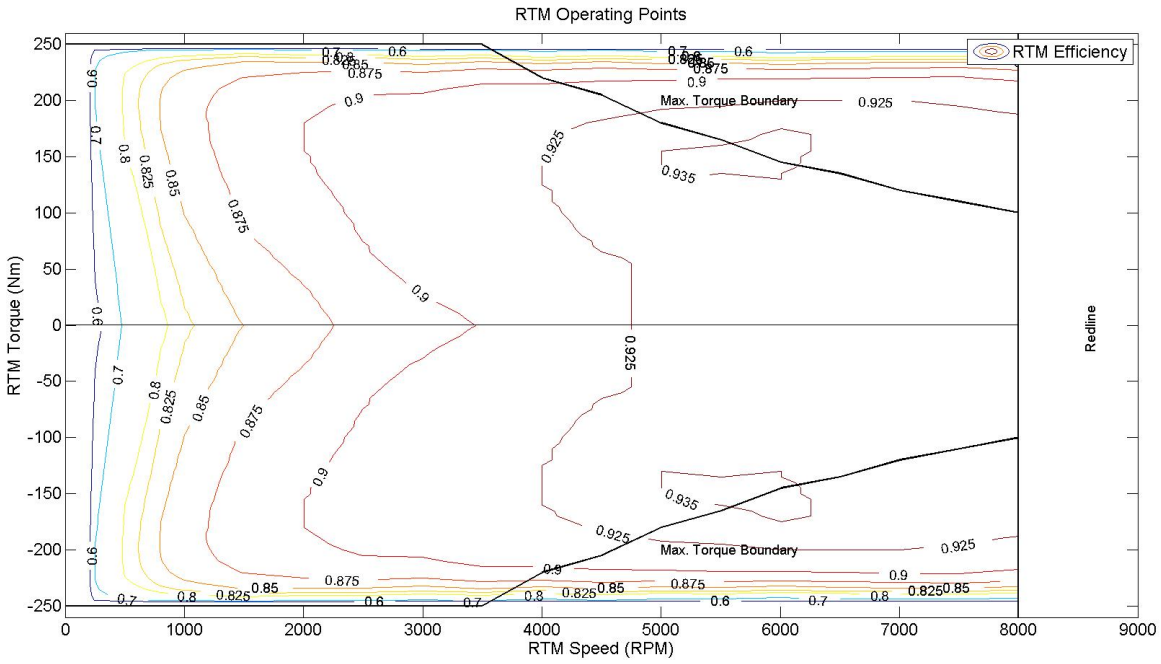

Figure 3.9: Magna E-Drive Efficiency map with torque limits

It should be noted that the torque response of the motor models is assumed to be instantaneous, since no transient operating data was available at the time of model development. This leads to occasional discontinuities in the data, which are managed by interpolation. Additionally, the efficiency data provided by the motor manufacturers was captured at only one nominal input voltage, and thus the small variations in motor power loss due to input voltage transients will not be captured by the model.

3.4.5 Driver

In any forward-looking simulation, the vehicle model must be made to follow the required driving cycle as closely as possible. In effect, the driver's behaviour must also be simulated. If only considering the longitudinal operation of the vehicle, the

driver requests additional positive torque from the powertrain until the vehicle speed is equal to the desired value, and requests negative torque if the vehicle speed must be decreased.

An effective automated driver model combines open-loop feed-forward with a closed-loop PI controller, tuned to follow a setpoint provided by a driving cycle. The open-loop component roughly calculates the output torque requirement at each simulation step, and the closed-loop component acts to minimize the speed error and compensate for disturbances. The resulting output is the driver torque request, which represents the accelerator pedal position provided by a real driver. A diagram of this process is given in Fig. 3.10.

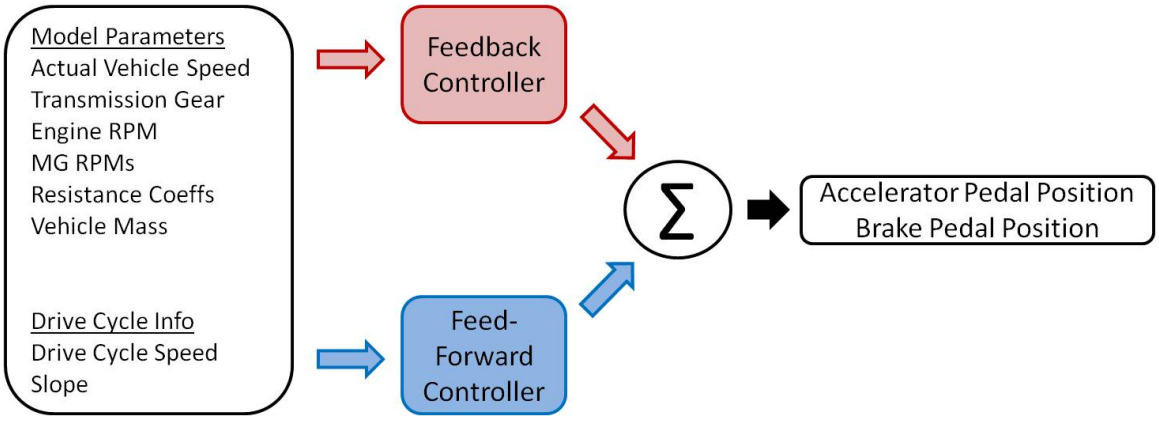

Figure 3.10: Driver Model Flowchart

It should be noted that a backward-looking model, such as that used by the DP algorithm in this work, does not require a simulated driver, since the vehicle is implicitly assumed to follow the driving cycle correctly and exactly.

3.5 Modelling Development Summary

This chapter provides a basic overview of hybrid vehicle powertrain modelling techniques, and describes the chosen component model fidelities and methods used in this work. In general, the powertrain components are represented by empirical models generated from measured data, while a simplified physics-based model is used to represent longitudinal vehicle dynamics.

Chapter 4

Benchmark Analysis Methodology

As mentioned previously, one of the team's design goals for the EC2 architecture was to allow powertrain flexibility in meeting the driver's torque requests, and to exploit that flexibility to achieve lower fuel consumption. A range of control strategies can be implemented to manage powertrain operation, and each strategy will impose its own 'flavour' of operation in terms of regime selection and torque distribution. This would result in a wide range of fuel consumption values and operating profiles.

It is apparent that in order to objectively evaluate the performance of the selected EC2 architecture, it is necessary to derive operational benchmarks for the architecture that are independent of any specific control methodology. The key information to be obtained is

1) The optimal control policies (component operating profiles) for the different operating regimes over city and highway driving.

2) The resulting optimal fuel economy from application of those policies.

3) The optimal power split between electrical and mechanical power.

4) The ideal operating regime to employ under specific driving conditions.

Gathering this data will provide an insightful look at the architecture's performance potential. The information can also be used to guide the development and implementation of applicable real-time control strategies. Hence, benchmark analysis is the focus of this work.

In order to perform this analysis, it is necessary to consider all events over the driving cycle, as operational decisions made at or near the beginning of the cycle will

influence the ending conditions. For example, heavy ESS discharge at the start will require a corresponding heavy charge towards the cycle finish. When simulating in real-time, it is not possible to take these future events into consideration, due to the forward-looking nature of real-time control systems. It is also not possible to search through all feasible paths in real-time, due to the computational burden of the search. As mentioned in Section 2.3, global optimization methods are ideal for performing this type of analysis.

The identification of optimal control policies is fundamentally a constrained global optimization problem. Since the policies must be globally optimal over complete drive cycles, drive cycle time is an additional variable in the optimization, making the optimization formulable as a path search problem. Thus, of all the methods available, dynamic programming was selected for the benchmarking analysis. Other global optimization methods such as genetic algorithms or simulated annealing are well-suited for design optimization, but less so for the automotive benchmark problem formulated as a path search.

Dynamic programming offers dramatically reduced computation time compared to some other graph search methods. As this is a global optimization program that is being run off-line, the computation time is not as critical a constraint on the problem as it would be if operating in real-time. Despite this, it is desirable to keep the computation time reasonable and manageable to improve usability of the program. DP was selected over other comparable graph and tree search methods due to the relative complexity of the cost calculation and the overlapping nature of the subproblems. Analysing a powertrain operating over a drive cycle is more complex than a typical path search problem, with the transition costs evolving depending on the particular path selected.

DP also implicitly excludes infeasible solution paths. This characteristic is particularly important, as the use of a backward-looking model, with its implicit assumption of correctly following a driving cycle, can yield many infeasible solutions where particular components or operating states cannot meet the imposed load.

4.1 Fundamentals of Dynamic Programming

This section introduces the fundamental mathematical concepts behind Dynamic Programming. DP was developed by mathematician Richard Bellman in the early 1950's, as part of work on stochastic multi-stage decision-making algorithms for the United

States Air Force [46]. Over the years it has become a widely-used approach in scheduling, inventory management, and process planning, as well as problems requiring multi-stage decision-making [47].

4.1.1 The Basic Problem

Given a discrete-time deterministic system, the states evolve over time as described by the transition function

$$x_{k+1} = f(x_k, u_k), \qquad \{k = 0, ..., N\} \tag{4.1}$$

where x_k, the state variable at time k, exists in a space S_k, and u_k, the control input at time k that modifies x_k, exists in a space C_k.

The set of control inputs, termed a *policy*, consists of a sequence of functions

$$\pi = \{u_0, u_1..., u_{N-1}\} \tag{4.2}$$

where each u_k is constrained to take values in a subset of C_k, depending on the current state x_k. The specific constraints applied are part of the formulation of a particular DP problem, and serve to eliminate infeasible control inputs. The set C_k is called an admissible policy.

Each transition of x_k between different values incurs a cost. The cost represents the effort of moving from one state to another, and serves to differentiate the paths. In a literal path-search problem, the cost represents the distance required for each path step. The cost can also be used to enforce constraints, by highlighting infeasible state transitions and ensuring they will not be selected during the path search. The cost function describing the state transition and implementing constraints is given by $g_k(x_k, u_k)$.

The cost of a given path is additive over time, since it accumulates over each stage of the problem. The total cost of a path can be expressed as

$$g_N(x_N) + \sum_{k=0}^{N-1} g_k(x_k, u_k) \tag{4.3}$$

Thus, the cost of following a given policy π starting from state x_0 is

$$J_\pi(x_0) = g_N(x_N) + \sum_{k=0}^{N-1} g_k(x_k, \pi_k) \tag{4.4}$$

An optimal policy can be found, denoted π^*, that minimizes this cost, such that

$$J_{\pi^*}(x_0) = min \ \{J_\pi(x_0)\} \tag{4.5}$$

4.1.2 The Principle of Optimality

The cornerstone of DP is Bellman's Principle of Optimality, which states that "An optimal policy has the property that whatever the initial state and initial decision are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision" [48].

This property suggests that if an optimal policy

$$\pi^* = \{u_0^*, u_1^*..., u_{N-1}\} \tag{4.6}$$

passes through the state x_i at time $k = i$, and we desired to find the optimal policy to get from x_i to x_N, we would do so by minimizing the truncated cost function

$$J_\pi(x_i) = g_N(x_N) + \sum_{k=i}^{N-1} g_k(x_k, u_k) \tag{4.7}$$

and we would find that the optimal policy is simply the truncated policy

$$\pi_i^* = \{u_i^*, u_{i+1}^*..., u_{N-1}\} \tag{4.8}$$

Practically, this means that the overall optimal control policy can be derived by sequentially determining the set of optimal policies for a series of smaller sub-problems. Put to a road trip analogy, if the fastest driving route from Vancouver to Toronto passes through Calgary, the principle of optimality implies that the Calgary-to-Toronto section of the overall drive is also the fastest route from Calgary to Toronto.

4.1.3 DP by Backwards Induction

There are several different approaches to the practical implementation of a DP algorithm, the selection of which depends on the particular characteristics of the problem to be solved. In optimal control problems where *a priori* knowledge of the expected

operating conditions and desired path is available, the optimal policy can be found with DP using a backward induction approach, by looking at the problem from the final (terminal) point and calculating the cost for each successive stage, working backwards in time towards the initial point. In the context of automotive design, standardized drive cycles provide the desired path. It is important to note that due to the random nature of on-road driving, it is not possible to know true on-road driving conditions *a priori*.

First, the final stage ($k = N-1$) is considered, and an optimal policy is determined for this step. This is referred to as the 'tail sub-problem'. Next, the optimal policy for the tail sub-problem involving the final two stages ($k = N-2 : N-1$) is determined, and the process is continued until the policy for the full problem (k=N-1:1) has been identified. More application-specific information about the choice of backwards induction is given in Section 4.3.

This type of problem can be efficiently solved using a recursive algorithm, as follows. Starting at the final stage ($k = N-1$), the minimum cost is found to be

$$J^*_{N-1}(x_{N-1}) = min \ \ \{g_{N-1}(x_{N-1}, u_{N-1})\} \tag{4.9}$$

The minimum cost of the next stage ($k = N-2$) is therefore

$$J^*_{N-2}(x_{N-2}) = min \ \ \{g_{N-2}(x_{N-2}, u_{N-2}) + J^*_{N-1}\} \tag{4.10}$$

Working backwards from here towards the initial stage $k = 0$, the total cost can therefore be found by the recursive equation

$$J^*_k(x_k) = min \ \ \{g_k(x_k, u_k) + J^*_{k+1}\}, \qquad k = \{0, ..., N-1\} \tag{4.11}$$

where the cost $J^*_0(x_0)$ is the optimal cost of the overall control policy π^*.

4.2 Statement of the Optimization Problem

The benchmark analysis problem is, at its core, a global optimization problem. The basic objective function can be expressed as

$$J = min \sum_{t=0}^{T} \dot{m}_{fuel}(t) \tag{4.12}$$

where t is the discrete time step of the drive cycle, $\dot{m}_{fuel}$ is the fuel consumption rate of the engine (quasi-static over each timestep t), and 0 and T are the beginning and end points, respectively. It must be emphasized that this function is not describing the minimization of instantaneous fuel consumption at each step. Rather, the minimum referred to here is the sum total of fuel consumed over the full drive cycle. This is to allow for opportunity charging and discharging - the ability of the hybrid powertrain to either produce more engine power to adjust the ICE operating point for improved efficiency, or discharge the ESS via the electric motor(s) to provide propulsive power.

As the benchmark results must be directly comparable to outputs obtained via other control methods, or other powertrain architectures, CS operation must be maintained. Therefore, the change in ESS SOC between the beginning and end of the drive cycle must be kept to as near zero as possible, regardless of the level of SOC variation during the drive. Note that this restriction is required for comparison, and would not normally be present in real-world operation. This leads to the constraint

$$\sum_{t=0}^{T} P_{batt}(t) \approx 0 \tag{4.13}$$

Although the state of charge of the ESS must also be controlled, as the ESS can be damaged by over-charging or over-discharging, this is not an explicit constraint in the optimization problem, as it is addressed in the DP algorithm formulation, as described in Section 4.3.

A key constraint is imposed by the driver power request P_{req}; the powertrain cannot produce any more power at the wheels than the driver commands. To reduce the complexity of the problem, the net power to or from the ESS (P_{elec}) is chosen as the independent variable. This is effectively the amount of electric propulsion or regenerative braking used. The mechanical power (either from the ICE/BAS or the mechanical brakes, depending on whether the vehicle is accelerating or decelerating) forms the dependent variable.

In order to represent the different operating regimes, some variations of this constraint occur. The parallel regime operates in a 'parallel through the road' configuration, meaning the ICE is powering the front wheels, and the RTM is driving the rear wheels. The power split in this case is mechanical in nature, with the sum of the ICE and RTM output powers being equal to P_{req}. This yields the constraints

$$P_{ICE}(t) = P_{req}(t) - P_{RTM}(t), \quad P_{req}(t) > 0 \tag{4.14}$$

$$P_{Brake}(t) = P_{req}(t) - P_{RTM}(t), \quad P_{req}(t) < 0 \tag{4.15}$$

where

$$P_{RTM}(t) = \begin{cases} P_{elec}(t)/\eta_{RTM}, & P_{RTM}(t) > 0 \\ P_{elec}(t) * \eta_{RTM}, & P_{RTM}(t) < 0 \end{cases} \tag{4.16}$$

where η_{RTM} is the RTM efficiency. In the series regime, the RTM provides all the power at the wheels. The mechanical brakes are only used when the RTM cannot provide enough stopping power, due to high SOC or power limitations. The power split in the series regime is electrical in nature, occurring in the vehicle's electrical bus. Energy to or from the battery (P_{elec}), and energy generated via the BAS and ICE (P_{Gen}), must sum to be equal to the input power of the RTM (P_{RTM_in}). This gives the constraints

$$P_{RTM}(t) = P_{req}(t) \tag{4.17}$$

$$P_{Gen}(t) = P_{RTM_{in}}(t) - P_{elec}(t), \quad P_{req}(t) > 0 \tag{4.18}$$

$$P_{Brake}(t) = P_{RTM_{in}}(t) - P_{elec}(t), \quad P_{req}(t) < 0 \tag{4.19}$$

where

$$P_{RTM_{in}}(t) = \begin{cases} P_{RTM}(t)/\eta_{RTM}, & P_{RTM}(t) > 0 \\ P_{RTM}(t) * \eta_{RTM}, & P_{RTM}(t) < 0 \end{cases} \tag{4.20}$$

The trivial solution for optimal fuel economy is to never consume any fuel, by keeping the ICE off and discharging the ESS to power the vehicle. However, to ensure CS operation, the SOC needs to be kept from diverging too far from the setpoint towards the end of the simulation, such that the ESS can be recharged to its initial value within the remaining drive cycle time. While ensuring pure CS operation is not a core real-world requirement, it is required for the benchmark analysis in order to facilitate fuel consumption comparisons across different hybrid architectures. Strictly enforcing CS behaviour generally results in higher fuel consumption, as more charging is needed over a cycle. A penalty function is used to enforce this, and takes different forms depending on the operating regime and current drive cycle conditions. Referred to as the 'regeneration penalty', the function is given by

$$\gamma\alpha|P_{elec}(t)| + (1-\gamma)\beta|P_{brake}(t)|$$
$$\text{where } \gamma = \begin{cases} 1, P_{req}(t) = 0 \\ 0, P_{req}(t) < 0 \end{cases} \tag{4.21}$$

where α and β are tunable weighting parameters that are adjusted to ensure CS operation in all regimes and all driving cycles. The regeneration penalty adds additional cost to use of the mechanical brakes, and to charging through the BAS when stationary. Some tuning of these functions is required, to balance the additional cost imposed by the penalty with the base transition costs.

Additional constraints are imposed on the optimization due to the physical limitations of the vehicle powertrain. There are finite limits to the amount of instantaneous power the ESS, electric machines, and ICE can supply at a given vehicle speed, stemming from the limits on both torque output and rotational speed of the components. These are hard limits, meaning the simulation result is infeasible if it must exceed them. To deal with these cases, the algorithm checks the component operating points and power levels as they are calculated, and discards any solutions that are tagged as infeasible.

All of the preceding requirements and constraints lead to the formalized optimization problem formulation given in Eqs. 4.22 and 4.23.

$$minimize \sum_{t=0}^{T} \dot{m}_{fuel}(t) \tag{4.22}$$

$$
\begin{aligned}
&\text{subject to} \sum_{t=0}^{T} P_{batt}(t) \approx 0, \\
&P_{elec}(t) + (\gamma P_{ICE}(t) + (1-\gamma) P_{brake}(t)) = P_{req}(t), \\
&\gamma\alpha|P_{elec}(t)| + (1-\gamma)\beta|P_{brake}(t)|, \\
&\text{where } \gamma = \begin{cases} 1, P_{req}(t) \geq 0 \\ 0, P_{req}(t) < 0 \end{cases} \\
&E_{batt\_min} \leq E_{batt}(t) \leq E_{batt\_max}, \\
&P_{batt\_min} \leq P_{batt}(t) \leq P_{batt\_max}, \\
&P_{RTM\_min} \leq P_{RTM}(t) \leq P_{RTM\_max}, \\
&\omega_{RTM\_min} \leq \omega_{RTM}(t) \leq \omega_{RTM\_max}, \\
&P_{BAS\_min} \leq P_{BAS}(t) \leq P_{BAS\_max}, \\
&\omega_{BAS\_min} \leq \omega_{BAS}(t) \leq \omega_{BAS\_max}, \\
&P_{ICE}(t) \leq P_{ICE\_max}, \\
&\omega_{ICE}(t) \leq \omega_{ICE\_max},
\end{aligned}
\tag{4.23}
$$

4.3 Dynamic Programming Algorithm Implementation

In order to determine performance benchmarks for the EC2 architecture, optimized ICE and electric machine power profiles over a drive cycle are obtained through the use of a dynamic programming algorithm. This algorithm determines the globally optimal solution to the objective function described in equation 4.23. A multi-stage approach using the method of coarse partitions was employed, to allow for reasonable execution time while still maximizing the resolution of the results.

4.3.1 Model Setup for Benchmarking

The selection of DP to perform the benchmark analysis poses an interesting challenge, from a modelling perspective. Although a range of existing EcoCAR team low- and high-fidelity vehicle and component models are available, as outlined in Chapter 3, the majority are forward-looking empirical and physics-based models implemented in Simulink for controls development purposes. By their nature, and due to the software

platform difference, these are not suitable for use in a backwards-looking analysis in their current state. The benchmark algorithm is also more readily implemented in MATLAB than Simulink, due to the latter's vector-based processing speed.

