University of Bath

UNIVERSITY OF BATH

PHD

Optimal Control of Hybrid Electric Vehicles for Real-World Driving Patterns

Vagg, Christopher

Award date:
2015

Awarding institution:
University of Bath

Link to publication

Optimal Control of Hybrid Electric Vehicles for Real-World Driving Patterns

Christopher Vagg

A thesis submitted for the degree of Doctor of Philosophy

University of Bath

Department of Mechanical Engineering

August 2014

Christopher Vagg

This thesis is dedicated to my incredible wife, Lara.

Table of Contents

Acknowledgements

Abstract

Optimal control of energy flows in a Hybrid Electric Vehicle (HEV) is crucial to maximising the benefits of hybridisation. The problem is complex because the optimal solution depends on future power demands, which are often unknown. Stochastic Dynamic Programming (SDP) is among the most advanced control optimisation algorithms proposed and incorporates a stochastic representation of the future. The potential of a fully developed SDP controller has not yet been demonstrated on a real vehicle; this work presents what is believed to be the most concerted and complete attempt to do so.

In characterising typical driving patterns of the target vehicles this work included the development and trial of an eco-driving driver assistance system; this aims to reduce fuel consumption by encouraging reduced rates of acceleration and efficient use of the gears via visual and audible feedback. Field trials were undertaken using 15 light commercial vehicles over four weeks covering a total of 39,300 km. Average fuel savings of 7.6% and up to 12% were demonstrated. Data from the trials were used to assess the degree to which various legislative test cycles represent the vehicles' real-world use and the LA92 cycle was found to be the closest statistical match.

Various practical considerations in SDP controller development are addressed such as the choice of discount factor and how charge sustaining characteristics of the policy can be examined and adjusted. These contributions are collated into a method for robust implementation of the SDP algorithm.

Most reported HEV controllers neglect the significant complications resulting from extensive use of the electrical powertrain at high power, such as increased heat generation and battery stress. In this work a novel cost function incorporates the square of battery C-rate as an indicator of electrical powertrain stress, with the aim of lessening the affliction of real-world concerns such as temperatures and battery health. Controllers were tested in simulation and then implemented on a test vehicle; the challenges encountered in doing so are discussed. Testing was performed on a chassis dynamometer using the LA92 test cycle and the novel cost function was found to enable the SDP algorithm to reduce electrical powertrain stress by 13% without sacrificing any fuel savings, which is likely to be beneficial to battery health.

Significant Contributions

The work presented in this thesis is believed to contain the following novelty and significant contributions to the existing body of knowledge in the field:

1. Design of an eco-driving driver assistance system for the light commercial vehicle (LCV) market, and the demonstration of 7.6% fuel savings.
2. A comparison of real-world LCV driving data to various legislative test cycles for the selection of a stochastically representative drive cycle.
3. A proposed methodology for the robust implementation of the SDP algorithm in its application to HEV control optimisation.
4. Use of probabilistic charge sustenance plots to allow the nominal charge sustaining behaviour of a control strategy to be examined.
5. Incorporation of the square of battery C-rate into the optimisation cost function, allowing the trade-off between fuel saving and electrical powertrain stress to be examined. Demonstration of this non-linear trade-off in a simulation environment.
6. Design of a SDP controller based on real world driving data, implementation on a real vehicle and testing on a chassis dynamometer over a representative drive cycle. Demonstration of a 13% reduction in electrical powertrain stress without compromising fuel savings.
7. Adaptation of the dynamometer-suitable SDP controller for real-world use by transformation of the state space, and demonstration of its operation on public roads.

List of Acronyms

BMS	Battery Management System
CoV	Coefficient of Variance
DAS	Driver Advisory System
DP	Dynamic Programming
ECMS	Equivalent Consumption Minimisation Strategy
ECU	Engine Control Unit
EM	Electric Machine
ESS	Energy Storage System
EUDC	Extra-Urban Driving Cycle
GSI	Gear Shift Indicator
HEV	Hybrid Electric Vehicle
ICE	Internal Combustion Engine
IPS	Inertial Power Surrogate
KERS	Kinetic Energy Recovery System
LCV	Light Commercial Vehicle
NEDC	New European Driving Cycle
OBD	On-Board Diagnostics
OEM	Original Equipment Manufacturer
RPA	Relative Positive Acceleration
SAFD	Speed-Acceleration Frequency Distribution
SDP	Stochastic Dynamic Programming
SEI	Solid-Electrolyte Interface
SOC	State Of Charge
SOH	State Of Health

Chapter 1   Introduction

This chapter offers an introduction to the subject area of hybrid electric vehicles, motivations for their adoption and an overview of the Ashwoods Automotive retrofit conversion system used in the work. The scope of the research is described and the author's individual contributions are distinguished from those of others with respect to the various work packages completed in the project as a whole.

1.1 Overview

Awareness of environmental and sustainability issues has grown considerably in recent decades and as such there is an ever growing pressure on motor vehicle manufacturers to produce vehicles that are more energy efficient, recyclable, and less polluting. Coupled with this, the rising cost of crude oil and taxation incentives for 'clean' transportation add financial motivation to the shift toward environmentally friendly transportation. Aside from any personal convictions toward the ethos of sustainability both individuals and companies, now more than ever, are encouraged to buy in to low carbon vehicles for at least two motives: to make a statement concerning the individual or company's ethical values by appearing 'green', and to make financial savings either through reduced fuel bills or taxation incentives.

Whilst the thermal efficiency of the Internal Combustion Engine (ICE), by which road vehicles have been almost exclusively powered, has increased significantly in the previous two decades, such engines based on the principle of a 'heat pump' are ultimately limited by the Carnot efficiency. This defines a theoretical maximum efficiency dependent on the temperature rise of combustion gasses. As such, efficiency gains are becoming harder to find and manufacturers are designing increasingly elaborate engines (incorporating complex electronic fuel injection strategies, exhaust energy recovery, exhaust gas recirculation and variable valve timing, for example) in the pursuit of ever smaller improvements. Taylor [1] predicts that improvements in ICE efficiency during the present decade will see vehicle fuel consumption fall by 6-15%, whilst integrating electric motors to hybridise vehicle powertrains and to recover kinetic energy could increase this figure to 21-28%.

Hybridisation of vehicle powertrains can take many forms, but fundamentally seeks to deliver tractive power from more than one energy store. Usually one of the energy converters may be operated bi-directionally and so the vehicle's fuel economy is improved by its ability to recover kinetic energy during periods of braking or coasting and to convert this back into stored energy for later use. Fuel economy may be further improved by the ability to trade energy between the two (or more) storage systems, allowing opportunities for the energy converters to be operated near their points of maximum efficiency, even if this results in an inequality between their combined power generation and the vehicle's instantaneous power requirement.

Furthermore, some hybrid vehicles allow for one of the energy storage systems to be charged off-line, for example from the electricity grid, which may be able to deliver less carbon-intensive and/or cheaper energy.

In most cases the primary energy store remains a petroleum-based fuel, with a traditional ICE used to convert this into tractive work. The secondary energy store and conversion technology tend to vary considerably more, both in size and in usage; technologies commonly suggested for this application include flywheels and hydraulic accumulators, though by far the most common is an electric battery. A hybrid vehicle incorporating an electric powertrain is termed a Hybrid Electric Vehicle (HEV) and uses an Electric Machine (EM) to convert electrical energy into mechanical energy and vice versa; the term 'electric machine' is used to emphasise the ability of the machine to operate as a bi-directional energy converter in both motoring and generating regimes.

Some of the companies for whom fuel economy is a particularly pertinent issue are those operating large fleets of commercial vehicles in an urban environment. The drive cycles of these vehicles combined with potentially very large fleets mean that fuel costs for such companies can be extremely significant. Although vehicle manufacturers have introduced HEVs, and these continue to gain increasing market acceptance, their application has been almost exclusively focused toward the passenger car market. This is a peculiarity when it is considered that in many cases the typical drive cycle of Light Commercial Vehicles (LCVs) – delivery vans, for example – is far better matched to the characteristics of the hybrid powertrain than many passenger car trips. HEVs have historically been subject to cynicism from some parties, and it has been questioned whether early designs actually made any improvement at all. Some scepticism may be a result of manufacturers claiming fuel saving figures derived from laboratory testing which do not reflect the savings achieved in the real world; this is part of a much bigger issue pertaining to manufacturers optimising vehicle performance to legislative drive cycles which do not correspond to real world conditions.

This research will examine the potential for supervisory control of HEV energy management which is optimised for real-world use, so as to ensure that the

hybridised powertrain delivers the best possible fuel consumption savings when operating during normal use. With this target an important part of the work is to characterise 'normal use' of vehicles in the real world, and as well as this to acknowledge some difficulties which affect HEVs operating in the real world which are often overlooked in simulation and dynamometer testing.

This work was undertaken in collaboration with Ashwoods Automotive Ltd (Exeter, UK) – specialists in retrofit hybrid electric powertrains. Although the techniques and knowledge gained through this work are applicable to a broad range of road vehicles its scope is specifically the application of retrofit hybrid-electric systems to LCVs. The significance of this market will be explained in the following sections.