Since the primary concern of this work is in studying the energy consumption and power output of the various powertrain components, the most critical information contained in the component models is that which defines component efficiency. The EcoCAR team's empirical component models, specifically the ICE, and the RTM and BAS motors, are all built around a core of manufacturer-provided efficiency data. For this work, the powertrain component efficiency and physical operating limit data was extracted from the corresponding Simulink models and converted to 1D and 2D maps in MATLAB, to enabled direct use with the MATLAB-based DP algorithm.

4.3.2 DP Problem Formulation

To reduce computational complexity, and thus simulation time and memory use, the powertrain and vehicle are modelled as a discrete-time system. The driving cycle is discretized to a time resolution of 1 second per step. Each time step, referred to as a 'stage' in DP terminology, is represented by the variable k. The amount of energy in the ESS (essentially the SOC) is selected as the state variable. The selection of this particular quantity is done to take advantage of an important aspect of the backwards-induction method of DP. In this case, the final state of the system at stage $k = T$ becomes an initial condition of the DP problem, and by selecting an identical value at the initial state (at stage $k = 1$), CS operation can be enforced. Also, it becomes simple to avoid exceeding the upper and lower limits of ESS SOC, as they are inherently defined by the boundaries and range of the state space. The state space is initially discretized using a step size of 1 Watt-hour (Wh), and is represented by $E(k)$. This step size was determined experimentally, by examining the execution time, quality of results, and memory use resulting from a range of possible step sizes to achieve the best and most balanced performance in all categories.

If using the full ESS capacity of 16.2 kWh, this means there are 8100 discrete values of the state variable, which imposes a considerable computational burden. This was reduced through the following assumptions, based on knowledge of the standard automotive driving cycles. Consider that typical CS operation occurs around a specific ESS SOC, normally ranging from 15% to 30% of maximum (see Fig. 2.1 for an illustration of this). It is unusual to charge the ESS significantly above that

setpoint, due to the excessive fuel consumption required and/or the low likelihood of a sustained, aggressive regenerative braking opportunity. Therefore, the upper limit of the state space can be set slightly above the nominal CS setpoint. It is also unrealistic to expect the ESS to discharge right down to 0% SOC, as the driving cycles do not require a sustained power output of enough magnitude. The upper and lower limits were set experimentally, based on observation of state value trajectories during the DP development process, to 22% and 14% respectively, yielding a total state space size of 1296 discrete levels.

The primary control input that changes the state of the system is the electrical power drawn from, or added to, the ESS, represented by $P_{elec}(k)$. Every change in the state variable $E(k)$ (or change in SOC) is caused by a certain instantaneous power demand. This leads to the following state transition function:

$$E(k+1) = -P_{elec}(k)\Delta t + E(k) \tag{4.24}$$

where the Δt is the time step of the simulation. The negative sign is inserted to preserve the power flow direction convention of negative and positive power flows representing ESS discharge and charging, respectively. The 1Wh step size of $E(k)$ results in a $P_{elec}(k)$ step value of 3600W.

The cost function determines the fuel and electric power inputs required to make a given simulation step, depending on vehicle operating regime and conditions as given by the model. As the goal of the simulation is to minimize fuel consumption while keeping powertrain components within the bounds of their safe operating limits, the cost-to-go of a simulation step is therefore the sum of $\dot{m}_{fuel}$ and a set of static and distance-based penalty functions designed to further implement the constraints of Eq 4.23. The two hybrid operating regimes (series and parallel) are simulated separately, the results are compared during post-processing, and the final optimal ESS trajectory and control inputs are determined for each operating regime.

The following steps describe the operation of the developed DP algorithm. Note that i indicates the state grid position at stage k, and j indicates the grid position at stage $k+1$.

1) Discretize the state space into a grid. Discretize the requested power P_{req}. Starting at the final stage ($k = T$), select a desired 'initial' value for E_k on the grid ('E_T').

2) Examine each transition from state i at stage $k = T-1$ to E_T at $k = T$. Cal-

culate vehicle powertrain operating conditions over the transition (component speeds and efficiencies, transmission gear, etc.). Determine the cost of each transition, as follows:

a) Calculate the instantaneous electric power $P_{elec}(k)$ required to achieve that transition, and from that the RTM power $P_{RTM}(k)$, using the RTM efficiency map.

b) Calculate the instantaneous ICE power $P_{ICE}(k)$ (if accelerating) or mechanical braking power $P_{brake}(k)$ (if decelerating) required to satisfy the constraints of Eqs. 4.14 to 4.19.

c) Determine the resulting fuel cost ($\dot{m}_{fuel}$) by interpolating the brake-specific fuel consumption map.

d) Apply penalty functions to satisfy the physical constraints given in Eq. 4.23, as applicable.

e) Record the total cost of the transition, $g_{ij}(k)$.

3) Working backwards from stage $k = T - 2$ to $k = 1$, examine the transitions from each possible state value i at stage k to each possible state value j at stage $k+1$, repeating the actions of Steps 2a-2e. This generates the 'cost-to-go' matrix.

(a) At each stage, determine and record the total path cost $J(k)$, which is the sum of the minimum cost to go from state i at stage k to stage $k+1$ and the cumulative cost to transition from $k+1$ to the final stage $k = T$. Along with this total cost, record the corresponding state value j at stage $k+1$ that constitutes the next step in the optimal path.

4) With all transition costs computed and recorded, perform a forward-looking analysis to find the optimal trajectory, as follows:

(a) At stage $k = 1$, select a desired initial state value E_1. Using the total path cost recorded at Step 3a, determine the optimal next state E_2.

(b) Determine the (optimal) values of $P_{RTM}(k)$, $P_{ICE}(k)$, and $P_{brake}(k)$ that result in the transition from E_1 to E_2.

(c) Continue stepping forward, following the lowest total path cost at each stage, from stage $k = 1$ to $k = T$.

5) The optimal control policies can now be identified from the optimal trajectories of P_{RTM}, P_{ICE}, and P_{brake}. For the series regime, the optimal ICE control policies can be determined separately, by finding the optimal speed and torque at each time step.

4.3.3 Combined ICE and BAS Operation

Finding the operating points of the BAS and ICE in the series regime required some additional assumptions and manipulations. In the DP algorithm, the BAS output (electrical) power is determined, followed by the ICE output power, to find the fuel required. However, the component output power, as determined through efficiency, is a function of torque and speed. The speeds of the RTM, and the ICE in the parallel regime, are constrained by the wheel speed and transmission gear, meaning the problem can be reduced to having a single degree of freedom. In the series regime, the ICE is unconstrained with respect to both torque and speed, making it more time-consuming to identify operating points. To reduce computational complexity, some mechanical constraints were used to reduce the dimensionality of the problem.

First, the BAS and ICE are coupled by a belt system, meaning that BAS speed must always be linearly proportional to ICE speed, related by the belt ratio (given by K_{belt}). Second, both components are assumed to operate together as a generator system; at each simulation time step, all torque produced by the engine must be fully absorbed by the BAS, in order for the ICE speed to remain constant. These constraints, given in Eq. 4.25, mean that when functioning as a generator, the ICE and BAS can be considered as a single unit.

$$\omega_{BAS}(t) = K_{belt}\omega_{ICE}(t), \tag{4.25a}$$

$$\tau_{BAS} = (1/K_{belt})\tau_{ICE}, \tag{4.25b}$$

To achieve this, the BAS and ICE efficiency maps were overlaid and combined, with the grid values set to correspond to the ICE's torque and speed. Effectively, the combined map gives ICE fuel input needed to provide the required power while

compensating for the additional losses imposed by the BAS and coupling belt. This combined 'generator' efficiency map was then used in a separate offline optimization to generate a lookup table linking generator power requests to optimal ICE torque and speed values. This lookup table is used in the series-regime DP computations to find the optimal ICE operating point at each time step. A comparison of the ICE and derived generator efficiency maps is shown in Fig. 4.1.

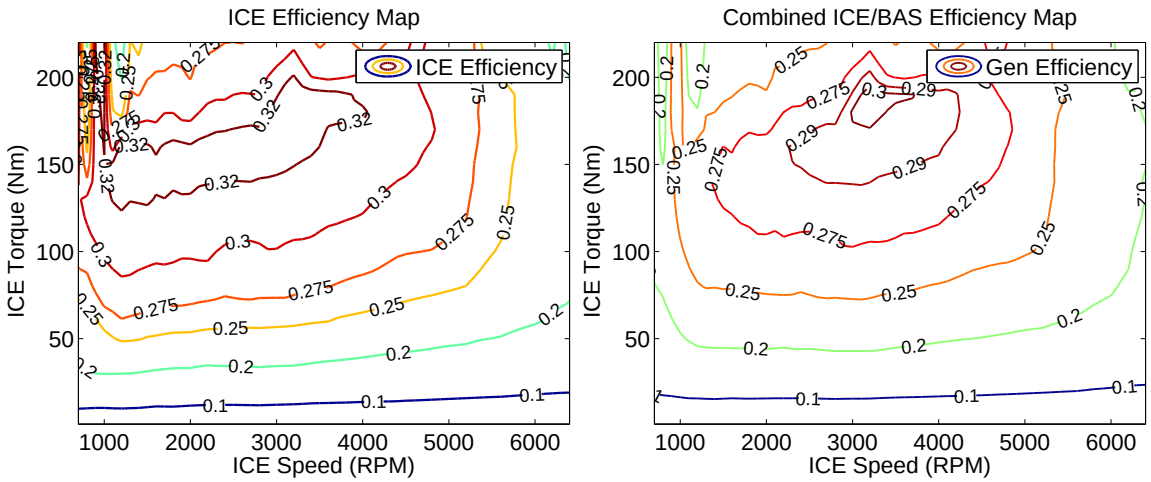

(a) LE9 efficiency map

(b) Combined 'Generator' efficiency map

Figure 4.1: ICE and 'Generator' efficiency maps

4.3.4 Solution Refinement

Due to the relatively coarse discretization of the state variable, the number of feasible state transitions between stages is low. For example, with ΔE_{ESS}=1 Wh, from any starting state, only 32 of the 1296 possible state transitions are feasible. Many valid solutions may exist between consecutive feasible paths, but are rendered inaccessible by the selection of state variable step size. If this situation occurs many times over a simulation, the resulting fuel economy value and performance curves will be sub-optimal. However, countering this with a decrease in the state variable step size (with a corresponding increase in feasible states), results in a drastic rise in computation time and memory usage. In order to achieve higher resolution results with respect to the state variable, yet still ensure reasonable run times, a progressive solution refinement approach was adopted.

The first refinement implementation utilized a method known as 'Coarse Partitions'. A first 'coarse' solution is obtained that defines the rough trajectory of the

state variable over the drive cycle. A new state space is created, centered on the coarse solution and bounded by +/-1 coarse step, that is divided into fine steps. A second pass performed that examines the solution within that region, yielding a finer resolution. This method was eventually rejected due to the inability to correctly sequence the solution series across the boundaries of consecutive time steps. With this method, two consecutive time steps will have the same number of states, but those states may represent different actual ESS SOC levels. Therefore, when selecting the optimal trajectory at a particular step k, the SOC value at k+1 may not be represented correctly when shifting to that step.

The current implementation determines the same initial 'coarse' solution as above. Next, the maximum and minimum state values are used as the upper and lower bounds of a new state space, which is again divided into 1296 states. This yields a new minimum P_{elec} resolution that is finer than the initial scan, with the level of improvement dependant on the new state range. The DP algorithm then calculates the optimal path within this new refined space.

In effect, the initial coarse search is performed to identify the boundaries of the region in which the optimal path is located, and the fine search refines this solution for a more accurate result. This progressive scanning can in theory be repeated several times to achieve further refinement, with the solution eventually converging. Only a single fine search is performed in this work, as the final resolution is adequate for this application and subsequent refinement does not change the fuel consumption levels significantly.

4.3.5 Performance Considerations and Improvements

A common compromise in modelling and simulation is between computation time and fidelity of results. In the case of dynamic programming, the term 'the curse of dimensionality' is used to describe the large rise in computation time and memory requirements as the size of the search space and number of problem states increases. To minimize this, it is essential to keep these variables at the minimum size that gives adequate results. This work was performed on a Lenovo ThinkPad laptop, with a 2.4 GHz Quad-Core CPU and 8GB of RAM; although very capable for many tasks, computing resources were relatively limited given the scope of the problem. Therefore, significant effort was spent in improving the algorithm execution efficiency.

In order to speed up computation time across the whole simulation, the constraints

of Eq. 4.23 were used to filter out infeasible states and avoid performing any calculations on them. For example, if no driving power is being requested ($P_{req}(k) = 0$), or a given transition would result in RTM power exceeding limits, no search is performed, and the transition is labelled as infeasible. This simple exclusion of infeasible points provided a 50-80% decrease in the number of evaluations required, depending on driving cycle. MATLABs Parallel Computing Toolbox was also used to further speed up computation, by enabling the state transition cost calculations to be conducted in parallel by the 4 available CPU cores. This reduced the time to calculate the coarse and fine transition costs (Steps 2 and 3 from Section 4.3.2) by a factor of four.

A final performance improvement was made by revising all calculation stages to utilize vector operations wherever possible. Initial revisions of the DP algorithm had the cost-to-go matrix populated element-by-element; while a logical first step when implementing the algorithm from pseudo-code, this method is very inefficient. The problem formulation allows for a simple approach to be used, where the cost-to-go can be calculated simultaneously from a given initial state to a vector of all possible next states. Given the number of states and MATLAB's inherent vector operation efficiency, this improvement yielded the largest reduction in computation time; initial implementations of the DP algorithm ran a 300s driving cycle in around 4.5 hours, while use of vector operations reduced that time to 0.43 hours, or 25 minutes.

It should be noted that in some DP problem formulations, the cost-to-go matrix can be re-used under different situations and conditions once calculated. In this work, the state transition costs are significantly influenced by the simulation parameters. The RTM and ICE efficiencies depend on vehicle speed (given by the driving cycle) and operating regime, and the degree of power split between ESS and ICE is influenced by the permissible level of regenerative braking. This means that a new cost-to-go matrix must be calculated for each driving cycle and regime - essentially for each of the simulations performed.

Careful memory allocation and array manipulation was also exercised to ensure the DP algorithm did not overrun the available RAM and lead to thrashing. Large amounts of data are generated when calculating the cost-to-go matrix; for example, the full cost-to-go matrix for the HWFET driving cycle would normally contain 1296x1296x766 32-bit floating-point elements, taking up 5146 MB (5.15 GB) of memory. This was reduced through use of MATLAB's Sparse Array object, and exploitation of the large number of infeasible state transitions. Sparse Arrays store only the array location and contents of non-zero elements, allowing massive arrays contain-

ing frequent zero elements to be stored in a memory-efficient manner. An encoding scheme was used that marked actual zero-cost elements and converted infeasible elements to zero, which enabled the massive cost-to-go arrays to be stored efficiently. Using this scheme, the cost-to-go matrix in the above example requires only 370 MB to store.

Note that this does not directly allow larger state spaces to be used or higher-resolution evaluations to be performed. The cost data still needs to be manipulated in MATLAB, which imposes on an upper limit of what can be managed. Additionally, this method relies on large numbers of infeasible transitions to be recorded, which is aided in part by the relatively low resolution of the scan. A final point to note is that the DP simulations are performed off-line, not in real-time in a vehicle, and so the computation time and memory consumption reductions were implemented to facilitate greater ease of use of the benchmark program, and to implement reasonable wait times, not to satisfy any in-vehicle operational requirements.

4.4 Methodology Summary

This chapter lays out the optimization problem formulation used in this work to determine the vehicle performance and optimal control policies, and describes the DP algorithm implemented to solve it. The core of the problem is the minimization of fuel consumption across a driving cycle, using *a priori* knowledge of the cycle to allow globally-optimal use of powertrain components. The problem is constrained by the physical limitations of the vehicle components, and the need to maintain CS operation.

The DP algorithm uses a backwards-induction approach over a discretized search space, to save memory and computation time. The ESS SOC is used as the state variable; this selection and the backwards-induction formulation mean that initial and final conditions can be imposed, enforcing CS operation. A tunable penalty function keeps SOC from diverging too far from the initial value, to ensure a valid CS solution is available. The initial 'coarse' solution is improved by a repeated search of a condensed search space, yielding the final result.

Chapter 5

Results and Discussion

This chapter describes the results of the DP analysis, including initial validation of successful operation and the fuel consumption values under the series and parallel regimes. Use of the BAS in the parallel regime and ICE/BAS operating regime, though possible in this architecture, are not considered in this work. Computational performance is quantified and evaluated with respect to the final resolution of results. Finally, an evaluation is performed of the results obtained when using the DP-generated control outputs to drive a Simulink powertrain model.

5.1 DP Algorithm Validation

Due to the complexity of hybrid electric vehicle powertrain systems, generation of the powertrain model and determination of an optimal performance solution analytically were not possible without making significant assumptions. Therefore, the validation of the results produced by the DP algorithm was carried out in simulation. Since automotive driving cycles are fairly long and complex, a simple testing cycle was developed where optimal solutions could be reasonably predicted. During the validation process, the upper and lower power limits and the state of the vehicle systems were modified from their real-world values, to force the algorithm to deal with specific conditions. This section describes the validation process and results.

5.1.1 Validation Cycle

A simple driving cycle was utilized for algorithm validation, consisting of a segment of constant acceleration from 0 to 100km/hr, a short cruise segment at 100km/hr, and a

segment of constant deceleration from 100 to 0km/hr. A plot showing the validation cycle, with accompanying power and torque requests at the wheels, is shown in Fig. 5.1.

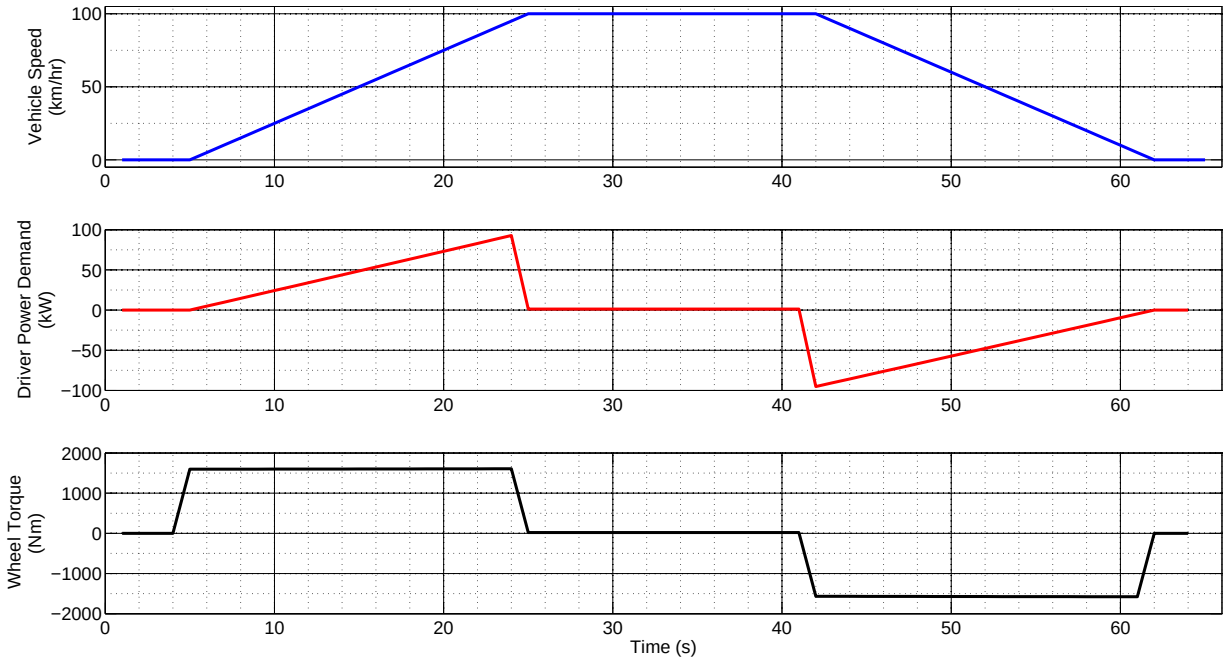

Figure 5.1: Validation cycle for DP algorithm

Note that the power and torque requests during the cruising phase of the validation cycle are very small but non-zero; the cruising power requirement is around 1.3kW.

5.1.2 Validation Results

The optimal solution to this cycle can be determined by inspection and simple analysis. Assuming regenerative braking to be 100% efficient at recovering the vehicle's kinetic energy, the trivial solution, using *a priori* knowledge of the cycle, would be to discharge the ESS through the RTM to accelerate the vehicle, and fully utilize regenerative braking during the deceleration phase to recover the discharged energy and restore the SOC to its initial state. Realistically, the finite efficiencies, as well as the upper and lower power limits of the powertrain components, render this solution infeasible.

If the RTM power is limited, the RTM will still be used, but the ICE or mechanical brakes must supplement it or take over when P_{req} exceeds the RTM limits. Fig. 5.2 shows the optimal solution with RTM power limited to +/− 50kW. After t=15, the RTM maximum power limit is exceeded, and the ICE provides the balance of power required to meet the driver power request. At t=41, the deceleration segment commences and braking force is applied by both the RTM and mechanical brakes, since braking solely with the RTM at this point would violate the RTM minimum power constraint. It is also worth noting that the ICE alone is used to propel the vehicle during the cruising phase of the cycle; due to the limitation on RTM power, the RTM is already regeneratively braking to the fullest extent possible, and thus there is no further regen opportunity to offset additional ESS discharge through electric propulsion. Due to its very high energy capacity, the ESS only discharges by 1.3% of capacity, or 211Wh, during the validation cycle.