1.2 Transport in the Context of Climate Change and CO_2

In the United Kingdom the transport sector generates 21% of total greenhouse gas emissions – the single biggest contributor aside from the energy generation sector – and of these 92% are owing to road transport [2]. These figures are reflected in countries around the world and, in combination with problems arising from city smog and energy security, have motivated a coordinated attempt to reduce pollutant emissions from road vehicles over time. The greenhouse gasses emitted by road vehicles constitute almost exclusively carbon dioxide (CO_2), which is a product of hydrocarbon combustion and so is proportional to fuel consumption.

Since the introduction of European emission standards (Euro 1, Euro 2, etc.) in 1992 and similar standards worldwide tailpipe emissions have been controlled increasingly stringently. Every vehicle sold in the European Union must go through Type Approval, which includes rigorous emissions test procedures. Regulation (EC) No 692/2008 [3] outlines the implementation of Euro 5 and Euro 6 emission regulations, while Regulation No 83 [4] deals with specific testing procedures and protocols for all passenger and light commercial vehicles, including electric and hybrid-electric vehicles.

Trends in public attitudes toward transport, including climate change and CO_2 emissions, are monitored in the UK by the Department for Transport. Results of a

2010 survey [5] showed that 82% of participants believe that climate change will impact the UK, and that 69% are concerned about this. The statement "Transport is one of the major contributors to climate change" was correctly identified as true by 66% of participants. This concern for the environment affects what consumers look for in the market place, and 56% of people said that low carbon emissions would be "high on my list of must haves" when buying a new car, though environmental friendliness may not be the only motive for this response. With fuel costs continuing to rise in recent years 55% of people cited running costs as a reason for wanting a car with lower carbon emissions, while 43% mentioned the environment. Another survey [6] which is conducted every year, but has a smaller sample, suggests that levels of belief and concern about climate change have actually fallen slightly in recent years, though still remain high, with 76% of respondents at least fairly concerned about climate change. In this survey 38% of people considered environmental friendliness important when buying a car whilst 76% deemed cost important, of whom 53% considered fuel costs significant.

As well as public opinion and fuel costs, government schemes such as the Carbon Reduction Commitment Energy Efficiency Scheme add another incentive for reducing emissions. This mandatory scheme requires companies not covered by other legislation to report on their annual CO_2 footprint as calculated according to The Carbon Trust Standard Rules which includes fuel used in company vehicles [7]. As well as these figures being published in a league table, therefore adding a dimension of competition and company reputation, companies are also charged for their carbon emissions, adding further fiscal incentives. Environmental concerns, the cost of fuel, and government incentives are all strong drivers for reduced fuel consumption and CO_2 emissions in the automotive transport sector.

1.3 Significance of the Existing Vehicle Fleet

Once sold, vehicles tend to stay on the roads for around a decade, with the mean age of vehicles being in the order of 8 years [8, 9]. In Great Britain 59% of all cars are over 6 years old, and 78% are over three years old [8]. Light commercial vehicles (LCVs) tend to be slightly newer than cars, with the comparative figures being 57%

and 75% respectively [9], this perhaps owing to their more intensive use and therefore faster aging.

As a result of the slow turnaround of the vehicle fleet any new legislative change or technological advance which was to result in, for example, a 50% reduction in fuel consumption would take in the order of a decade to saturate the vehicle fleet and for its potential savings to be realised in the real world. For this reason technologies with potential to reduce the fuel consumption of existing vehicles have an important role to play in the short- and medium-term, while the knowledge gained during the development of such devices may also be transferable to new vehicle design.

Over 3.3 million LCVs are registered in the UK [9], covering an estimated fleet mileage of 43 billion miles per year [10] (about 12,800 miles per vehicle per year). LCVs may be used in a variety of applications, however among the most common are delivery (including courier and mail services) and technical maintenance or call-out services. Driving patterns for these applications typically involve a high proportion of time on urban and suburban roads as well as frequent stop-start behaviour, which are often regarded as the ideal conditions for HEVs due to the potential to recover substantial braking energy. Despite being strong candidates for hybridisation none of the leading LCV manufacturers currently produce hybrid or electric versions of their vehicles in any volume at a market competitive price; it therefore seems likely that even a decade from now the majority of the active LCV fleet will not be OEM produced hybrids.

Since the duty cycle of LCVs means that they are likely to benefit from hybridisation, and given also that their intensive use means they consume a considerable volume of fuel annually, the business case for their hybridisation is strong. Having observed the potential for fuel saving by hybridisation of LCVs Ashwoods Automotive Ltd have developed a retrofit hybrid system aimed at this market, which will be described in the following section.