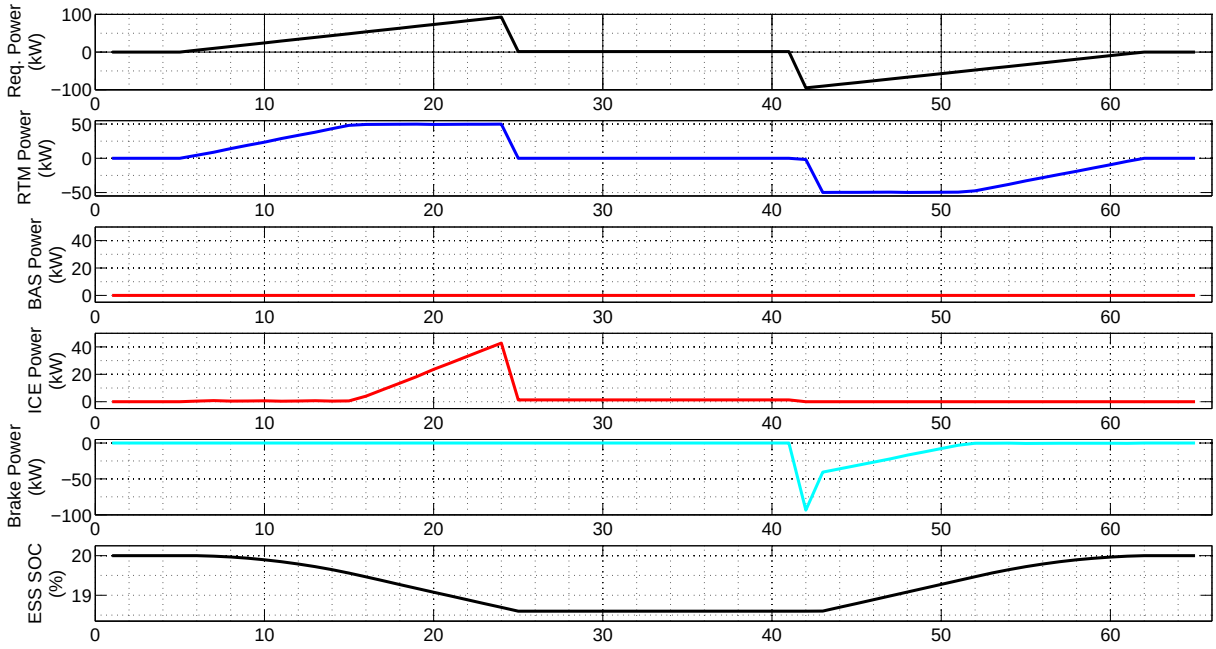

Figure 5.2: Validation cycle results, Parallel, RTM power limited to 50kW

Fig. 5.3 illustrates a different case, where full RTM power is available, but the cycle is started at near-minimum SOC. The ICE is used to propel the vehicle at the start of the cycle, instead of the RTM. The DP algorithm exploits its *a priori* knowledge of the approaching minimum ESS SOC limit to command additional charging through

the BAS to offset the upcoming discharge. The RTM is then brought in to bring the vehicle up to cruising speed as the ICE ramps down. The minimum SOC is just reached at the outset of the cruising phase. During braking, the mechanical brakes are applied to begin the deceleration, with the RTM subsequently used to regeneratively brake and restore the SOC to its initial value.

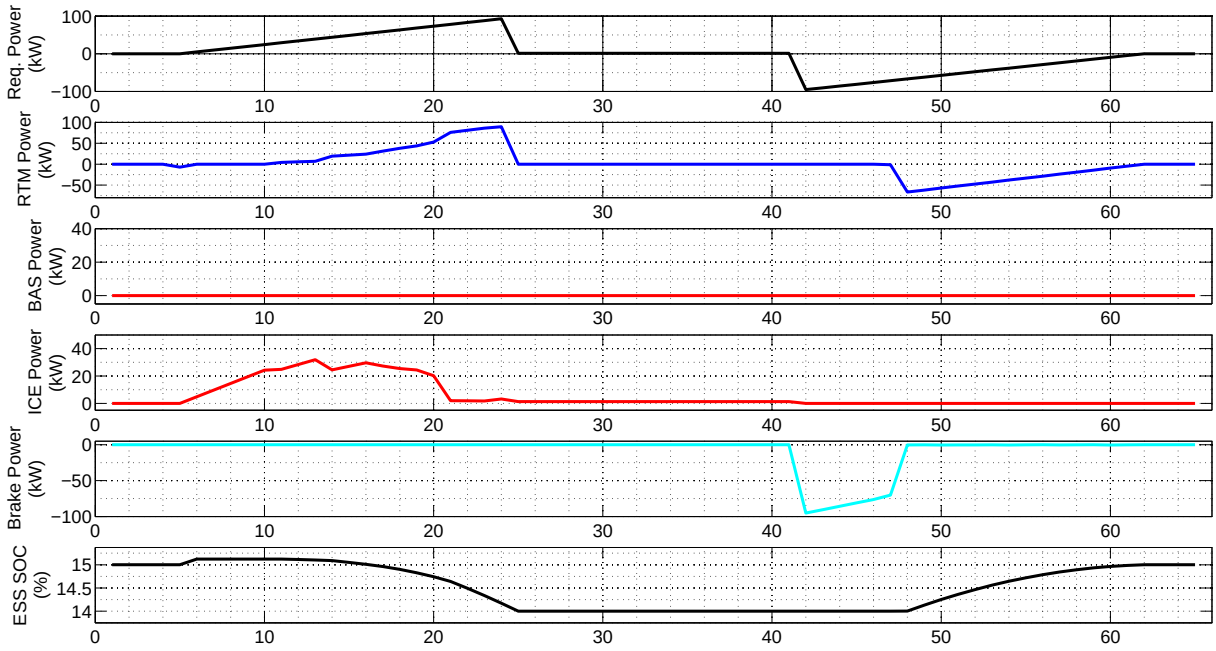

Figure 5.3: Validation cycle results, Parallel, low initial SOC

Fig. 5.4 illustrates the validation cycle while operating in series regime. Operation is as close to the trivial solution described above as the physical constraints and problem formulation will allow. With only the RTM as a traction power source, the RTM power must follow exactly the requested power during acceleration. At t=41 the mechanical brakes are very briefly applied, because the requested braking power is slightly larger than the lower limit of the RTM. This means the initial ESS discharge is slightly greater in magnitude than the recoverable energy during deceleration, resulting in the low-magnitude ICE output charging through the BAS.

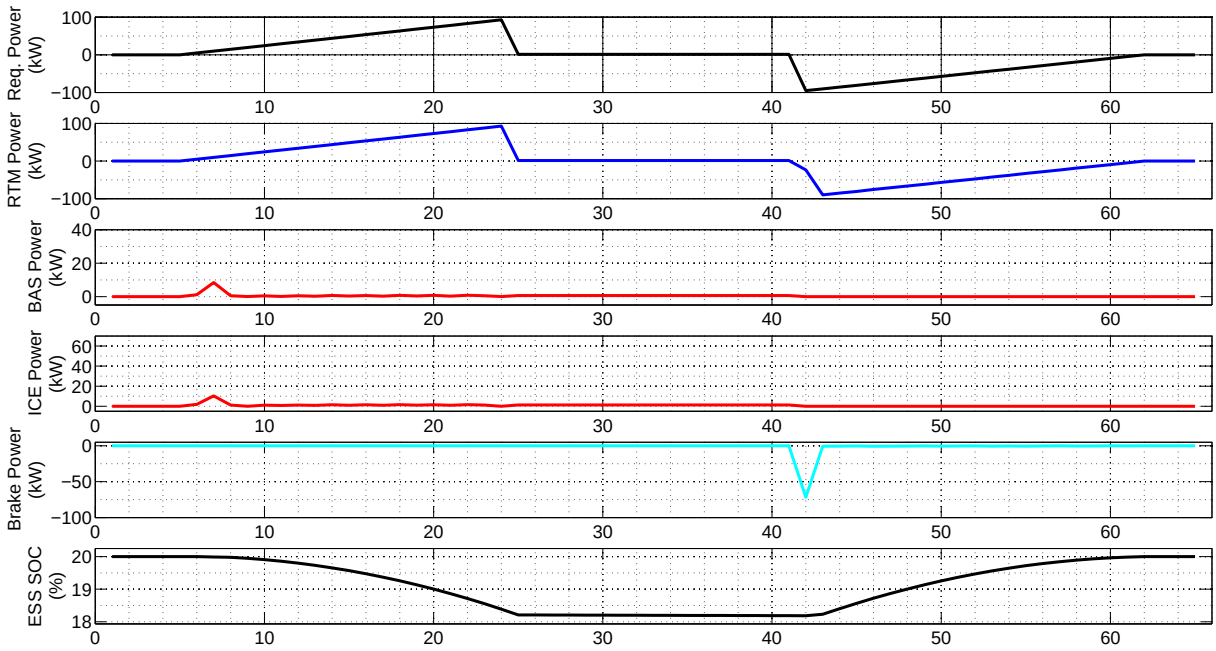

Figure 5.4: Validation cycle results, Series

Fig. 5.5 illustrates the validation cycle while operating in series regime with a low initial SOC. Since the ESS will be depleted to its lower limit (here set to be 14%) before the end of the acceleration phase of the cycle, the ICE offsets the expected discharge by aggressively charging the ESS through the BAS earlier in the cycle. The mechanical brakes are used more heavily during braking in this case, as the SOC target is close enough to the lower limit that only a small duration of regen braking is required to replenish it.

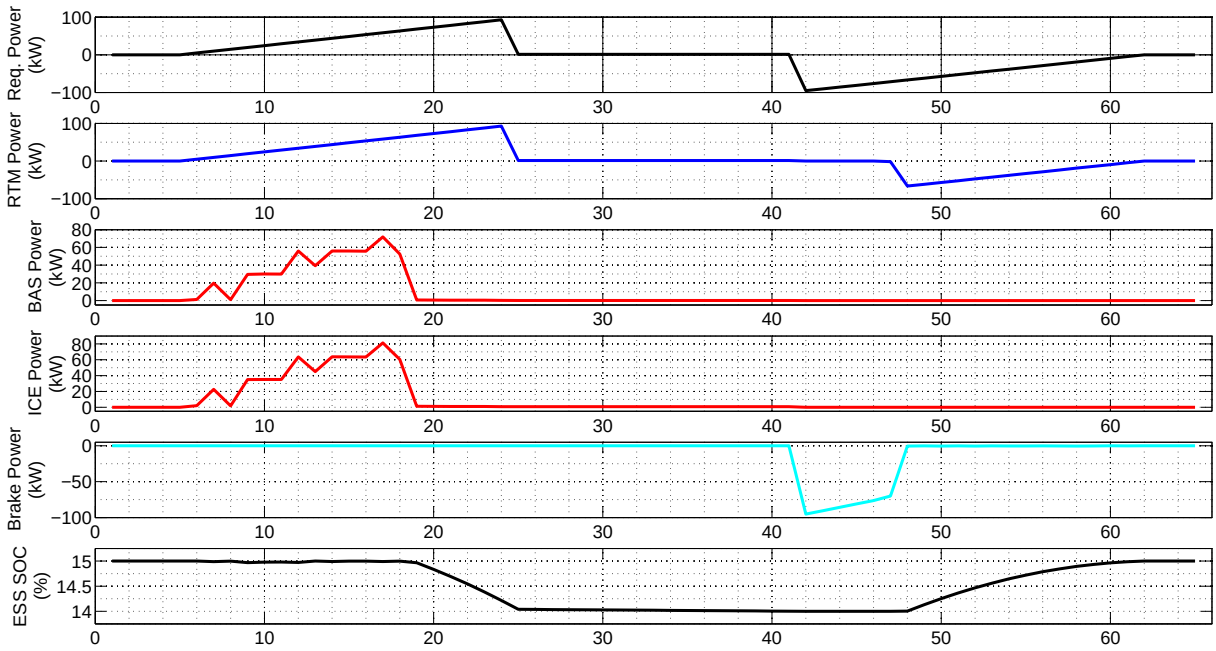

Figure 5.5: Validation cycle results, Series, low initial SOC

In Fig. 5.5, and in many of the plots of subsequent cycles, the ICE and BAS power plots are not smooth, but feature 'squiggles', or up/down value changes at each time step. These result from a combination of the DP algorithm's independent/dependent variables and the choice of quantization levels, and will be explained here. Recall from Section 4.2 that to simplify the computation, the ESS power is selected as the independent variable in the simulation. ESS power also serves as the control input that changes the quantized state variable, ESS state of charge. By the constraints of Eq. 4.14, 4.15, 4.18 and 4.19, the RTM, ICE (and BAS in series regime), and mechanical brake power values are dependent on the ESS power and driver requested power.

Quantization of states and power levels is essential to keeping DP algorithm memory use to a manageable level. Quantization of ESS SOC and resulting quantization of the ESS power means that the values selected as the dependent variable will also be quantized. As the (continuous) requested power crosses one quantization level and increases toward the next, the dependent variable will make a corresponding jump down close to 0. The coarse time discretization of one second per step also contributes to the roughness, as value changes appear sudden when plotted at this level.

It is not possible to test the limited RTM power case with the series regime, as the RTM is required to provide all vehicle propulsive power, and it can be readily observed that more than 50kW is required for the vehicle to fully follow the validation cycle. The results from Figs. 5.2 to 5.5 demonstrate that the DP algorithm developed for this work provides valid optimal solutions for the given driving cycle.

Minor tuning of the penalty function was performed during algorithm validation, to ensure CS operation. Tuning was carried out by adjusting penalty coefficient values until the penalties term was in the right order of magnitude to balance fuel consumption. Once tuned, the penalty function successfully enforced CS behaviour in all subsequent cycles.

5.2 Optimal Fuel Economy Analysis

5.2.1 Driving Cycles Used

Driving cycles provide short representations of the operating conditions typically encountered by light-duty vehicles under different circumstances, enabling repeatable and timely testing to be performed. The US EPA, regarded as the North American authority on automotive testing standards, develops and uses a range of driving cycles to test vehicle performance and fuel economy to a common standard. Four such standard driving cycles are used in EcoCAR2, and in this work. These cycles are described in this section, and are shown along with the associated wheel power and wheel torque profiles required by the EcoCAR 2 vehicle.

The Urban Dynamometer Driving Cycle, or 'UDDS', simulates urban driving conditions in 'stop-and-go' traffic. As it is long and repetitive, the first 505 seconds of this cycle are most commonly used, being representative of the overall conditions. The UDDS cycle is shown in Fig. 5.6.

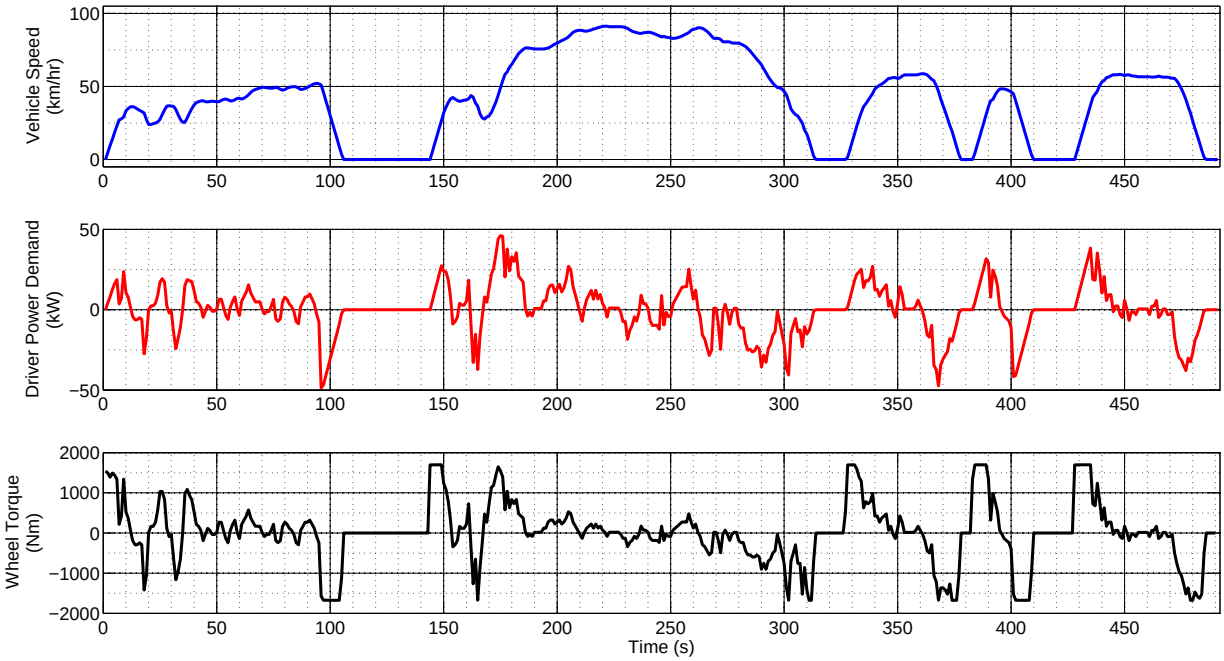

Figure 5.6: Urban Dynamometer Driving Cycle (UDDS)

The Highway Fuel Economy Driving Schedule, or 'HWFET', represents long-distance highway cruising conditions with speeds kept under 90km/hr, and is shown in Fig. 5.7.

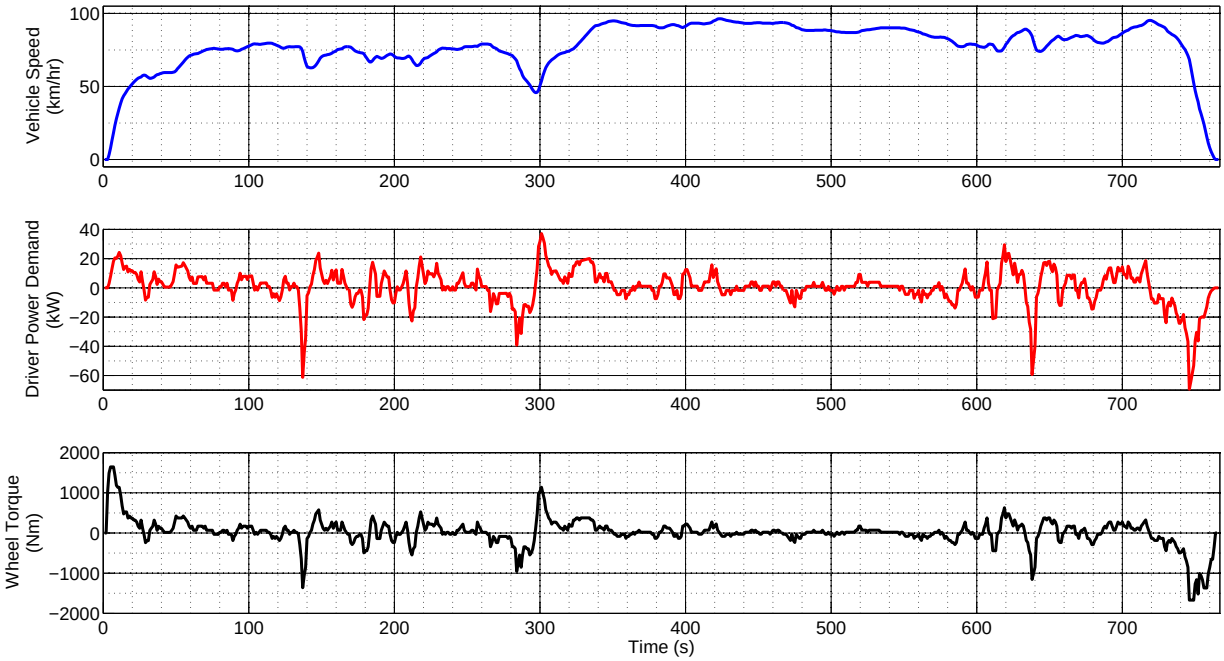

Figure 5.7: Highway Fuel Economy Driving Schedule (HWFET)

The US06 cycle is a newer, more aggressive cycle that includes both urban and highway segments. US06 was specifically implemented to include sharper accelerations and higher speeds, to provoke greater power and torque demands in vehicles. In this work, this cycle is split into separate highway and city segments that are run independently, in order to quantify and evaluate city and highway operation separately. The US06 City cycle is shown in Fig. 5.8, and the US06 Highway cycle is given in Fig. 5.9.