1.4 Introduction to Ashwoods' Hybrid Electric Vehicle

Ashwoods Automotive's hybrid electric conversion system consists of a battery pack and an EM coupled to the propeller shaft as shown in Figure 1-1. Energy may be recovered from the rear axle during braking by using the EM in its generating regime to charge the batteries; this energy may later be used to assist the engine by delivering a portion of the driver's power demand. Systems of this kind are often called Kinetic Energy Recovery Systems (KERS). It is also possible to perform some energy trading by increasing the load on the ICE to charge the batteries; this moves the ICE to a higher torque operating point where it is typically more efficient, though this efficiency gain may be outweighed by the round-trip mechanical-electrical-mechanical energy conversion losses. This architecture is known as 'parallel' or 'torque-assist' because the EM acts to deliver tractive torque alongside the conventional ICE, with both conventional and electric power paths having a mechanical connection to the wheels; other common HEV architectures include 'series' and 'power-split'.
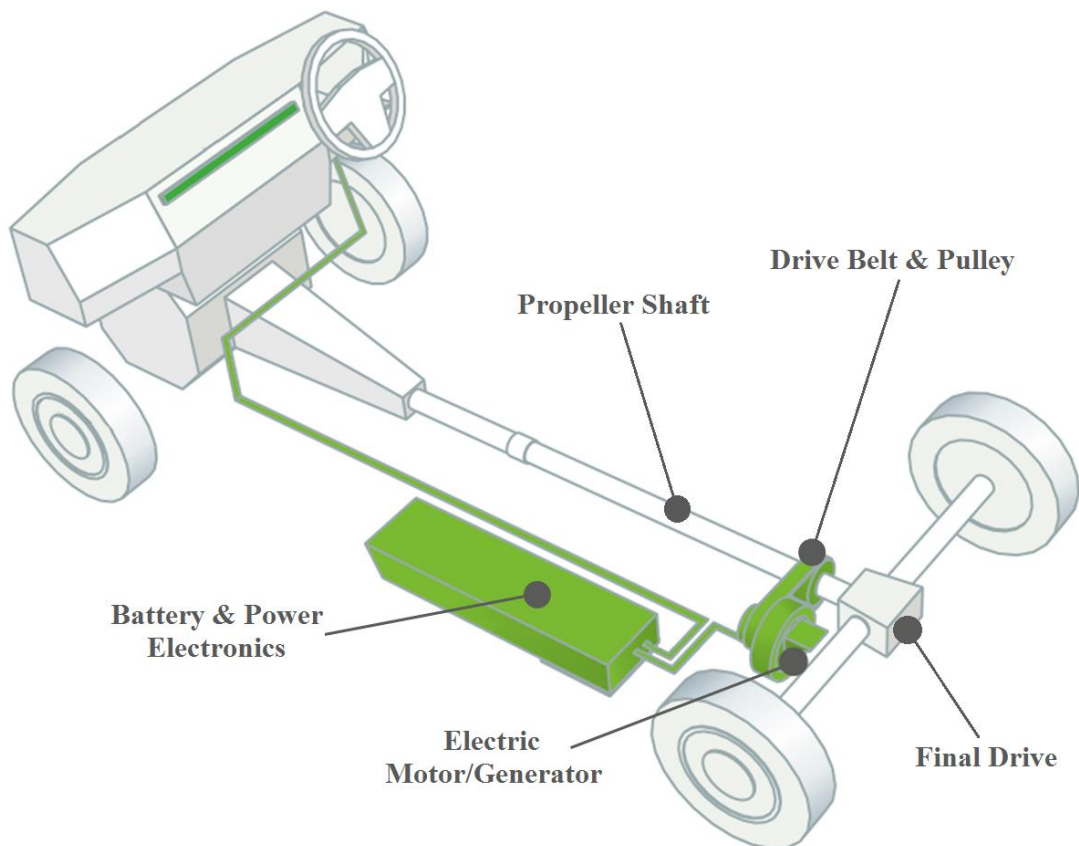
Figure 1-1: Architecture of Ashwoods Automotive's retrofit Hybrid Electric Vehicle. Retrofit components are highlighted in green, while OEM components are shown in grey.

In series architectures tractive force is provided exclusively by the EM, while the ICE is used for the sole purpose of generating electricity to maintain the level of energy stored in the battery pack and so has no mechanical link to the wheels. Series architectures usually exploit a downsized ICE which may be operated near its maximum power, where efficiency is best, to supply the average power requirement of the drive cycle. The battery pack acts as an energy buffer and stores enough energy for the EM to meet transient power demands. Series hybrids are well suited to urban driving conditions but suffer from multiple energy conversion efficiency losses. Power-split hybrids such as the Toyota Prius (described in [11]) exhibit some extremely attractive traits, allowing exploitation of the strengths of parallel and series architectures while avoiding their weaknesses. These systems allow power flows to be blended through the use of, for example, epicyclic gearboxes where the wheels, ICE and EM are connected to the sun gear, ring gear and planet carrier in some order, therefore allowing great flexibility in how the required wheel speed and torque are met.

Whilst both series and power-split architectures have their strengths both require heavy integration with the conventional powertrain which is problematic in the context of a retrofit system. Series architectures invariably lead to removal of the vehicle transmission, at least two EMs (and therefore two power electronic inverters), and probably engine downsizing. A power-split architecture would probably replace the standard vehicle transmission with an epicyclic gearbox, and this architecture also often uses two EMs. Although there are many configurations which could be adopted with great technical scope the complication and cost involved in these conversions eliminates them as viable retrofit options.

As a result of the problems identified with other hybrid architectures Ashwoods Automotive have identified the parallel HEV architecture as the only one which may be easily and cheaply implemented as a retrofit system. The simplicity of this design also means that most manufacturers will still honour their standard warranty, since no components of the standard vehicle are actually modified – this is a vital issue when selling into the commercial vehicle sector. Furthermore the system is entirely fail-safe, any malfunction simply reverts the vehicle to its standard (non-hybrid) condition in which it is drivable as though the hybrid system was not present. Even

in the most dramatic failure mode where for some reason the EM were to become seized the drive belt simply breaks, preventing any damage to the standard vehicle.

Having established the reasons for the mechanical and architectural design of the hybrid system this research will focus on its control. It may be expected that energy will be recovered by the EM at a relatively consistent rate depending on the frequency of braking events in the drive cycle. Given this availability of energy the question arising is when this should be re-deployed to maximum effect and when, if ever, the energy recovered should be augmented with energy generated from the ICE. The object of this research is therefore to develop an optimal energy management policy which maximises the potential gains of the hybridisation.

For the control of most hybrid vehicles it may be assumed that the driver's power demand power is shared between the ICE and EM, that is to say that for a fixed accelerator pedal position the controller is capable of increasing the power delivered by the EM and correspondingly decreasing the power delivered by the ICE to meet the driver's demand exactly. Whilst this would almost certainly be a valid assumption for a vehicle designed by an Original Equipment Manufacturer (OEM) with a heavily integrated hybrid powertrain, it is not at all the case for a retrofit vehicle where no supervisory control is exercised over the ICE. For this reason the addition of the electric powertrain essentially supplements the power available from the ICE, making the vehicle more powerful. It might be hoped that in response to this drivers would reduce the accelerator pedal position such that less power is delivered by the ICE thereby saving fuel, however this is by no means guaranteed. Furthermore the optimal control policy for the hybrid vehicle depends greatly on the typical drive cycle which it is used for, and so for these reasons this work also includes significant research into driver behaviour characterisation and modification, as will be described later.

1.5 Work Split

Since the extent of this work is quite large it is important that the personal contributions of the author are clearly presented. Activities undertaken in the course