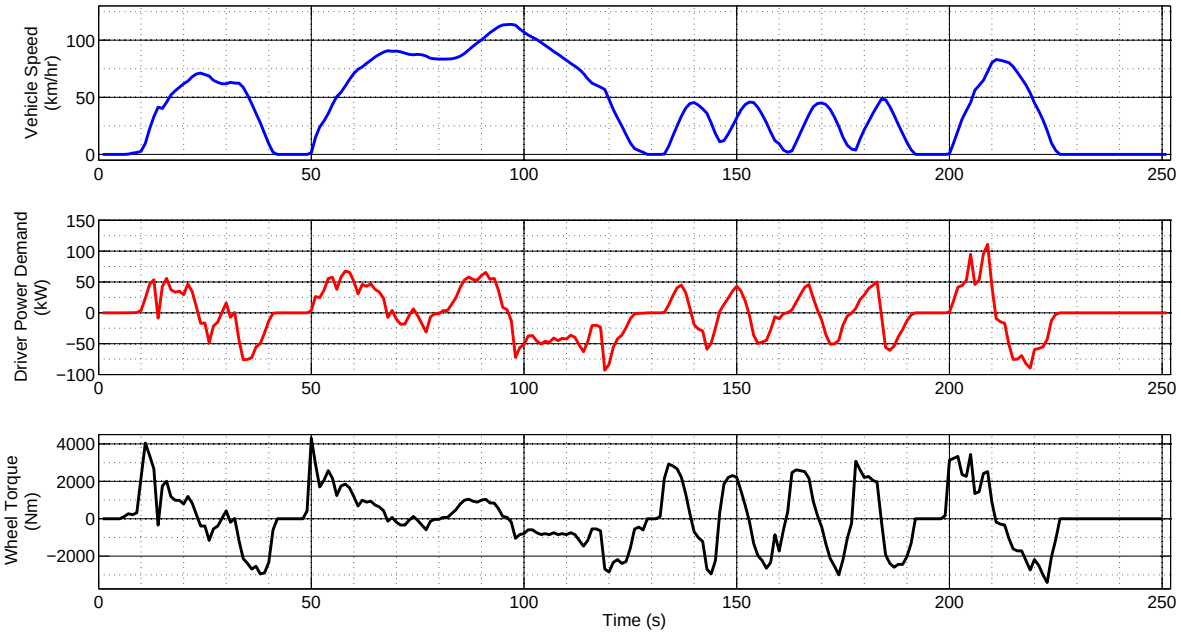

Figure 5.8: US06 City Cycle

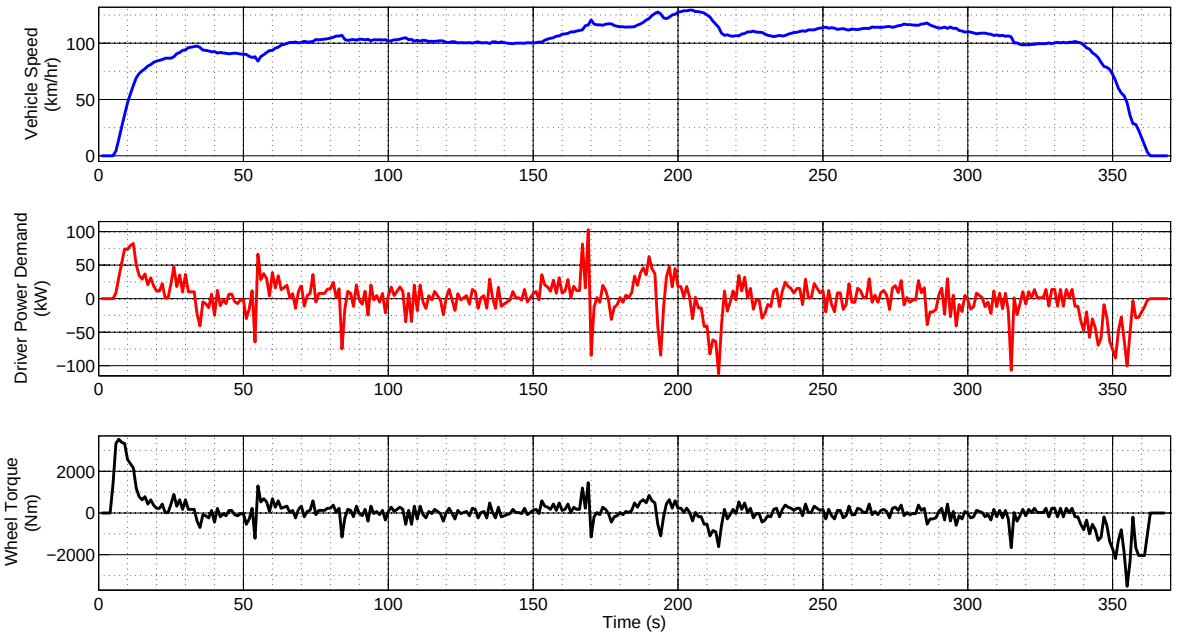

Figure 5.9: US06 Highway Cycle

Table 5.1 gives the distances travelled in each of the driving cycles used.

Cycle	**Distance (km)**
UDDS	5.78
HWFET	16.50
US06 City	2.85
US06 Highway	10.05

Table 5.1: Driving Cycle Distances

5.2.2 Fuel Economy Analysis Results

The CS fuel economy rating used in EcoCAR2, and in this work, is a weighted sum of the fuel used in the four standard driving cycles, as given by Eq. 5.2.2:

$$FE_{Total} = 0.29 * FE_{UDDS} + 0.12 * FE_{HWFET} + 0.14 * FE_{US06City} + 0.45 * FE_{US06Hwy} \quad (5.1)$$

This approach is an abbreviated method that produces results approximating the vehicle energy use seen in the 2008-enacted EPA five-cycle fuel economy testing regime, which is the current standard in North America for determining the 'window-sticker' fuel economy values of vehicles. In this work, fuel economy values obtained with this method are referred to as the 'EcoCAR2 4-cycle' value, or 'FE(4-cycle)'. Also, since E85 is the fuel used by the EcoCAR2 architecture, the combined FE value is given in Litres of gasoline-equivalent per 100km, or Lge/100km, a unit representing the amount of gasoline required to provide equivalent energy to the alternative fuel. Use of this unit allows direct comparison between this and any other gasoline-powered vehicle.

Using the DP algorithm and MATLAB model, the optimum CS fuel economy value was obtained for each of the four driving cycles, and the EcoCAR2 4-cycle value was calculated. The CS fuel economy values are given in Table 5.2.

Both regimes provide substantially improved fuel economy over the stock Malibu. However, the parallel regime outperforms the series in fuel economy on the US06 City cycle, and the series regime offers better fuel economy than the parallel on the

Drive Cycle	**FE(4-cycle)**	$(Lge/100km)$
	Series	Parallel
US06 City Segment	8.72	6.99
UDDS (505)	4.73	4.58
US06 Highway Segment	4.31	3.89
HWFET	3.12	3.89
EcoCAR2 4-cycle	**4.91**	**4.74**
Stock 2013 Malibu Eco	8.83	

Table 5.2: Fuel Economy Values

remaining cycles. To obtain more insight into this behaviour, a closer examination of the results from the US06 City and US06 Highway cycles is given, they being most representative of 'real-world' city and highway driving.

US06 City Cycle

The US06 City cycle is the most demanding in terms of power requirements, due to its aggressive accelerations. The results for this cycle are shown in Figs. 5.10 and 5.11.

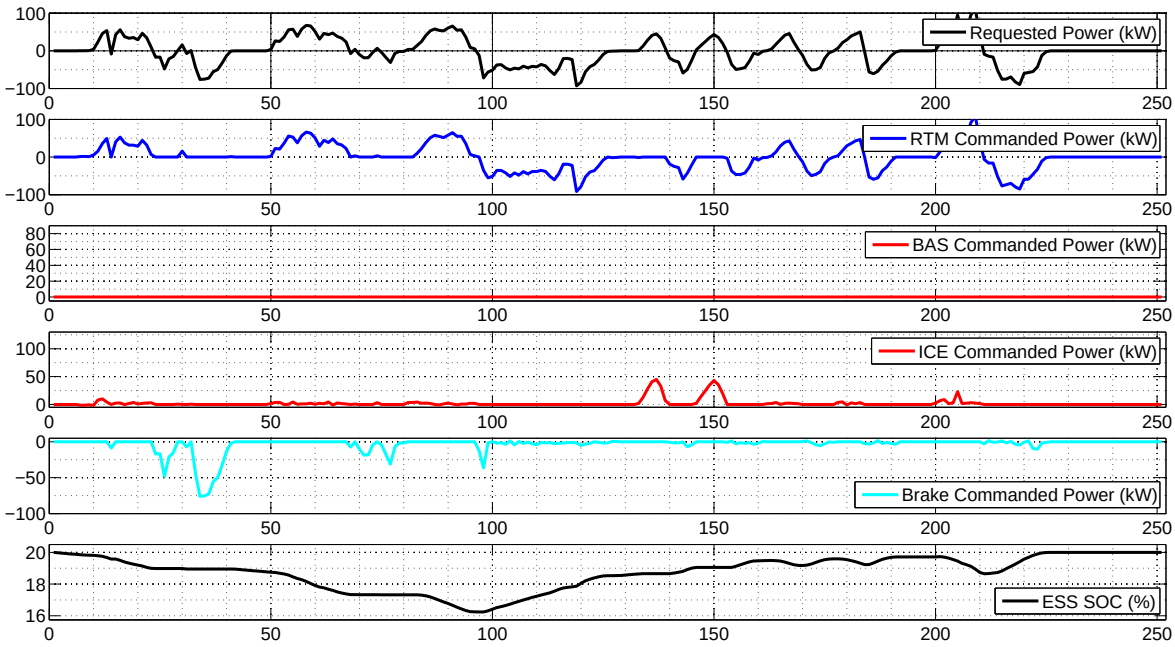

Figure 5.10: Power Outputs, US06 City, Parallel regime

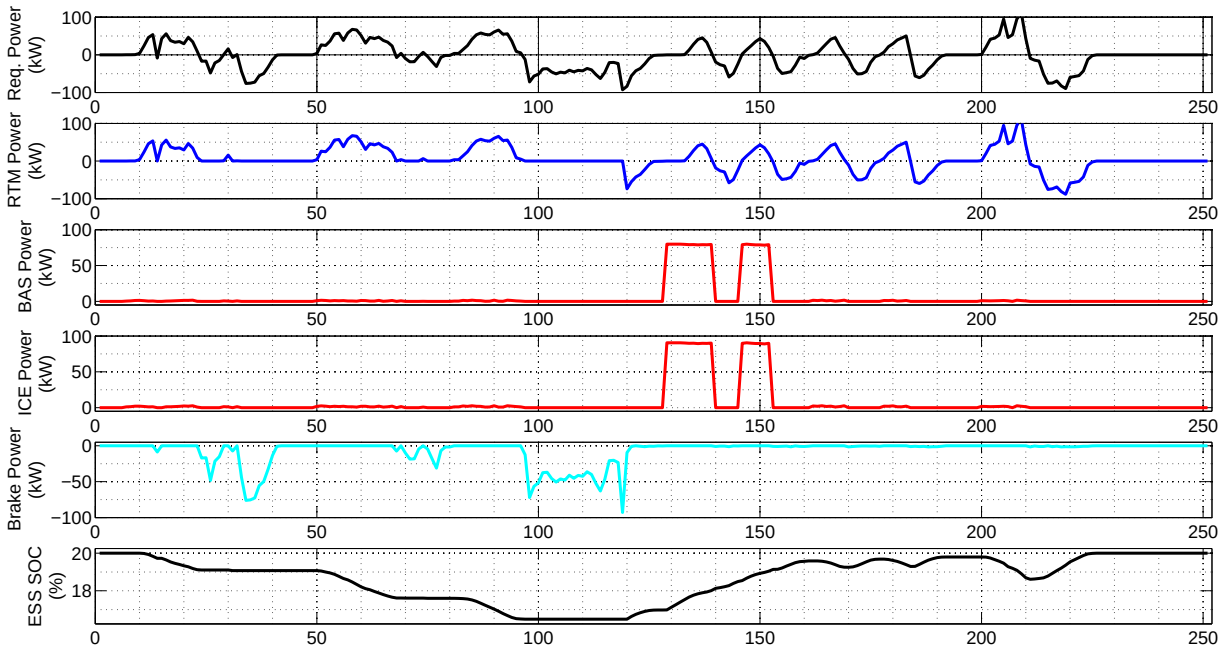

Figure 5.11: Power Outputs, US06 City, Series regime

RTM use is similar in both regimes, though naturally more extensive when operating as a series. Regenerative braking, as one of the main efficiency-enhancing abilities of a hybrid powertrain, is used to great effect in this cycle by both regimes, as the vehicle takes advantage of the repeated stops to replenish the ESS. The availability of this energy means that ICE power is only used to supplement the RTM during acceleration. This type of operation is an expected outcome of the team's design strategy; recall from Section 2.2.2 that a large energy-storage system was selected to displace fuel use as much as possible.

The primary difference in powertrain operation is in the use of the ICE. The parallel regime uses very little ICE power, except during a set of mid-cycle accelerations. In the series regime, the ICE is used to charge the ESS, through the BAS, in short-duration, high-power events. Although substantially more ICE power is used by the series regime, the fuel consumption increase over the parallel is moderate, due to the difference in ICE operating conditions.

Figs. 5.12 and 5.13 show the ICE operating points over the US06 City cycle in parallel and series regime. In the parallel regime, the ICE produces low power outputs, and ICE operating points are uniformly clustered in a low-speed, low-torque

operating area of low efficiency. This results from the physical constraint imposed by the transmission; ICE-to-road coupling through the transmission prevents the supervisory controller from adjusting ICE speed.

Fig. 5.13 illustrates the operating points of the ICE/BAS system, using the combined efficiency map described in Section 4.3.3. The red curves highlight lines of constant electrical output power, with 80kW being the maximum. Due to the series regime's de-coupled ICE operation, there is no constraint on ICE speed, and operating points are found at single values of torque and speed. These is the result of the algorithm picking the most efficient torque/speed combination for a desired power level, which typically lie closest to the high-efficiency plateau in the center of the map. Consequently, in this case the operating points are mostly clustered at the 80kW contour, as close to the efficiency plateau as possible. The algorithm uses short, high-power charging events to maximize overall efficiency and offset the inherent fuel-to-electric conversion losses in series powertrain operation.

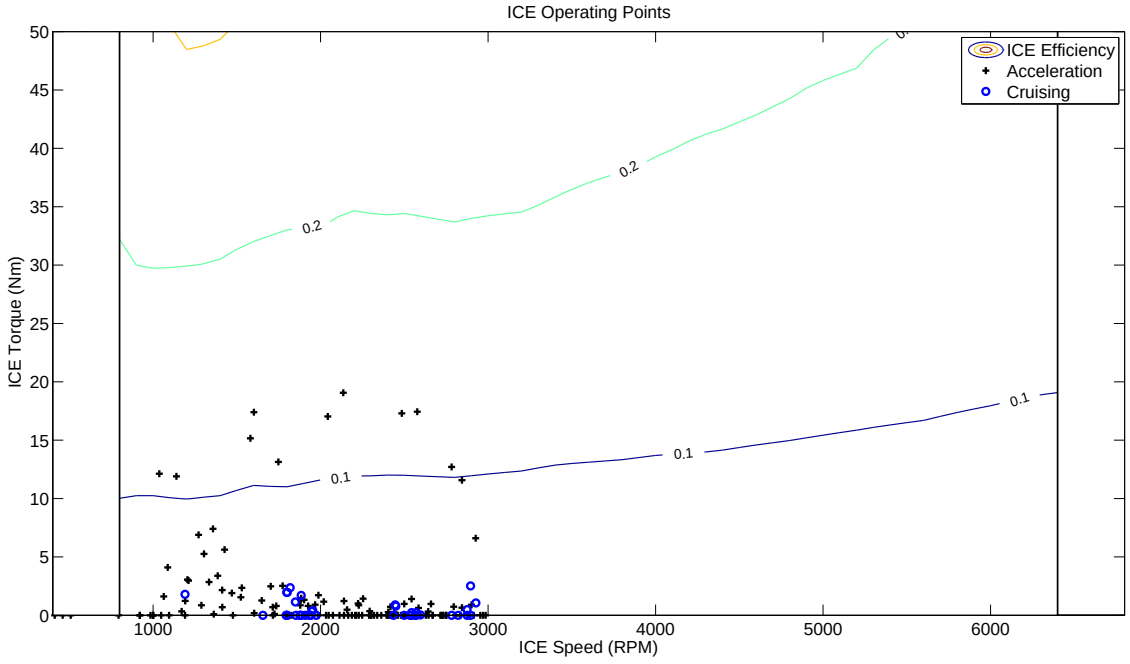

Figure 5.12: ICE operating points, US06 City, Parallel regime

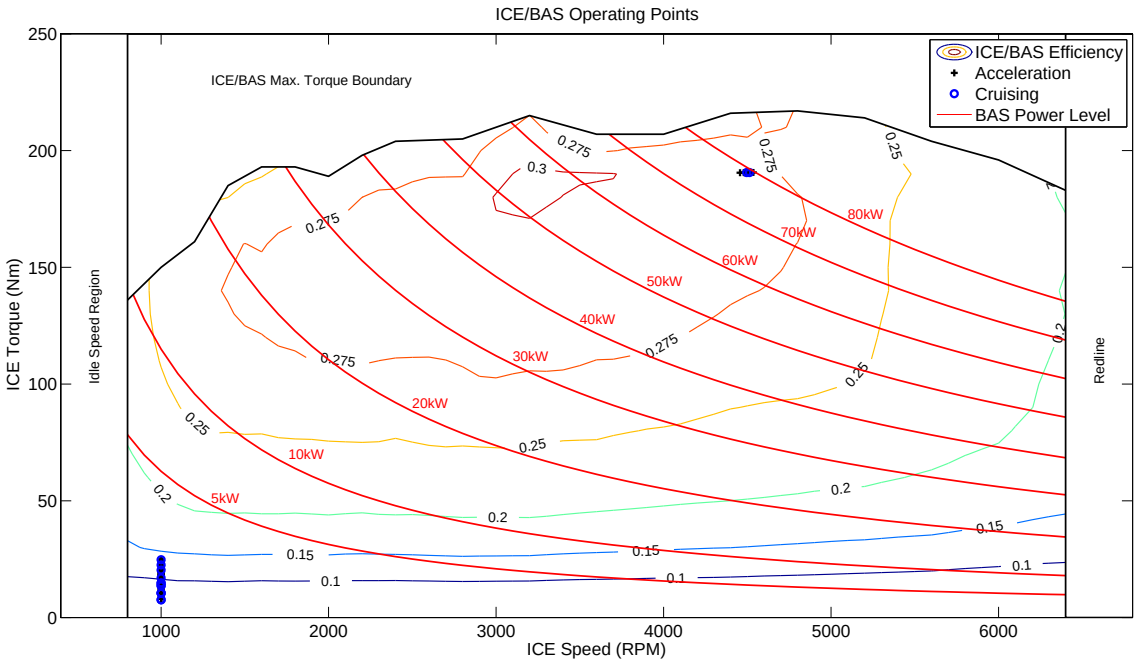

Figure 5.13: ICE operating points, US06 City, Series regime

The disparity of ICE efficiency between series and parallel regimes is further highlighted by the efficiency histograms of Figs. 5.14 and 5.15. The series regime's decoupling allows all ICE operating points to be set in a region of maximum efficiency for a given power level.

The RTM operating points can be seen in Figs. 5.16 and 5.17. The RTM speed is directly proportional to the wheel speed in both regimes, so only the RTM torque commands can be independently selected. RTM operating points for both regimes are similarly distributed for both series and parallel, which is not surprising; as noted above, the RTM is used heavily in both regimes to displace fuel use and maximize recovered energy. Operating points are distributed across the speed range, with acceleration points typically grouped at positive and negative high torque regions, and cruise points grouped at low torque and higher speed regions. In the parallel regime, acceleration points are slightly lower in torque magnitude, as the ICE is being used to supplement the peak acceleration torques, as seen in Fig. 5.10. As shown in Figs. 5.18 and 5.19, RTM operation over the majority of the drive cycle is in regions of high efficiency.

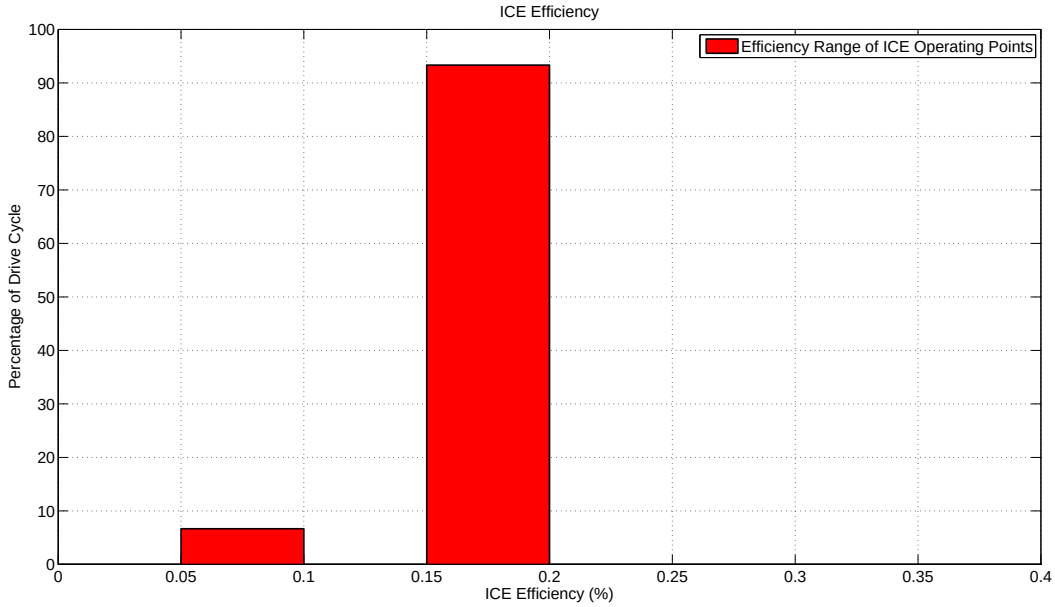

Figure 5.14: ICE Efficiency Histogram, US06 City, Parallel Regime

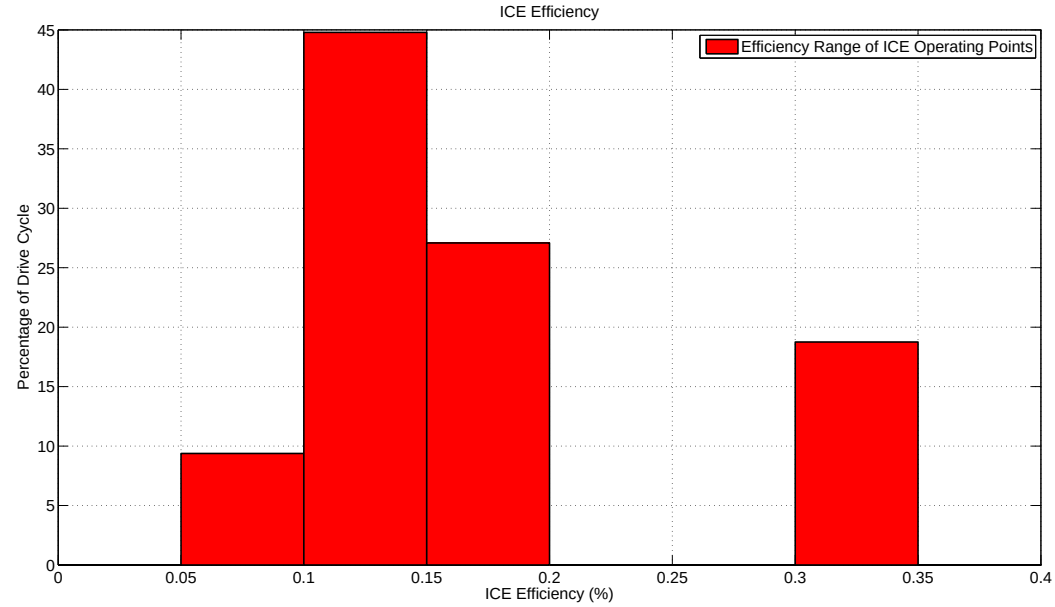

Figure 5.15: ICE Efficiency Histogram, US06 City, Series Regime

US06 Highway Cycle

The US06 Highway cycle requires the highest top speed of all the drive cycles, but the lack of frequent, substantial acceleration or deceleration events means that overall

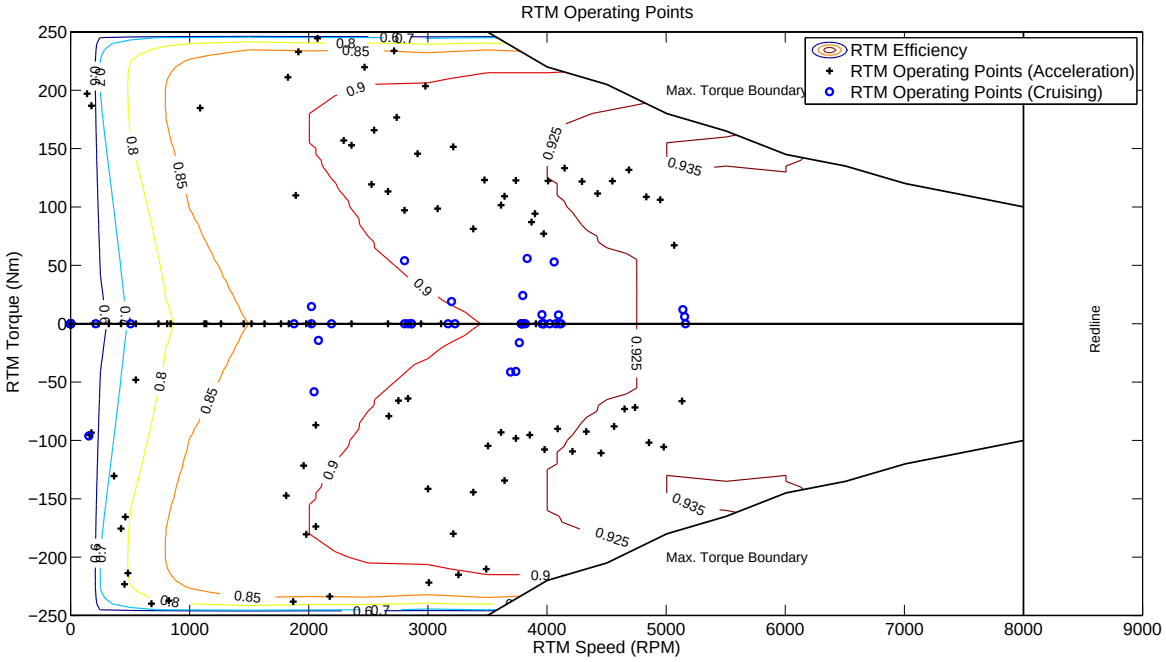

Figure 5.16: RTM operating points, US06 City, Parallel regime

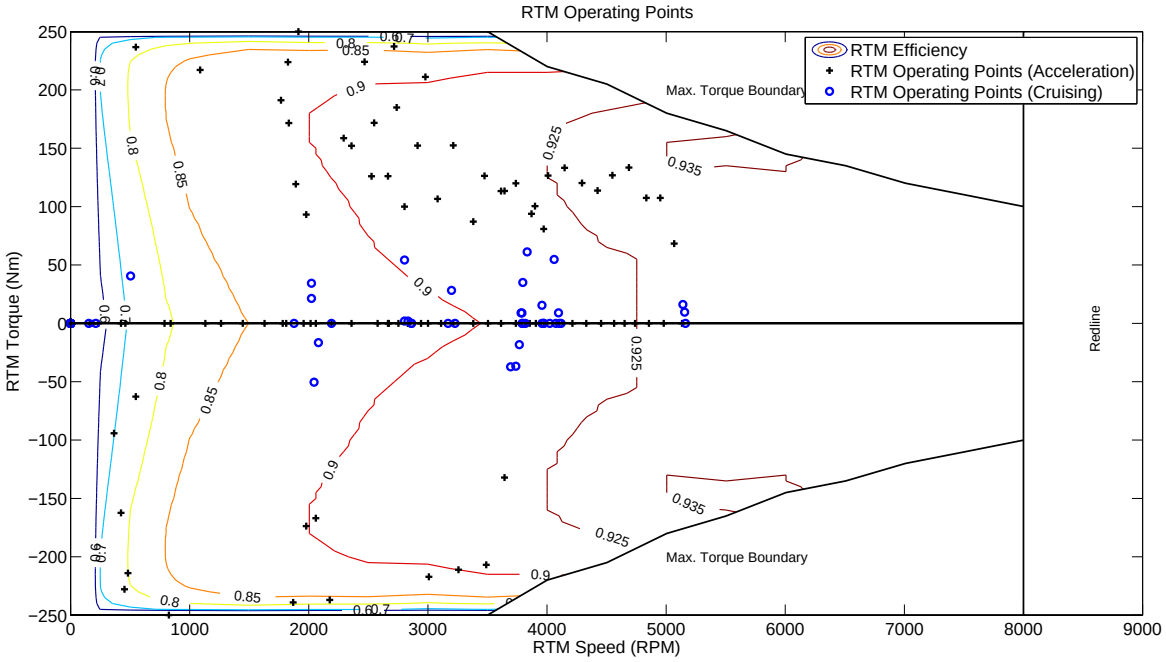

Figure 5.17: RTM operating points, US06 City, Series regime

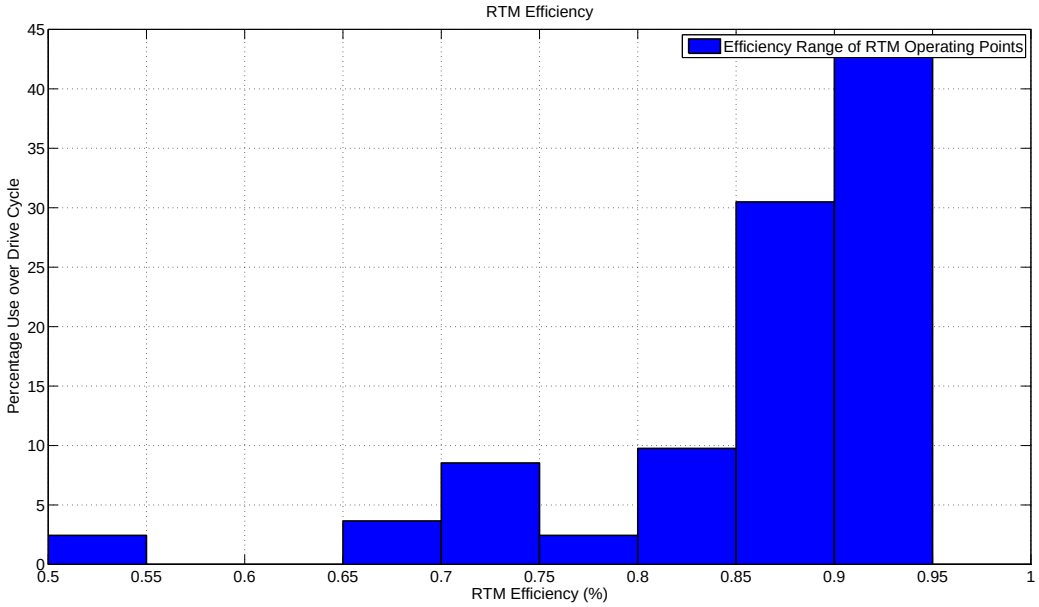

Figure 5.18: RTM Efficiency Histogram, US06 City, Parallel Regime

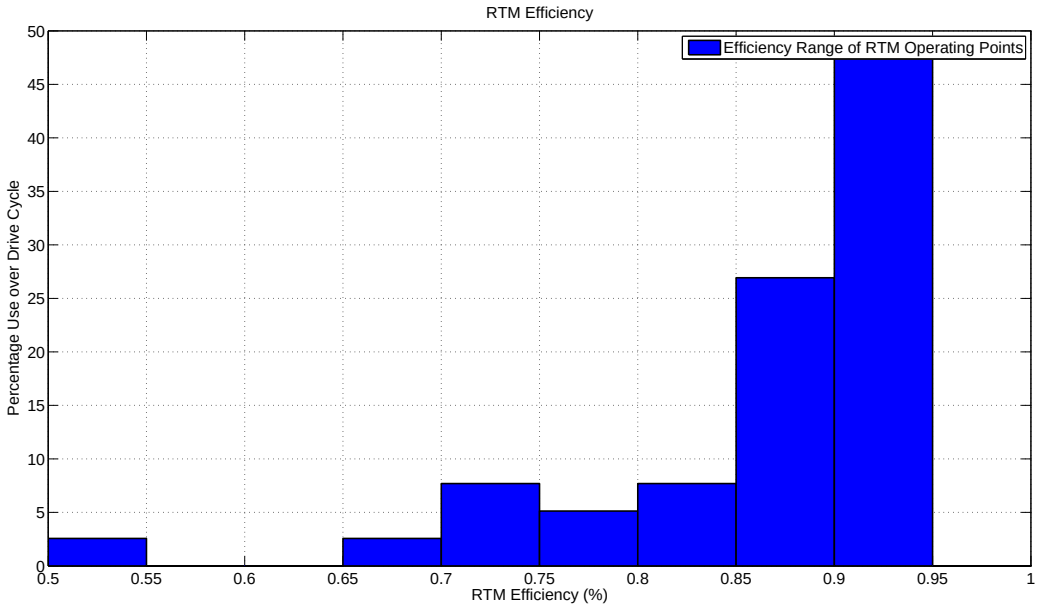

Figure 5.19: RTM Efficiency Histogram, US06 City, Series Regime

power requirements are quite low. The results for this cycle are shown in Figs. 5.20 and 5.21.

RTM use is again similar between the two operating regimes. There are not many

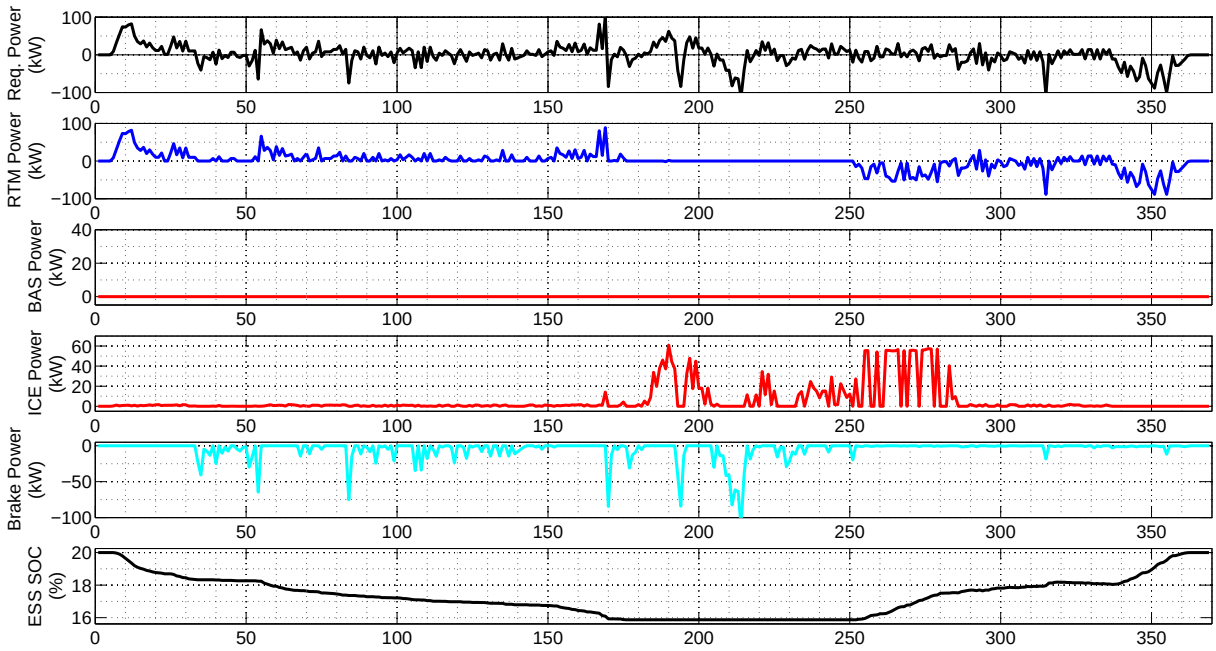

Figure 5.20: Power Outputs, US06 Highway, Parallel regime

opportunities for regenerative braking on this cycle, forcing the DP algorithm to reduce overall RTM use. Consequently, ICE power is used more than in the US06 City cycle, but the greater distance covered in this cycle results in a lower overall fuel consumption.

ICE operating points can be seen in Figs. 5.22 and 5.23. Operating points are more spread out, with a higher proportion of ICE operation during cruising than acceleration, compared with the US06 City cycle. The parallel-regime operating points are still clustered in a region of lower power demand and low efficiency, due to the low cruising power demand and the 6-speed transmission in the vehicle shifting to a high gear when torque requests are low. Recall that the transmission shift maps cannot be overridden in this architecture, and the transmission controller up-shifts during periods of low acceleration to reduce ICE torque. Series-regime ICE operation is again composed of short-duration, high-power charging events to obtain maximum efficiency. Several transitional operating points are located at low power levels, but in-line with the efficiency plateau occurring around 3100 RPM. The multiple 'on-off' behaviour of the ICE/BAS system is due to the low time resolution of the simulation, as can be seen from the frequent zero-crossings or near-zero power levels in the

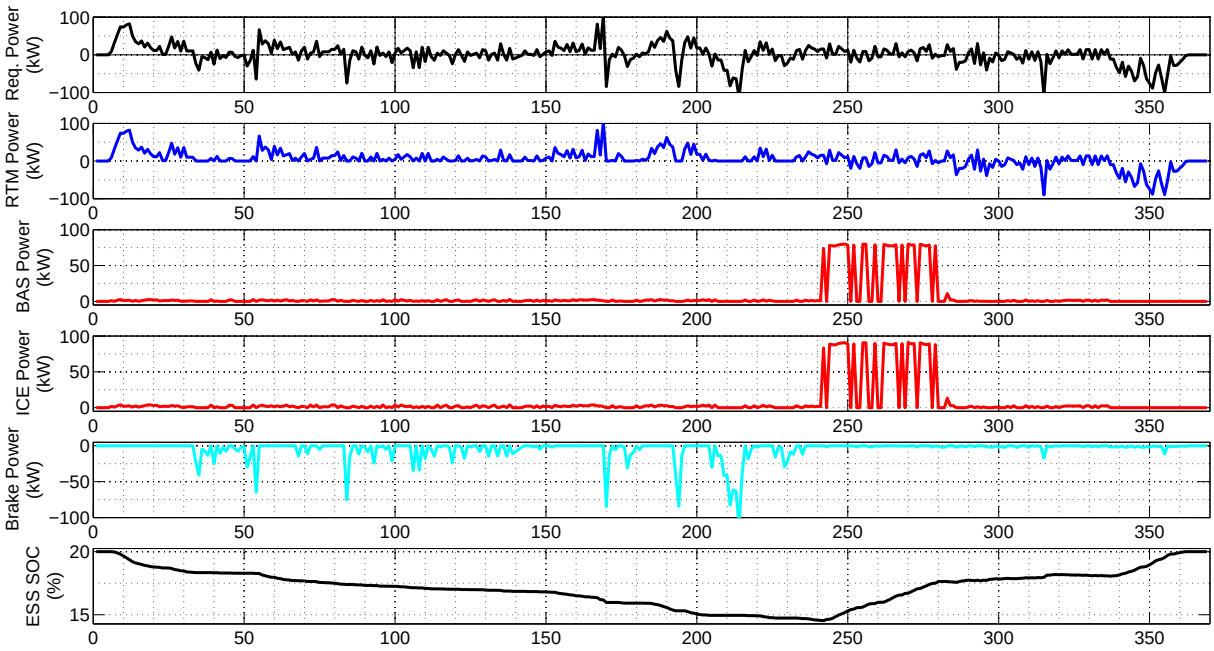

Figure 5.21: Power Outputs, US06 Highway, Series regime

'Requested Power' plot in Fig.5.23.

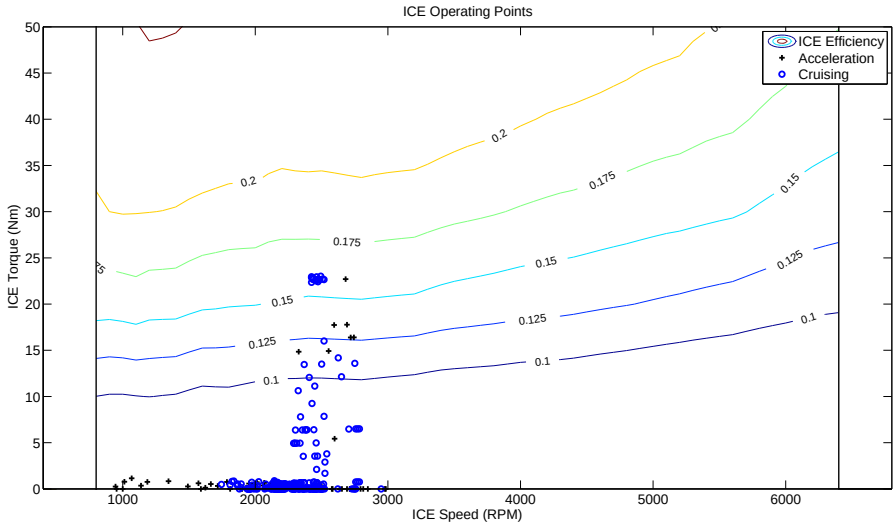

Figure 5.22: ICE Operating Points, US06 Highway, Parallel regime

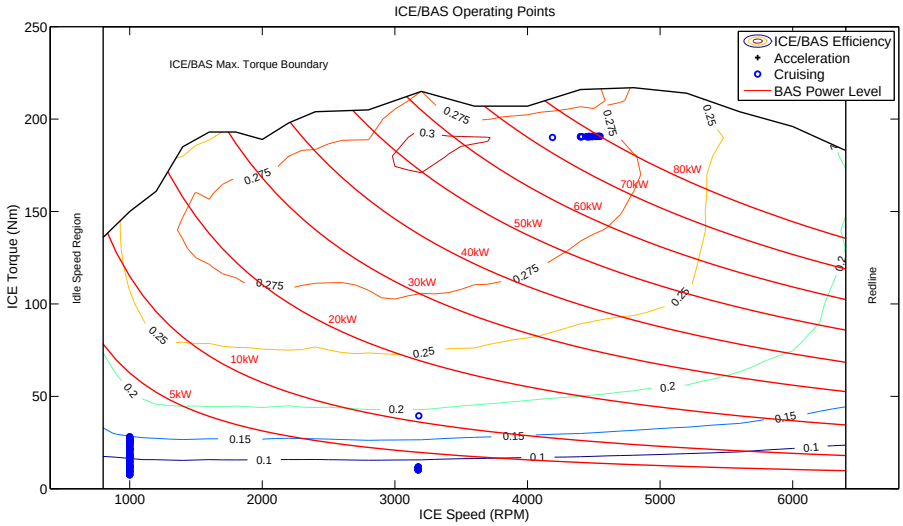

Figure 5.23: ICE/BAS Operating Points, US06 Highway, Series regime

Due to the high average cruising speed in this cycle, the RTM operates at higher rotational speeds, and therefore operating points are more closely clustered in the center regions of higher-efficiency for both series and parallel regimes. This is another expected result that validates the vehicle design strategy; the Magna RTM was selected on the basis of its efficiency under frequently-encountered operating conditions, such as the US06's 100km/hr cruising speed. This can be clearly seen in Figs. 5.24 and 5.25.