of the research included mapping, theoretical modelling, control design and testing of a real HEV, as well as the design and field testing of a driver feedback device (known commercially as *Lightfoot*) and analysis of the resulting data. Broadly speaking Ashwoods Automotive were responsible for hardware design, integration and low-level coding of drivers and protocols, whilst the author was responsible for the design and evaluation of top-level control strategies and logic. The work packages forming the project, their interdependences and the party responsible for their completion are shown in Figure 1-2.
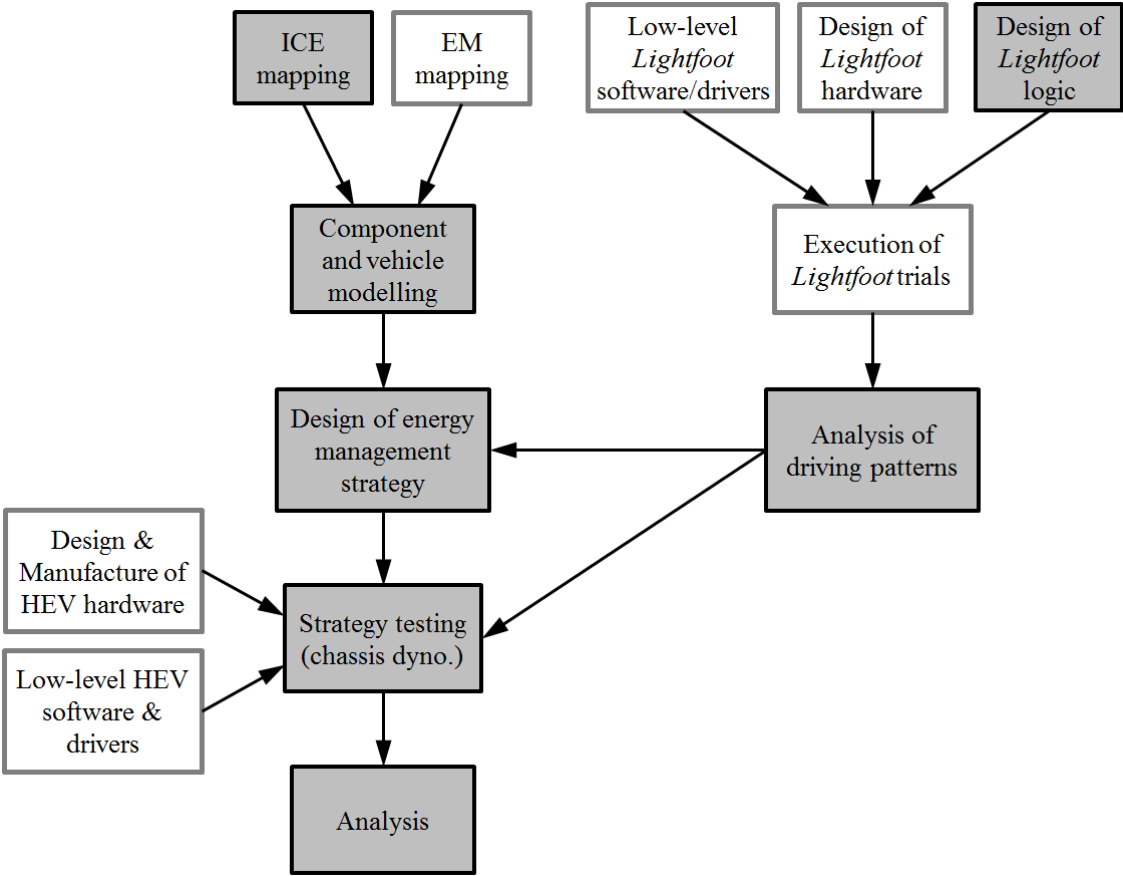
Figure 1-2: Work packages forming the project and their interdependences. Those completed by the author are grey (▇) whilst those completed by Ashwoods Automotive Ltd are white (☐).

1.6 Chapter Conclusions

Whether motivated by environmental concerns or the incentive to save money, reducing vehicle fuel consumption is now a priority for many individuals and companies. Despite the apparent aptness of HEVs to many LCV duty cycles OEMs have refrained from producing hybrid models of their commercial vehicles. For this reason Ashwoods Automotive Ltd produce a retrofit hybrid-electric conversion system.

This research will seek to understand how the energy management control system for the HEV should be designed so as to maximise its performance. In doing so it is important to understand vehicles' typical use in the real world and to ensure that they are driven appropriately; to this end the work includes the design and evaluation through field testing of a driver behaviour improvement tool. Data collected during testing of the driver behaviour tool are fed into the design of optimal control strategies and used to select an appropriate test cycle to evaluate HEV controllers.

Chapter 2   Literature Review

Given the scope of the project set out in Chapter 1 this review covers present state-of-the-art in hybrid vehicle control, seeking to identify opportunities for further research. It is observed that Stochastic Dynamic Programming has been shown to have great potential for predictive control, but has not been convincingly demonstrated in control of a real vehicle. Furthermore most optimal controllers reported in existing literature do not sufficiently account for practical limitations of hybrid electric systems, for example the consequences of persistent high power operation. Finally, it is noted that in the present context of a retrofit system, where the installed vehicle power is increased, some degree of driver behaviour monitoring is necessary and approaches to implementing this are reviewed.

In this chapter state-of-the-art energy management strategies for hybrid vehicles will be reviewed within the context of the control problem at hand. As will be discussed the best control strategy fundamentally depends on the drive cycle which a vehicle is typically used for; this motivated a full investigation into the typical driving patterns of the target vehicles, which also included the development of a tool designed to reduce fuel consumption by encouraging efficient driving behaviour. These topics are introduced and addressed in the later part of this review. Gaps in the present understanding are highlighted, which form the basis of the research aims.

The topics covered in this survey are non-exhaustive, partly for brevity, and also because there is a huge volume of literature examining HEV development, much of which is not relevant to this research. As an example, the relative strengths and weaknesses of different HEV architectures (parallel, series, power-split, etc.) are not examined. The object of this research is to examine retrofit technologies with commercial viability; for this reason the parallel torque-assist hybrid architecture was predefined outside the scope of this work, as described in Chapter 1. The emphasis of the review will therefore be on optimising the efficacy and performance of the system in the real world.