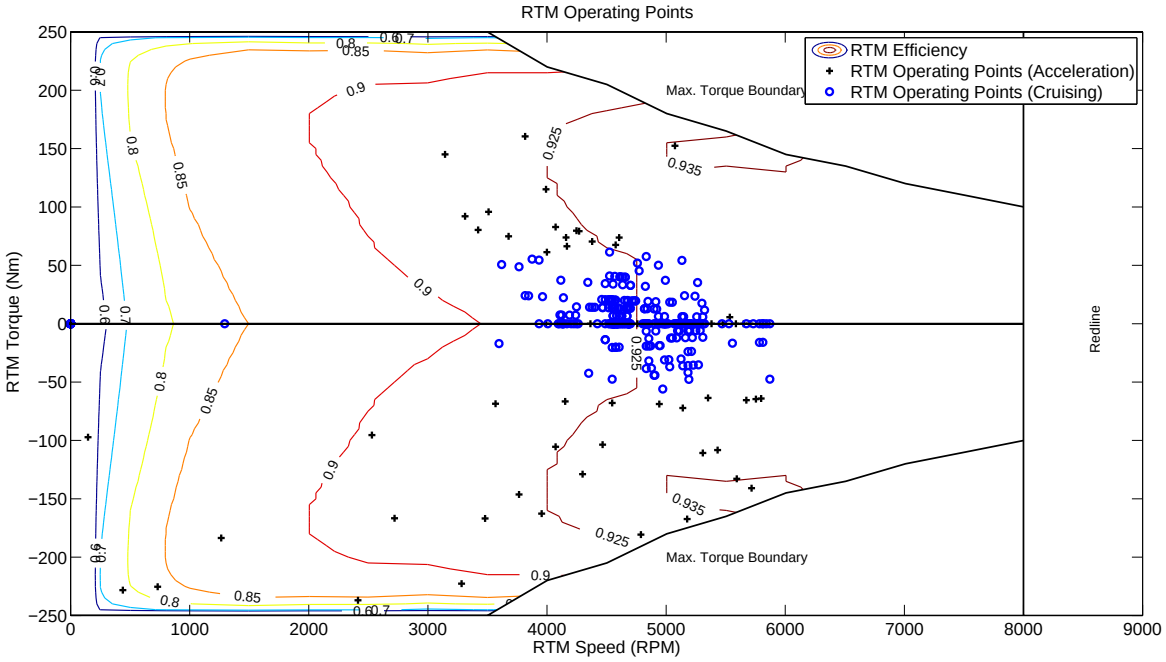

Figure 5.24: RTM Operating Points, US06 Highway, Parallel regime

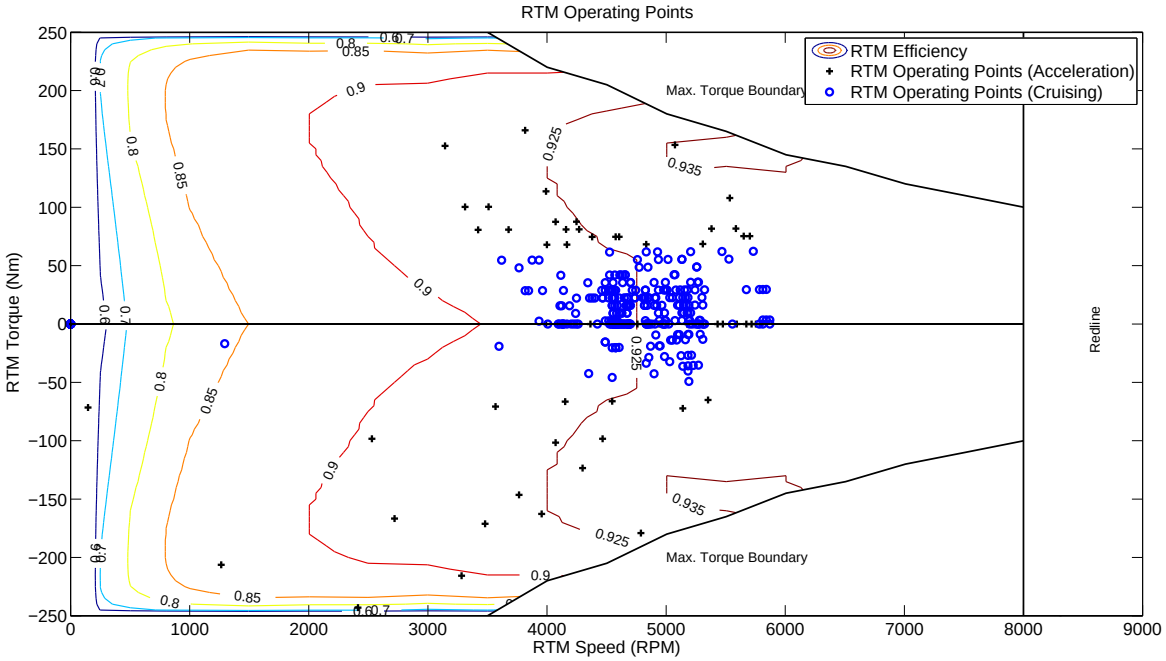

Figure 5.25: RTM Operating Points, US06 Highway, Series regime

The results of the 4-cycle analysis suggest that the series-regime efficiency gains of an optimally-run, decoupled ICE are not enough to offset the extra losses associated with charging through the BAS system, and that the direct-to-road ICE assistance provided in the parallel regime provides the best performance. The performance difference appears to be most significant in the more aggressive US06 driving cycles.

For the most aggressive US06 City cycle, series regime fuel consumption is significantly higher than parallel. This is because the algorithm's rejection of high-power ICE use means there is little opportunity to exploit optimal, decoupled ICE operation. Instead, the DP algorithm selects solutions that use fuel displacement through heavy RTM use as the best means to globally improve the fuel economy, using the vehicle's strong regenerative braking capability to operate primarily as an EV with occasional support from the ICE.

The less aggressive UDDS and HWFET cycles impose lower peak and average power demands, and RTM operating points are consequently located in regions of lower efficiency than for the more aggressive cycles. It therefore increases the cost to draw on ESS power relative to generating power from an optimally-run ICE, and the gap between series and parallel fuel consumption is narrowed. It should be noted that although the BAS efficiency map is based on component data provided by the manufacturer, the BAS-to-ICE coupling belt system is a custom unit designed and built in-house, and several assumptions have been made about its operating efficiency. Any modifications to that efficiency value could shift the series' fuel consumption to be equal to, or less than, that of the parallel regime.

A key implication of the DP algorithm's selection of heavy RTM use to reduce fuel consumption is that a limitation or reduction in regenerative braking will have a strong effect on the resulting fuel consumption - this idea is investigated in more detail in the next section.

5.2.3 Fuel Economy with Limited Regenerative Braking

The fuel economy results to this point have been determined assuming that the full regenerative braking ('regen') capabilities of the selected architecture can be utilized, and without consideration for the distribution of regenerative and mechanical braking efforts. In a real-world vehicle, the blending of these two braking forces aims to maximize the recovered energy, while ensuring vehicle stability and safe handling. Design of such a system requires extensive mechanical modification of the hydraulic

braking systems and poses a complex control problem that is well beyond the scope of this work. However, it is safe to assume that real-world regenerative braking levels will be much reduced compared to the rated power capabilities of the electric machines. Note that use of the BAS for regen is not considered, as the aim is to replenish the ESS using off-board energy.

To assess the effects of this real-world constraint on the vehicle's fuel consumption, the 4-cycle analysis was repeated, with regenerative braking restricted to 20% of the RTM's power capability, or -18kW. The value of 20% was selected as a compromise between allowing a reasonable level of regeneration while minimizing the effects on vehicle dynamics. The 4-cycle fuel economy results at 20% regenerative braking are shown in Table 5.3.

Drive Cycle	**FE(4-cycle)** $(Lge/100km)$	
	Series	Parallel
US06 City Segment	14.81	7.10
UDDS (505)	6.74	7.31
US06 Highway Segment	4.60	6.44
HWFET	3.89	5.36
EcoCAR2 4-cycle	**6.56**	**6.62**
Stock 2013 Malibu Eco	8.83	

Table 5.3: Fuel Economy Values with Regen Limited to 20%

The results for the US06 City cycle are shown in Figs. 5.26 and 5.27.

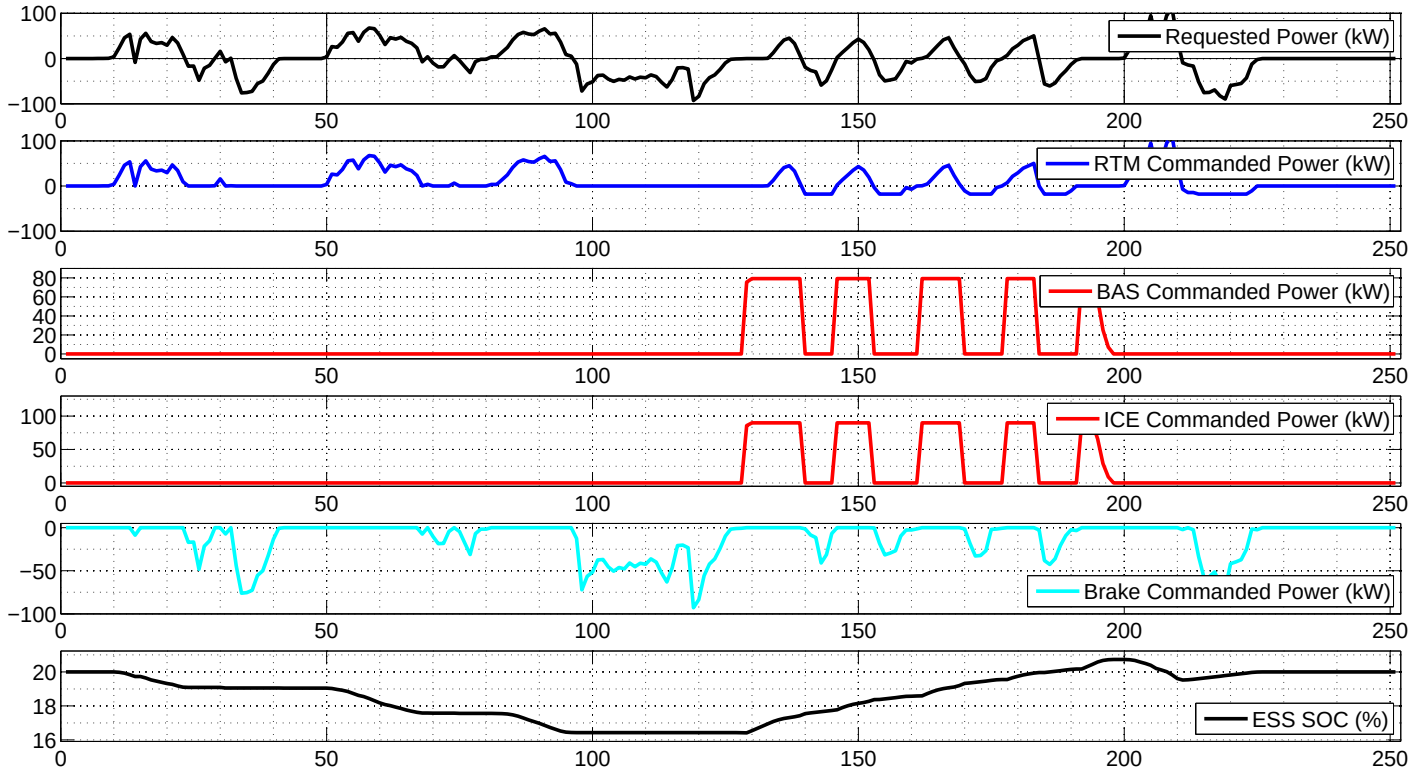

Figure 5.26: Power Outputs for US06 City, Series regime, 20% regen

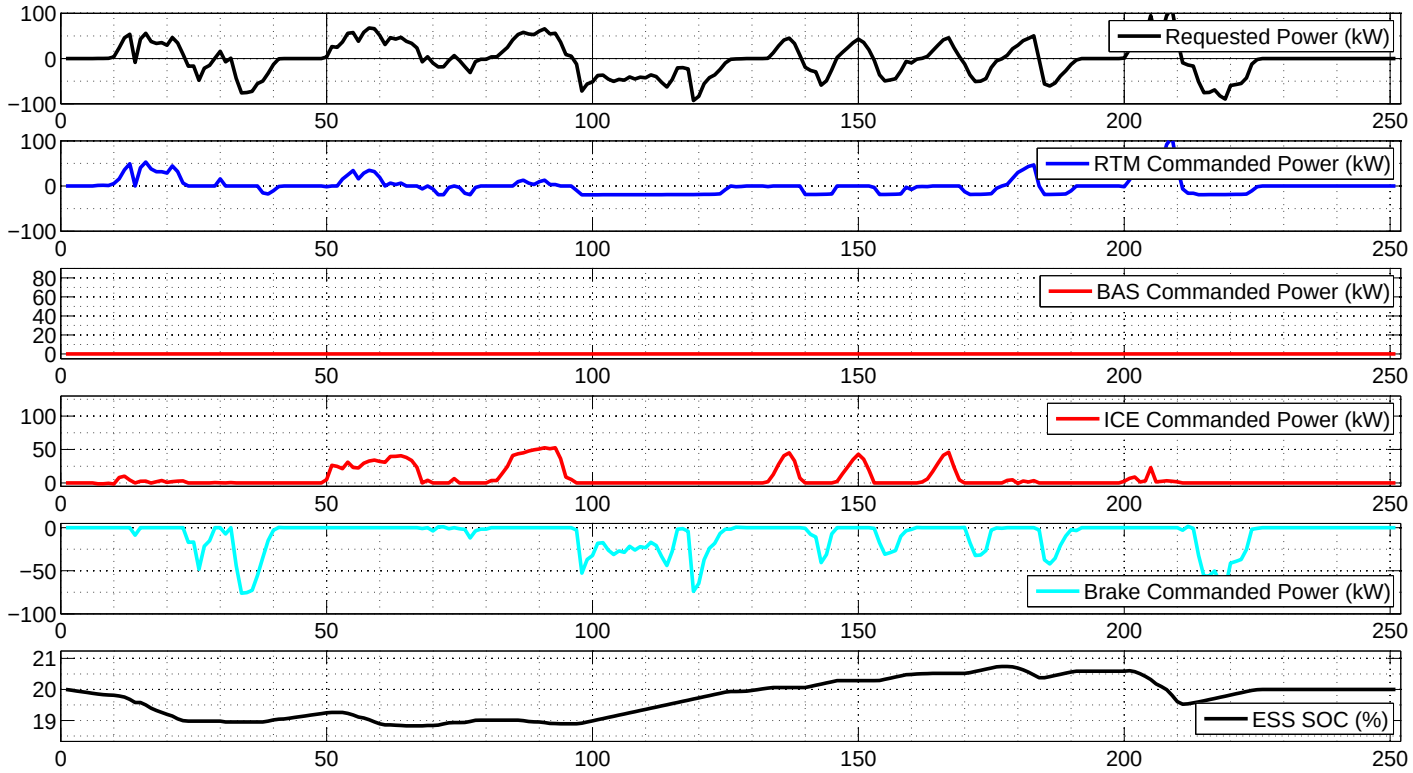

Figure 5.27: RTM Operating Points for US06 City, Parallel regime, 20% regen

Series powertrain operation is essentially similar to that of the full-regen case, but with more limited braking power from the RTM. This results in more charging events being required to restore SOC at the end of the cycle. When operating in the parallel regime, the ICE is used more heavily during accelerations, to reduce ESS discharge to a level that can be replenished through regen. The fuel consumption is higher over all the cycles, for all regimes, and the increase is especially pronounced on the US06 City cycle.

For the other three cycles, the series regime outperforms the parallel, offering reduced fuel consumption. This is due to the balance between electric drive capability and optimal ICE operation. In the full-regen scenario, it was far more expensive for the DP algorithm to choose solutions that relied more on the ICE, because of the ease of recovering spent ESS energy. With the regen capability restricted, the cost balance between consuming ESS energy and fuel energy shifts in favour of the optimally-run ICE-BAS system of the series regime.

Performance over the US06 City cycle does not follow this trend, because of the increased power requirement. In the parallel regime, the DP algorithm selects a path that minimizes divergence from the SOC setpoint, such that the limited regen can still recover enough energy to restore the ESS. In the series regime, the algorithm combines all BAS charging into shorter-duration, high-power events to maximize efficiency, which causes a corresponding deeper ESS discharge.

The almost 50% rise in fuel consumption observed with reduced regen highlights the importance of a well-designed and/or optimized regenerative braking strategy in the vehicle, and indeed in all hybrid vehicle architectures. On the other hand, even with the limited regen braking, the fuel economy improvement in either regime compared to the stock vehicle is still around 25%; this demonstrates that great potential for fuel savings can be found by implementing even mild or reduced-function hybrid powertrains.

5.3 DP Solution applied to Simulink Model

Although a simplified MATLAB model was used by the DP algorithm to determine optimal control policies, it is desirable to utilize a model more similar to those used for supervisory controller development. This way, the performance of a particular control strategy can be tested directly against the benchmark using the same model. Due to differing model structures, this is not straightforward, as described in Section 4.3.1.

In order to test the compatibility between the MATLAB-based DP model outputs and a higher-fidelity Simulink powertrain model, the control actions obtained from the DP algorithm were applied to the supervisory controller of the higher-fidelity Simulink model. The model used was a Simulink-only EcoCAR team model developed as a test platform for optimal control strategies.

5.3.1 Open-Loop Application of DP-Generated Torques

Initially, component torque requests were determined using the DP solution, which were applied as open-loop component torque requests within the Simulink vehicle model's supervisory controller. The DP Validation driving cycle shown in Fig. 5.1 was used, as the simplicity of the cycle allows for a more straightforward comparison of operation. The results are shown in Figs. 5.28 to 5.31. Note that DP component

power outputs are a direct result of the simulation, whereas Simulink component power outputs are calculated in post-processing from simulated component torque and speed values. Again, this is due to the backwards-vs-forwards simulating nature of the two models.

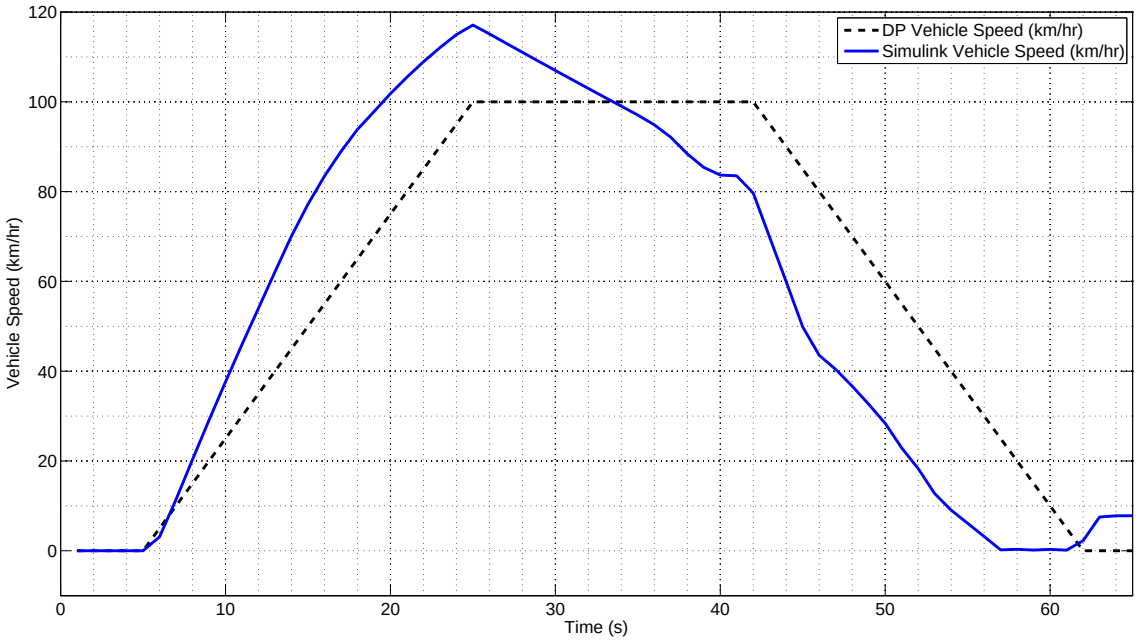

Figure 5.28: Vehicle speed comparison; DP vs Simulink, Parallel regime

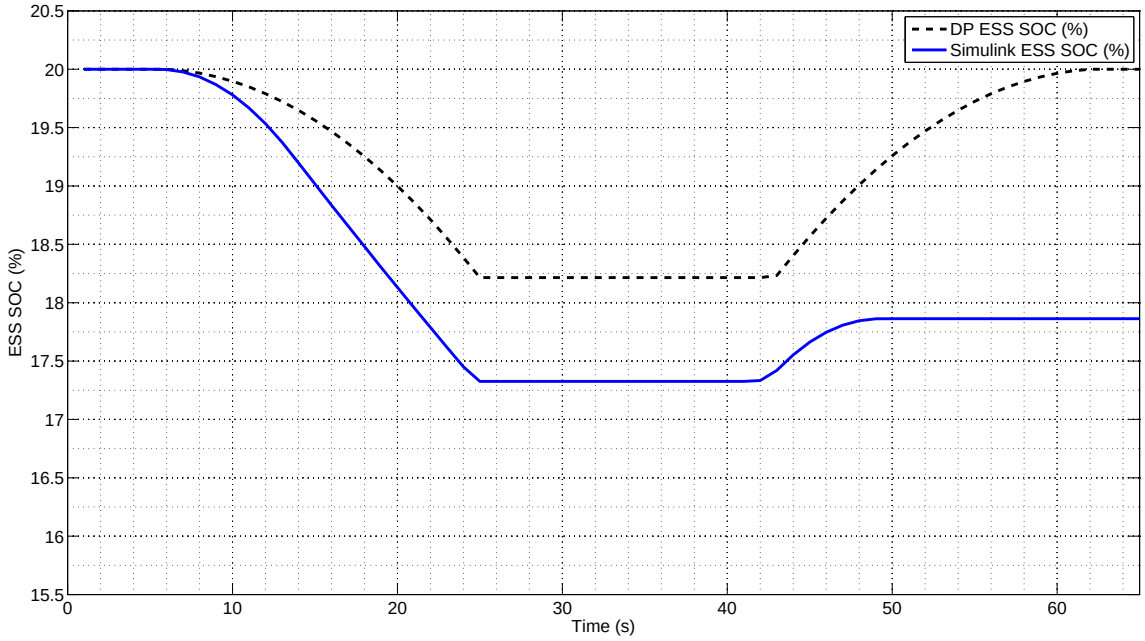

Figure 5.29: ESS SOC comparison; DP vs Simulink, Parallel regime

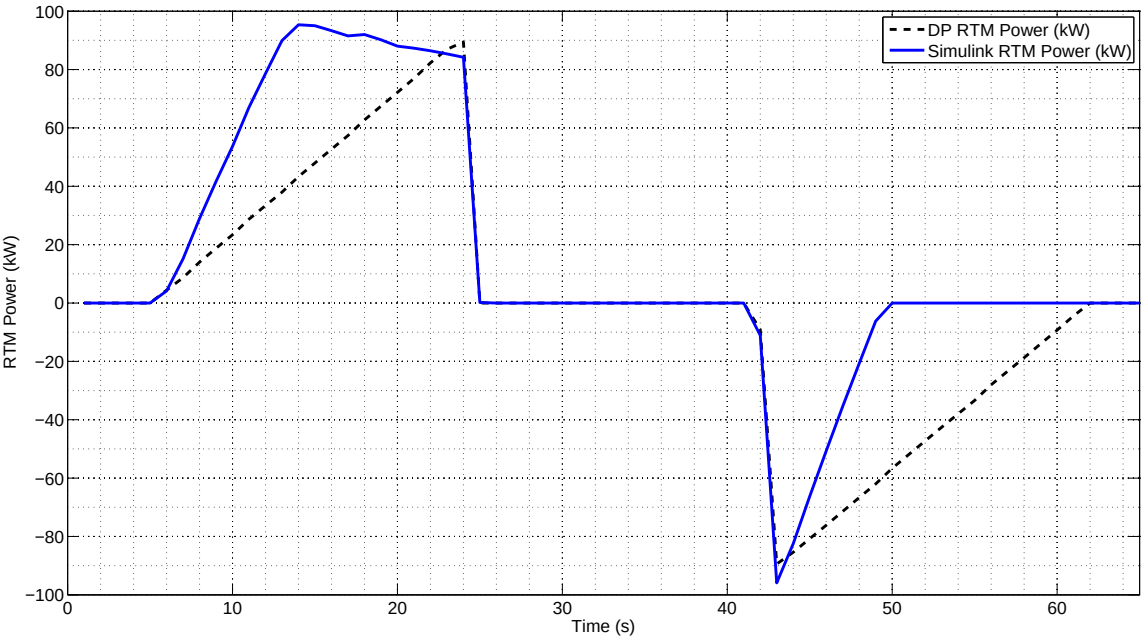

Figure 5.30: RTM Power comparison; DP vs Simulink, Parallel regime

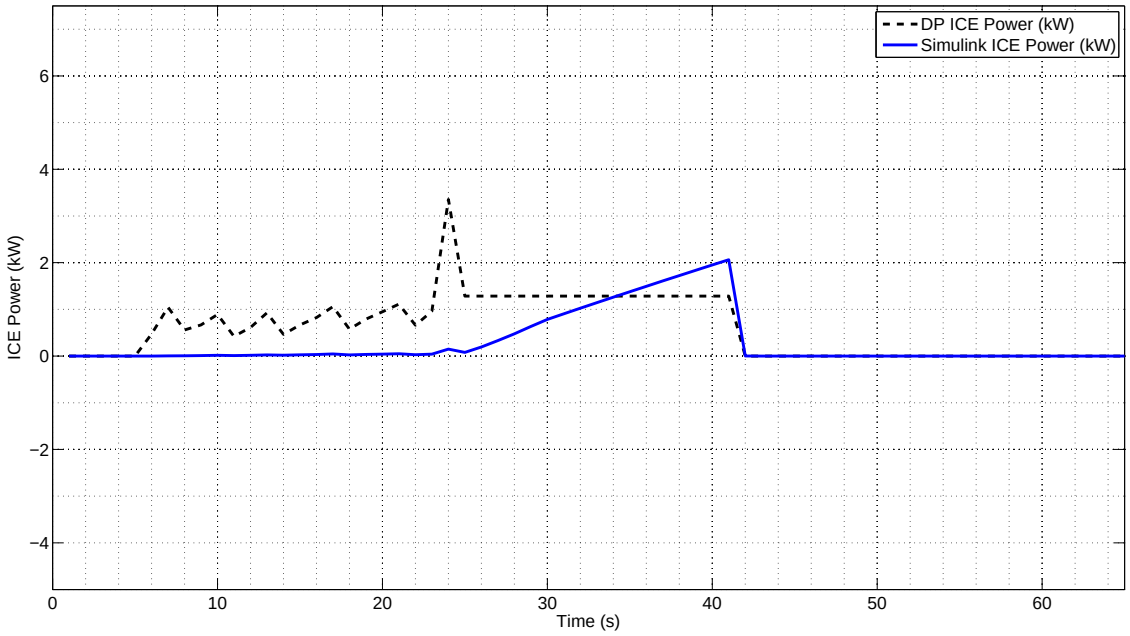

Figure 5.31: ICE Power comparison; DP vs Simulink, Parallel regime

Not surprisingly, the results diverge significantly. Using the DP-generated torque requests, the Simulink model accelerates more quickly than the DP model, and exceeds the desired 100km/hr cruising speed. During the cruising phase, the vehicle decelerates slowly; no RTM propulsive power is supplied, and looking at Fig. 5.31, the ICE does not supply any power until t=24, where it slowly ramps up but is still not significant. This was found to be caused by the particular implementation of the Simulink ICE and transmission models; the ICE cannot provide such low torque at the speeds it is operating it over this cycle.

During the deceleration phase, the vehicle's braking rate is higher than desired, and the regenerative braking captures less cumulative energy than in the DP model. This is due to the Simulink RTM model's more detailed torque limitations and asymmetrical efficiency curves, where regenerative torque is more limited than propulsive torque, as well as the reduced time spent with regenerative torque applied. Note also the Simulink model's return to an 8km/hr cruise after t=62. This is caused by creep torque, or parasitic drag from the idling ICE carrying through the transmission torque converter, which is accounted for in the Simulink model. The result of this component power discrepancy is a steeper ESS discharge during acceleration and a reduced recharge during braking, leading to non-CS behaviour, as can be seen in Fig.

5.29.

5.3.2 Closed-Loop (PI-Controlled) Application of Torques

To provide some mitigation for this, a PI controller was added in parallel with the DP-generated open-loop torque requests in the Simulink model supervisory controller. While the torque requests provided open-loop base setpoints, the PI controller closes the control loop and adds (or subtracts) torque based on an error signal, chosen to be the difference in Simulink and DP-sim vehicle speeds. PI tuning was performed empirically, with gains adjusted until acceleration and cruising speeds matched closely. The simulation results following this modification are shown in Figs. 5.32 to 5.35.

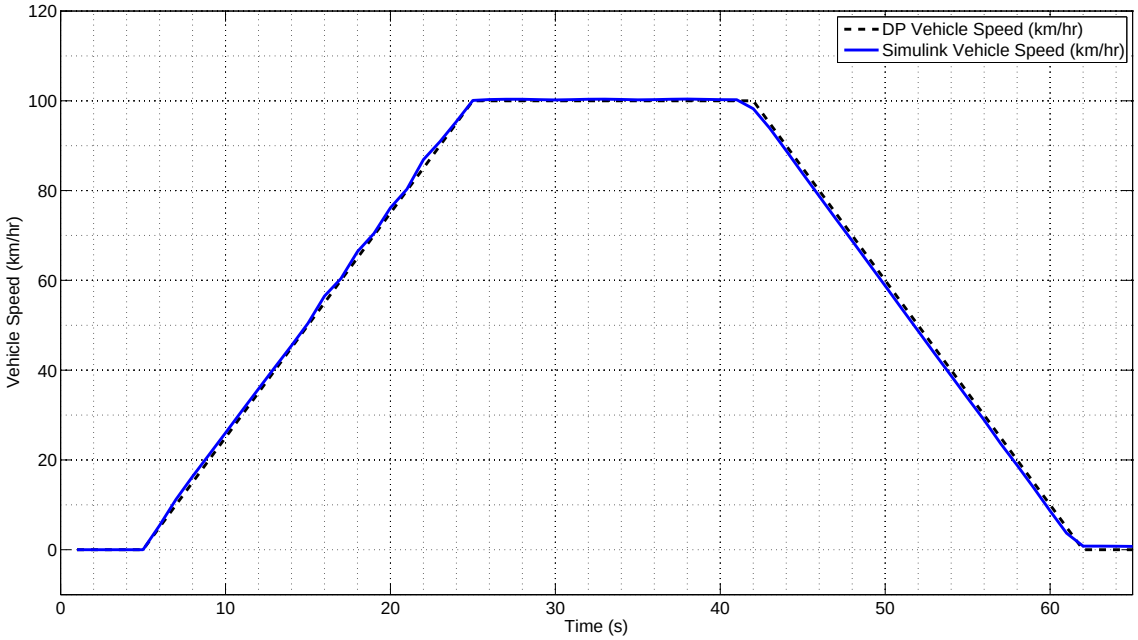

Figure 5.32: Vehicle speed comparison; DP vs Simulink, Parallel regime

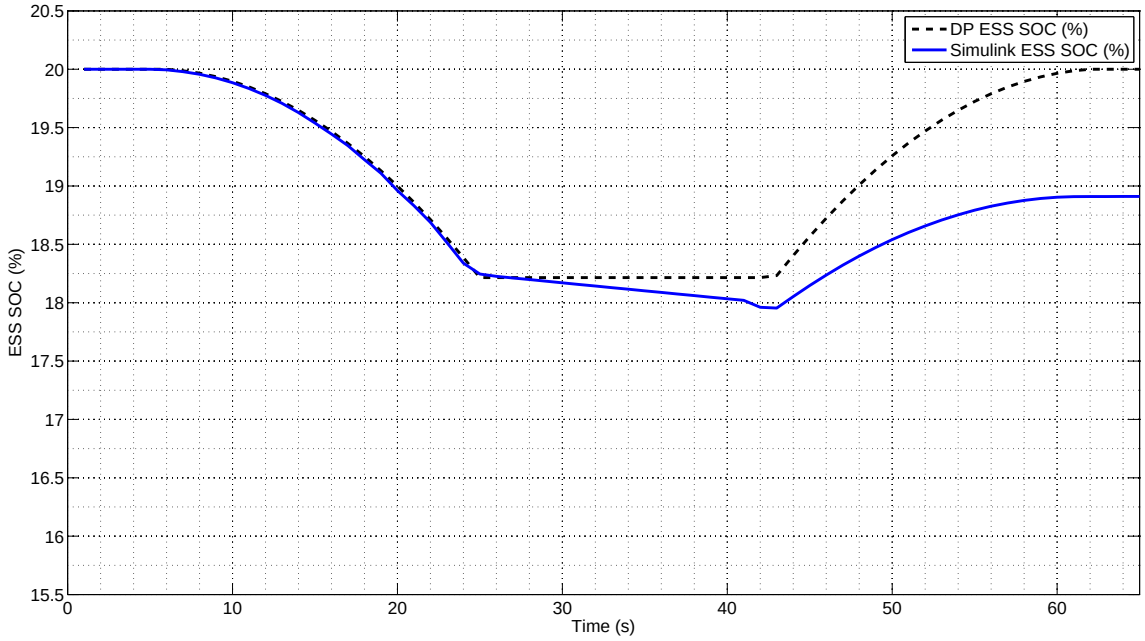

Figure 5.33: ESS SOC comparison; DP vs Simulink, Parallel regime

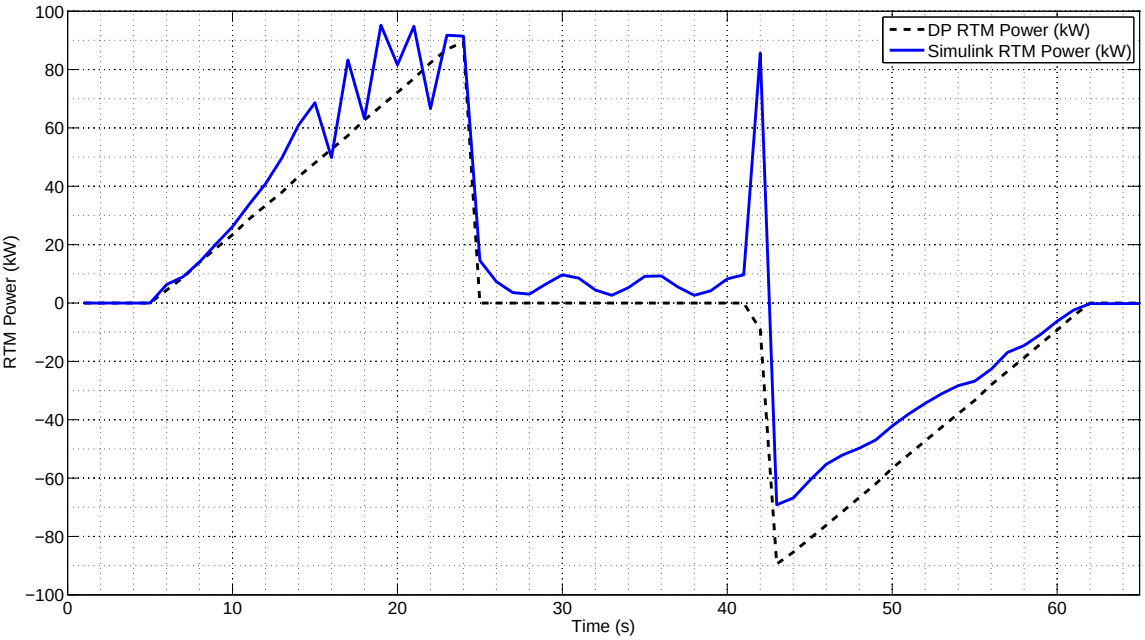

Figure 5.34: RTM Power comparison; DP vs Simulink, Parallel regime

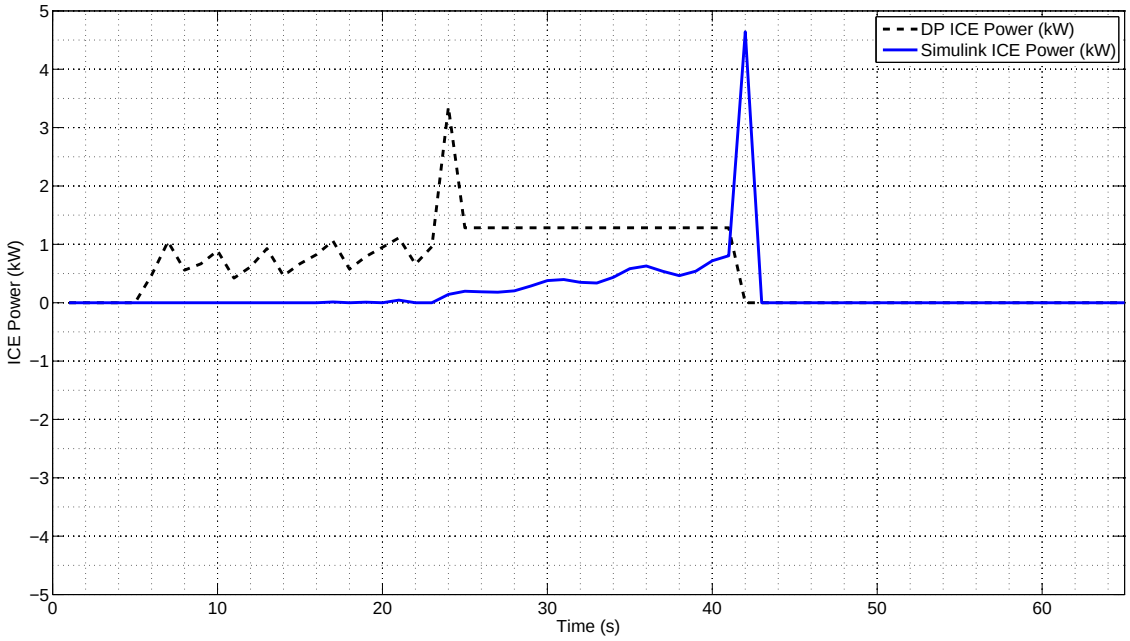

Figure 5.35: ICE Power comparison; DP vs Simulink, Parallel regime

With the assistance of the PI controller, the Simulink results more closely resemble the DP results. RTM power follows the general trend of the DP solution, with an additional oscillation caused by the PI loop. During acceleration, the PI controller adds additional RTM torque to match speed to the setpoint, resets when speed starts to exceed the desired value, and ramps up again; this process repeats several times. During the cruising phase, the PI controller causes RTM power to oscillate gently as cruising speed is maintained. Again, the ICE does not operate at the low torque levels requested, and only begins to ramp up during the cruising phase. The large spike in RTM and ICE powers at t=42s is an artifact of down-sampling from the 1ms Simulink time step to the 1s DP time step; several large, high-frequency oscillations exist in the Simulink data that are effectively averaged out but occasionally appear in the plots.

The ESS SOC effectively follows the DP curve through discharge (acceleration phase) and the initial segment of the cruising phase, before continuing to discharge. This is a result of PI-engagement of the RTM to propel the vehicle through the cruising phase. Finally, ESS regeneration is less during the braking phase, as RTM regen power is lower. This is again due to the implementation of the Simulink model, where RTM regenerative torque is more limited than propulsive torque.

Despite the influence of the PI controller, operation is still not charge-sustaining, meaning direct FE comparisons between the DP and Simulink results cannot be simply made. It is possible to address the charge imbalance by accounting for the additional fuel required to charge the ESS post-cycle, using an average value of fuel-to-electric conversion efficiency, as shown in Eq. 5.2.

$$\begin{aligned} FC_{SOCcorrected} &= \frac{m_{fuel} + m_{fuel_{charging}}}{d_{cycle}} * \rho_{E85} * 100 \\ \text{where } m_{fuel_{charging}} &= \frac{\frac{E_{consumed}}{0.296}}{LHV_{E85}} \end{aligned} \tag{5.2}$$

The efficiency value of 0.296 is derived from the fuel consumed by the BAS system during an 80kW charge. This yields a fuel consumption value of around 13L/100km for E85, or 9.55L/100km gasoline equivalent. Intuitively, this value seems unrealistically high for such a simple cycle, though it is in a similar range to the fuel consumption values obtained through the DP algorithm in the series regime on the US06 City cycle.

Although these results appear more promising, it must be remembered that the influence of the PI-controller changes the component torques and output powers from the DP-generated baseline. Thus, they can no longer be considered *optimal* control policies, and the benchmark results obtained with them are suspect.

A final comparison is given in Table 5.4 between the fuel consumption values for the Simulink model using different sets of control inputs and the original DP results. Performance when using the DP-generated, PI-modified optimal trajectories is compared against performance when using component torque commands generated by a practical rule-based controller.

Drive Cycle	**FE(4-cycle)** $(Lge/100km)$	
	Series	Parallel
US06 City (Simulink w/ DP)	26.43	16.41
US06 City (Simulink w/ rule-based)	14.33*	
US06 City (DP Result)	8.72	6.99
US06 Highway (Simulink w/ DP)	16.90	11.32
US06 Highway (Simulink w/ rule-based)	5.97*	
US06 Highway (DP Result)	4.31	3.89
Stock 2013 Malibu Eco	8.83	

Table 5.4: Fuel Consumption Comparison: DP and Rule-Based Control

* the rule-based controller does not distinctly define Series and Parallel operation, but rather selects the most appropriate component-level operation based on current conditions.