2.1 Hybrid Electric Vehicle Control

Achieving an optimal control strategy which maximises the potential benefits of a hybridised powertrain is non-trivial because the solution fundamentally depends on the future. For example, if the driver is soon to go down a long steep hill which will provide opportunity for kinetic energy recovery then it would make sense to deplete the stored energy reserves to achieve maximum immediate benefit, freeing up the capacity of the energy storage system (ESS) with the expectation that the energy will soon be replaced. Conversely, if it is unlikely that there will be a huge opportunity for energy recovery in the immediate future then the question becomes "When can the energy which is currently stored be deployed to greatest effect"?

In some cases it is possible that the future drive cycle is known in advance, if not precisely then at least to a good degree of certainty. Examples of such cases would

include railway applications, or Personal Rapid Transit systems. In road-based transportation advances are being made toward fully autonomous vehicles which would know their route in advance and could therefore plan an energy management policy. Perhaps more immediately conceivable are vehicles in which the driver's satellite navigation system can feed the powertrain supervisory controller with information about the likely future. For more familiar and mundane journeys such as the daily commute the driver may not be using a navigation system but, based on a database of similar journeys started from the present location at that time of day, considerable information about likely routes could probably be inferred by the controller. Nevertheless at present the future is entirely unknown to the powertrain control system in the vast majority of applications.

2.1.1 Control Classification

Salmasi [12] reviewed and classified approaches to the control of HEVs and proposed a classification structure which has been adopted by others [13]. This classification structure is presented in Figure 2-1 and each of the four categories of control strategy will be explained and discussed.

Rule-Based Control

Rule-Based controllers follow rigid logical protocols in order to decide on an appropriate power split for the hybrid vehicle powertrain. Based on inputs such as power demand, vehicle speed, acceleration, and state of charge (SOC) of the ESS the controller will behave in a predictable manner following logical statements to
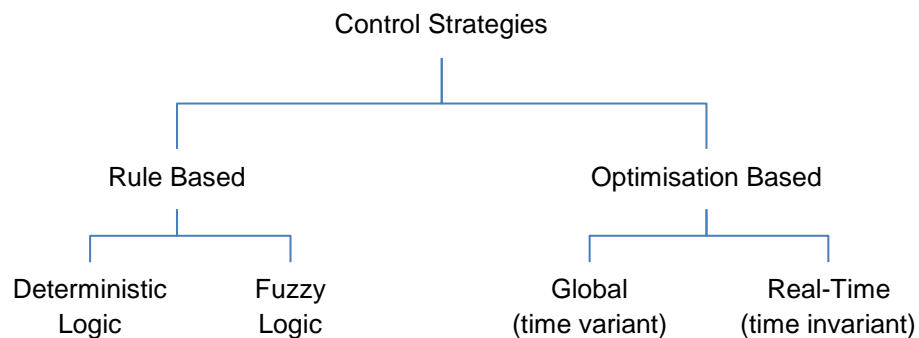
Figure 2-1: Classification of control strategies, adapted from [13].

demand either positive (assist) torque or negative (regeneration) torque from the EM. Rule-based controllers may be further sub-divided into Deterministic and Fuzzy Logic based classes.

Controllers categorised as 'Deterministic' are typified by the use of flowcharts, state flow diagrams and lookup tables to determine the vehicle state and make appropriate decisions. The controller outputs are therefore dictated by a set of discrete logical statements, probably designed by engineers based on heuristic logic using a set of desirable operating conditions and efficiency maps of the powertrain components.

Fuzzy Logic controllers gained popularity in an effort to move away from the rigidity of determinism. The central concept of fuzzy logic is to move away from 'crisp' variables (precise numeric values) and towards the use of linguistic measures (temperature could be cold, warm, or hot) with which humans are inherently more familiar [14, 15]. Rather than using raw values of input parameters, each is assigned a degree of membership to a fuzzy set, for example speed could be categorised as slow, medium, high or some combination of states. A variable's degree of membership to any possible fuzzy state is between 0 and 1, and the sum of the degrees of membership must equal 1 at all times; this input processing stage is known as 'fuzzification' (Figure 2-2).