All Simulink-based table results have been compensated for varying degrees of end-of-cycle charge imbalance, by applying the post-cycle BAS charging method of Eq. 5.2. The majority of the fuel consumption in the 'Simulink w/ DP' simulations occurs as a post-cycle charge; as with the earlier example using the Validation cycle, true CS operation is not maintained during the simulations. The rule-based controller offers far better performance than that provided by the PI-modified DP control policies, but it still does not approach the DP predictions. Notwithstanding the underlying model differences between the DP and Simulink models used here, the DP results provide a hard lower bound on fuel consumption.

Additionally, the Simulink-based results underline the concept that an improperly-applied optimal control policy can result in greater fuel consumption than a well-tuned, non-optimal, rule-based strategy. This is similar to the conclusions of Waldner in [29].

5.3.3 DP-to-Simulink Recommendations

These results highlight the major differences between the low- and high-fidelity models. First, the higher-fidelity Simulink model captures details of the powertrain that are not considered in the MATLAB model. As it is a forward-looking model, component rotational speeds in the Simulink models are a function of applied torque, and once speeds and torques are known, component efficiencies and input powers can be calculated. Resulting vehicle speed is calculated from the force applied at the wheels. This is not the same for the MATLAB model used by the DP algorithm, where component powers are determined through the DP search and torques are calculated in post-processing. The Simulink model also operates at a different time step than the MATLAB model (1ms versus 1s). To use the DP-derived control policies in Simulink, they must be heavily interpolated to find the missing values, which can introduce errors.

It is apparent that it is not feasible to directly apply the DP-generated component operating parameters to the higher-fidelity Simulink model without substantially modifying the optimal control policies to fit the Simulink environment. Using the DP algorithm developed for this work, it is also not feasible to include the Simulink

model directly in the analysis to calculate state transition costs. This would incur a massive increase in program execution time, and would require the integration of other complex software components such as steady-state detection algorithms and Simulink/MATLAB interfacing. Any potential increase in accuracy from this approach is very much outweighed by these factors.

Despite these present limitations, the DP algorithm remains a useful tool to identify the upper efficiency boundary of potential hybrid architectures and provide some degree of comparison with future optimal control strategies. To improve capability, high-performance computing resources could be applied to the problem. These would enable direct use of a higher-fidelity model with the existing DP algorithm.

A viable alternative, beyond the scope of this work, would be to analyse the high-fidelity Simulink model under varying conditions to generate a metamodel, or a reduced-complexity empirical model, that is representative of Simulink model behaviour while being more compatible with the DP algorithm. The metamodel would require greatly reduced computational resources to run, but would provide state transition costs that more closely reflect the parameters of the higher-fidelity, forward-looking Simulink model. This is in line with the method used by Murgovski *et al* in [45].

5.4 Summary of Results

This chapter presented the analysis results of the DP algorithm. The DP algorithm was first validated using a simplified testing cycle. The validation showed that the DP algorithm correctly identified optimal control actions while accounting for active constraints. The optimal control actions, and resulting fuel consumption values, were then computed for both parallel and series regimes over four standard driving cycles: US06 City and Highway, the UDDS, and HWFET cycles. Combined 4-Cycle fuel consumption values were 4.91L/100km for the series regime, and 4.74L/100km for the parallel regime, compared to the stock 2013 Malibu's 8.83L/100km.

Globally optimal operation involved the vehicle being propelled primarily through the RTM, with the ICE providing limited assistance after the half-way point of each cycle. Operation in the parallel regime appears most effective during more aggressive driving, where the ICE can work to provide propulsive assistance during hard accelerations. Operation in the series regime appears effective under gentler driving conditions, where limited opportunities exist for heavy use of regenerative braking to

replenish the ESS.

Vehicle operation was next assessed using a more limited and realistic degree of regenerative braking, by repeating the drive cycle simulations with total regen capability restricted to 20% of the RTM's maximum capability. Combined 4-Cycle fuel consumption increased to 6.56L/100km for the series regime, and 6.62L/100km for the parallel regime. Operation is similar to the full-regen case, but with an expected increase to ICE use. The series regime outperforms the parallel when regen braking is limited.

Finally, the DP-generated control policies were applied to a higher-fidelity Simulink model to test compatibility between the two models. The initial open-loop application of the control policies served highlight differences between the MATLAB and Simulink models, as the Simulink vehicle speed and ESS SOC diverged significantly from the DP solution. To mitigate this, a PI-controller was added to increase or decrease the torque provided by the DP control policies. Although this brought the Simulink vehicle's speed and SOC closer to the DP setpoints, the component torque requests were substantially modified from those of the optimal control policies.

Therefore, it is not feasible to directly apply the DP-generated component operating parameters to the higher-fidelity Simulink model in their current forms. It is also infeasible to include the Simulink model directly in the analysis to find state transition costs, due to the increased computational burden. Viable methods to overcome this problem are discussed in Section 6.3.

Chapter 6

Summary and Future Work

6.1 Summary of Work

Advanced hybrid vehicle powertrains, such as the Series-Parallel Multiple-Regime PHEV being implemented as UVic's EcoCAR2 competition entry, offer opportunities for fuel economy improvement through the use of optimization-based control strategies. Due to the multiple powertrain components and operating regimes, selecting component operating profiles to meet driver power demand is a multi-dimensional control problem, with powertrain performance and efficiency being strongly dependent on control methodology.

To quantify the performance improvement offered by a particular control strategy, it is essential to obtain performance benchmark values for the powertrain architecture that are independent of control strategy. These benchmarks provide a maximum performance target for comparison, provide data that can aid in the tuning of different control strategies, and inform the derivation of control rules and/or optimal values of specific control parameters, such as the fuel-electric equivalence factor used in the ECMS control strategy.

This work is motivated by the need to perform the above tasks for the advanced Series-Parallel Multiple-Regime PHEV architecture. The primary goal was the evaluation of the architecture from an optimal performance standpoint using global optimization, to gain an understanding of the ideal operation of this unique architecture. The second objective was to examine the differences between high- and low-fidelity powertrain models, and examine the effects and feasibility of incorporating a high-fidelity powertrain model into the global optimization process.

A simplified high-level MATLAB model was developed for the vehicle operating in series and parallel regimes in charge-sustaining mode. Powertrain subsystems were represented with empirical models, derived from measured component data. A global optimization algorithm was created, using a backwards-induction variation of Dynamic Programming, that performed a discrete-time analysis of the MATLAB model over a given driving cycle. The analysis was repeated at a range of different time and state space resolutions, to both assess the effects on the algorithm and to identify a good compromise between solution quality and reasonable execution time. A simple validation driving cycle was developed and used to demonstrate that the algorithm was correctly finding the optimal solution.

The DP algorithm was used to evaluate the operation and fuel consumption of the vehicle over a set of four standard driving cycles. Combined 4-Cycle fuel consumption values were 4.91L/100km for the series regime, and 4.74L/100km for the parallel regime, compared to the stock 2013 Malibu's 8.83L/100km. Limiting regenerative braking capabilities increased fuel consumption to combined 4-Cycle values of 6.56L/100km for the series regime, and 6.62L/100km for the parallel regime.

The parallel regime provided slightly more efficient operation, especially during periods of more aggressive driving. However, the series regime provided more efficient operation only under gentler driving conditions, and especially where opportunities for regenerative braking are more limited. In general, the decoupled ICE operation does not appear to provide sufficient efficiency improvement to offset the series regime's inherent additional energy conversion loss.

Performing simulations with the higher-fidelity Simulink model using control inputs generated from the DP analysis was found to be an ineffective approach, as the significant difference in model structures results in undesirable vehicle operation. While real-time corrections can be performed to the control inputs to enforce comparable vehicle operation, these efforts nullify the optimality of the DP-generated control policies.

The algorithm and model developed in this work were found to be a useful proof-of-concept for the benchmark analysis using DP, indicating that dynamic programming can identify globally optimal benchmark solutions. However, a different approach will be required to address the model disparities and allow direct analysis of higher-fidelity Simulink models with the DP algorithm.

6.2 Contributions

The core contribution of this work is the development of a proof-of-concept powertrain analysis algorithm based on dynamic programming. The variety of complex hybrid architectures means that a general-purpose tool is not available for this purpose. This work has verified the use of dynamic programming as a viable technique for powertrain benchmark analysis, and identified issues related to model fidelity and formulation that must be resolved to improve the accuracy of the benchmark results.

6.3 Future Work

This work provides an analysis of the Series-Parallel Multiple-Regime PHEV powertrain at a high-level, trading model fidelity for reduced computation time and simplicity of analysis. The initial combination of DP-derived control policies and a higher-fidelity Simulink model demonstrated the lack of direct compatibility between the two model levels, due to fundamental differences in modelling methodology.

In light of this, the next critical step is to incorporate a higher-fidelity model into the global optimization process, to ensure cross-compatibility of results. A very promising approach to achieve this would be through the use of metamodelling methods; either a black-box model would be defined through direct analysis of the high-fidelity Simulink model, or a series of lower-order polynomial models representing powertrain systems would be parameterized through curve-fitting. These would provide a much better approximation of powertrain behaviour than the simplified MATLAB model, while not requiring the same computational effort to analyse as the full high-fidelity Simulink model. The method presented in [45] is a good example of this, and should be further explored.

In either case, the existing DP cost function would need to be modified to utilize the metamodel to calculate fuel use, while the DP search functions could remain the same. Alternatively, high-performance computing resources could be leveraged to allow the direct inclusion of a high-fidelity model into the DP search.

Once the model compatibility issues are resolved, the analyses of this work should be repeated and results verified. From there, the DP control policies could be tested in-vehicle, to study the real-world effectiveness and implications to driveability of optimal operation. The DP outputs could also be used to formulate a rule-based control strategy for the vehicle, or to tune an ECMS or similar advanced real-time

control strategy.

The use of the BAS during parallel and ICE-only operation should be analysed to determine its effectiveness. In theory, the BAS can be used to add an additional torque load to the ICE, to shift its operating points to higher-efficiency regions. This was not considered in this work, as it adds an additional degree of freedom to the system that greatly increases simulation complexity and time. This type of operation may offer performance benefits, specifically as a method to maintain CS operation when limited regenerative braking is available. It would be desirable to incorporate this regime into the analysis and assess the effectiveness of these strategies.

Bibliography

[1] Natural Resources Canada: Review of Issues Affecting the Price of Crude Oil @ONLINE. Available: http://www.nrcan.gc.ca/energy/publications/sources/crude/issues-prices/1507, October 2010.

[2] National Energy Board. Canadas Energy Future: Energy Supply and Demand Projections to 2035 - Energy market Assessment @ONLINE. Available: http://www.neb-one.gc.ca/clf-nsi/rnrgynfmtn/nrgyrprt/nrgyftr/2011/nrgsppldmndprjctn2035-eng.html#s3, November 2011.

[3] Statistics Canada: Greenhouse Gas Emissions from Private Vehicles in Canada, 1990 to 2007 @ONLINE. Available: http://www.statcan.gc.ca/pub/16-001-m/16-001-m2010012-eng.htm, May 2010.

[4] Intergovernmental Panel on Climate Change: Climate Change 2013: The Physical Science Basis @ONLINE. Available: http://www.ipcc.ch/report/ar5/wg1/, May 2013.

[5] A. Emadi, K. Rajashekara, S. Williamson, and S. Lukic. Topological, Overview of Hybrid Electric and Fuel Cell Vehicular Power System Architectures and Configurations. *IEEE Trans. on Vehicular Tech.*, Volume 54, May 2005.

[6] D. Robinette and M. Powell. Optimizing 12-Volt Start-Stop for Conventional Powertrains. *SAE International Journal of Engines*, December 2011.

[7] Summary of Travel Trends: 2009 National Household Travel Survey @ONLINE. Available: http://nhts.ornl.gov/2009/pub/stt.pdf, June 2011.

[8] L. Brooke. *Chevrolet Volt: Development Story of the Pioneering Electrified Vehicle*. SAE International Vehicle Electrification Series, 2011.

[9] Capstone Turbine Solutions: HEV @ONLINE, March 2009.

[10] US Environmental Protection Agency - Find and Compare Cars @ONLINE, July 2013.

[11] K. Aoki et al. Development of Integrated Motor Assist Hybrid System: Development of the Insight, a Personal Hybrid Coupe. *SAE Technical Papers Series*, June 2000.

[12] J. Wishart, Y. Zhou, and Z. Dong. Review of Multi-Regime Hybrid Vehicle Powertrain Architecture. *Int. J. Hybrid and Electric Vehicles*, Volume 1, 2007.

[13] L. Zhou, J. Wise, S. Bowman, C. Crawford, and Z. Dong. Design, Modelling, and Hardware Implementation of a Next Generation Extended-Range Electric Vehicle. *IEEE Control Systems Magazine*, Volume 4, December 2010.

[14] A. Khaligh and L. Zhihao. Battery, Ultracapacitor, Fuel Cell, and Hybrid Energy Storage Systems for Electric, Hybrid Electric, Fuel Cell, and Plug-In Hybrid Electric Vehicles: State of the Art. *IEEE Trans. on Vehicular Tech.*, Volume 59 No. 6, July 2010.

[15] X. Chen, W. Shen, T.T.Vo, Z. Cao, and A. Kapoor. An Overview of Lithium-ion Batteries for Electric Vehicles. *IPEC Conf. on Power and Energy*, December 2012.

[16] A123 Systems Inc. Nanophosphate Basics: An Overview of the Structure, Properties and Benefits of A123 Systems Proprietary Lithium Ion Battery Technology. *White Paper: Nanophosphate Lithium Ion Battery Tech.*, July 2010.

[17] M. Jiang, J. Salmon, and A.M. Knight. Design of a Permanent Magnet Synchronous Machine for a Flywheel Energy Storage System within a Hybrid Electric Vehicle. *Electric Machines and Drives Conf.*, May 2009.

[18] X. Lu, K.L.V. Iyer, K. Mukherjee, and N.C. Kar. Battery, Ultracapacitor, Fuel Cell, and Hybrid Energy Storage Systems for Electric, Hybrid Electric, Fuel Cell, and Plug-In Hybrid Electric Vehicles: State of the Art. *IEEE Trans. on Vehicular Tech.*, Volume 59 No. 6, July 2010.

[19] Kinetic Energy Recovery Systems (KERS) @ONLINE. Available: http://www.formula1.com/inside_f1/understanding_the_sport/8763.html, March 2009.

[20] A. Emadi, Y.J.Lee, and K. Rajashekara. Power Electronics and Motor Drives in Electric, Hybrid Electric, and Plug-In Hybrid Electric Vehicles. *IEEE Trans. on Industrial Electronics*, Volume 55 No. 6, June 2008.

[21] T. M. Jahns and V. Blasko. Recent Advances in Power Electronics Technology for Industrial and Traction Machine Drives. *Proc. IEEE*, Volume 89, August 2002.

[22] Utility Factor Definitions for Plug-In Hybrid Electric Vehicles Using 2001 U.S. DOT National Household Travel Survey Data @ONLINE, March 2009.

[23] F. R. Salmasi. Control Strategies for Hybrid Electric Vehicles: Evolution, Classification, Comparison, and Future trends. *IEEE Trans. on Vehicular Tech.*, Volume 56, September 2007.

[24] A. Phillips, M. Jankovic, and K. E. Bailey. Vehicle System Controller Design for a Hybrid Electric Vehicle. *Proc. 2000 IEEE Int. Conf. on Control Applications*, December 2000.

[25] S. G. Wirasingha and A. Emadi. Classification and Review of Control Strategies for Plug-In Hybrid Electric Vehicles. *IEEE Trans. on Vehicular Tech.*, Volume 60, January 2011.

[26] N.J. Schouten, M.A. Salman, and N.A. Kheir. Fuzzy Logic Control for Parallel Hybrid Vehicles. *IEEE Trans. on Control Sys. Tech.*, Volume 10 No. 3, May 2002.

[27] S.G. Li, S.M. Sharkh, F.C. Walsh, and C.N. Zhang. Energy and Battery Management of a Plug-In Series Hybrid Electric Vehicle Using Fuzzy Logic. *IEEE Trans. on Vehicular Tech.*, Volume 60 No. 8, October 2011.

[28] T.A. Anderson, J.M. Barkman, and C. Mi. Design and Optimization of a Fuzzy-Rule Based Hybrid Electric Vehicle Controller. *IEEE Veh. Power and Propulsion Conf.*, September 2008.

[29] J. Waldner. Development of a 2-Mode AWD E-REV Powertrain and Real-Time Control System. M.a.sc thesis, University of Victoria, Victoria, British Columbia, Canada, 2011.

[30] G. Paginelli et al. Equivalent Consumption Minimization Strategy For Parallel Hybrid Powertrains. *IEEE Vehicular Technology Conference*, December 2002.

[31] C. Musardo, G. Rizzoni, and B. Staccia. A-ECMS: An Adaptive Algorithm for Hybrid Electric Vehicle Energy Management. *IEEE Conf. on Decision and Control*, December 2005.

[32] E.D. Tate and P.S. Boyd. Finding Ultimate Limits of Performance for Hybrid Electric Vehicles. *SAE Future Transportation Technology Conference*, December 2000.

[33] A. Piccolo et al. Optimisation of Energy Flow Management in Hybrid Electric Vehicles via Genetic Algorithms. *IEEE/ASME Int. Conf. on Advanced Intelligent Mechatronics Proceedings*, December 2001.

[34] M. J. Gielniak and Z. J. Shen. Power Management Strategy Based on Game Theory for Fuel Cell Hybrid Electric Vehicles. *IEEE Vehicular Technology Conference*, December 2004.

[35] H. Banvait, X. Lin, S. Anwar, and Y. Chen. Plug-In Hybrid Electric Vehicle Energy Management System using Particle Swarm Optimization. *World Electric Vehicle Journal*, Volume 3, December 2009.

[36] G. Paginelli et al. Simulation and Assessment of Power Control Strategies for a Parallel Hybrid Car. *Proc Instn Mech Engrs*, Volume 214 Part D, December 2000.

[37] S. Delprat, J. Lauber, T. M. Guerra, and J. Rimaux. Control of a Parallel Hybrid Powertrain: Optimal Control. *IEEE Trans. on Vehicular Tech.*, Volume 53 No. 3, December 2004.

[38] J. Si et al. Approximate Dynamic Programming for Continuous State and Control Problems. *17th Mediterranean Conf. on Control and Automation*, Volume 16 No. 6, June 2009.

[39] R. Wang and S.M. Lukic. Dynamic Programming Technique in Hybrid Electric Vehicle Optimization. *IEEE Int. Electric Vehicle Conf.*, March 2012.

[40] H.I. Dokuyucu and M. Cakmakci. Concurrent Design of Energy Management and Vehicle Stability Algorithms for a Parallel Hybrid Vehicle using Dynamic Programming. *American Control Conf. 2012*, June 2012.

[41] R.M. Patil, Z. Filipi, and H.K. Fathy. Comparison of Supervisory Control Strategies for Series Plug-In Hybrid Electric Vehicle Powertrains Through Dynamic Programming. *IEEE Trans. on Control Systems Tech.*, Volume PP No. 99, May 2013.

[42] J. Liu and H. Peng. Modeling and Control of a Power-Split Hybrid Vehicle. *IEEE Trans. on Control Systems Tech.*, Volume 16 No. 6, November 2008.

[43] G. Gary Wang and S. Shan. Review of Metamodeling Techniques in Support of Engineering Design Optimization. *Journal of Mechanical Design*, Volume 127, November 2007.

[44] N. Murgovski, J. Sjöberg, and J. Fredriksson. A Methodology and a Tool for Evaluating Hybrid Electric Powertrain Configurations. *Int. J. of Electric and Hybrid Vehicles*, Volume 3 No. 3, November 2011.

[45] Nikolce Murgovski. Optimal Powertrain Dimensioning and Potential Assessment of Hybrid Electric Vehicles. Ph.D. Thesis, Chalmers University of Technology, Göteborg, Sweden, 2012.

[46] R.E. Bellman and E.S. Lee. History and Development of Dynamic Programming. *IEEE Control Systems Magazine*, Volume 4, November 1984.

[47] D. P. Bertsekas. *Dynamic Programming and Optimal Control*. Athena Scientific, 2005.

[48] R.E. Bellman and S.E. Dreyfus. *Applied Dynamic Programming*. Princeton University Press, 1962.

[49] Argonne National Labs : Model-Based Design Definition Of Terms @ONLINE, March 2009.

[50] D. Gao, C. Mi, and A. Emadi. Modeling and Simulation of Electric and Hybrid Vehicles. *Proc. IEEE*, Volume 95, No. 4, April 2007.

[51] S. Halbach, P. Sharer, S. Pagerit, C. Folkerts, and A. Rousseau. Model Architecture, Methods, and Interfaces for Efficient Math-Based Design and Simulation of Automotive Control System. *SAE World Congress*, April 2010.

[52] The MathWorks in the Automotive Industry @ONLINE. Available: http:// www.terasoft.com.tw/product/doc/auto.pdf, April 2008.

[53] ASM Gasoline Engine Basic Simulation Package @ONLINE. Available: http: //www.dspace.com/en/inc/home/products/sw/automotive_simulation_models/ produkte_asm/asm_engine_models/asm_gasoline_engine_basic_sim.cfm, April 2008.

[54] V. Ramadesigan, P. Northropa, S. Dea, S. Santhanagopalanb, R. Braatzc, and V. Subramanian. Modeling and Simulation of Lithium-Ion Batteries from a Systems Engineering Perspective. *J. Electrochem. Soc.*, Volume 159, August 2012.