The control response to each fuzzy state may be different and so the second stage,
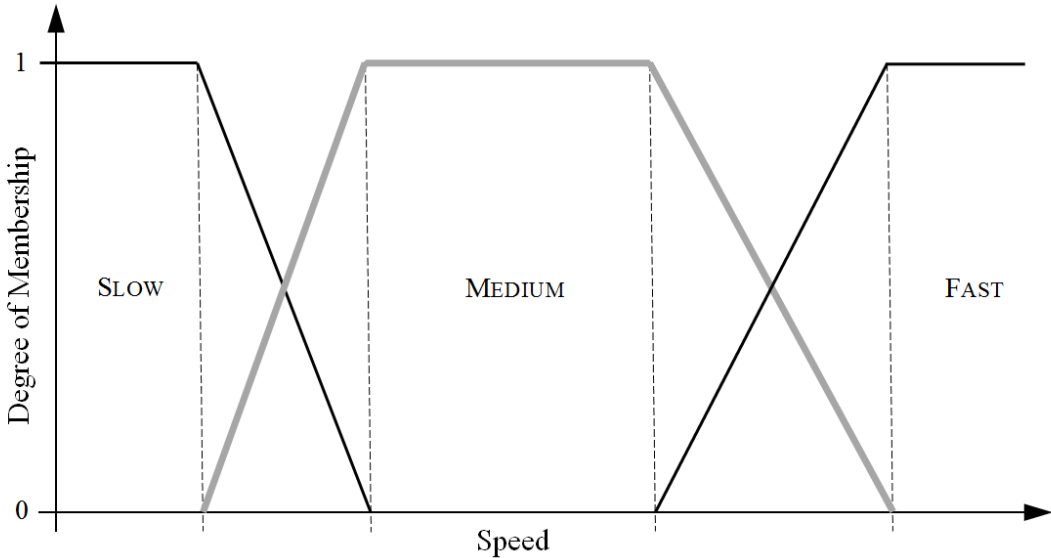
Figure 2-2: Fuzzy logic may be used to assign a crisp value of speed with degrees of membership to a fuzzy set. A level of abstraction is therefore added into the logic, which can simplify complex strategy design.

known as 'inference', uses fuzzy logic operators to establish the extent to which the 'antecedent' (condition) of each fuzzy rule is satisfied. Finally, through a 'defuzzification' process the 'consequents' (results) of each applicable fuzzy rule are combined and converted back into a 'crisp' output variable.

Salmasi regards fuzzy controllers as an extension of 'conventional' rule based controllers where the primary advantages are added robustness (for example to imprecise measurement) and ease of use, both in initial calibration and in adaptation to new or varying scenarios [12].

Lee and Sul [16] first showed the potential for the application of fuzzy control to HEVs, in this case with a view to reducing NOx emissions. The comparative ease with which a fuzzy controller could be implemented and tuned compared to deterministic control was highlighted and the control was shown to be robust. However, the controller was relatively simple as it only received engine speed and throttle demand as inputs and also sought only to optimise ICE performance. An interesting feature of the real-world parallel HEV implementation in this case is that the EM was integrated on the engine side of the transmission, taking the place of the flywheel. This is in contrast to the majority of parallel HEV implementations which integrate the EM downstream of the transmission. Pre-transmission integration of the EM has a significant advantage in that it greatly reduces the speed range that the motor operates over. The result of this is that for an EM of the same torque rating a pre-transmission motor can operate near its rated power much more of the time; the disadvantage is that this configuration requires heavily integrating the EM, adding complexity and cost to the system. In the present case the EM was actually operated almost exclusively in the field-weakening (constant power) region, meaning that the rated power of the motor could be used frequently, allowing maximum benefit.

Schouten et al. [17] built on the work of Lee and Sul [16] by adding SOC as an input variable to the fuzzy controller and optimising the HEV holistically, taking account of the efficiencies of the various powertrain components. This controller was implemented in a simulation environment and assumed the more common post-transmission EM configuration.

Numerous studies claim to show the benefits of fuzzy over conventional control [16, 17] however it is difficult to know whether these are really fair comparisons. It should be noted that the distinction between deterministic controllers and fuzzy controllers can be overstated. In most cases it is possible to implement the same control logic as a fuzzy controller using look-up tables, which are regarded as 'deterministic'. The primary advantage of the fuzzy controller lies in its higher level of abstraction, which allows easier implementation and tuning, and greater robustness. However this is also its weakness, since this same trait means that it is less precise.

Salmasi concluded that rule based approaches are simple and effective in real-time supervisory control, which led to their widespread use in early HEVs: both the Toyota Prius and Honda Insight are based on deterministic logic. However, these strategies are inherently inflexible and therefore do not perform so well over different types of drive cycle. Furthermore rule based approaches tend to optimise components in isolation rather than the system as a whole (holistically), and Wirasingha and Emadi [13] suggest that for these reasons developers have moved their attention towards optimisation based strategies.

Optimisation-Based Control

Design of rule based control strategies is almost exclusively based on heuristic logic, or on trial-and-error, to determine the best operating mode for the vehicle. Such techniques offer no guarantee of optimality in the strict mathematical sense of ensuring the control strategy is the best it could be. Furthermore trial-and-error development, for example evaluation and refinement of a controller's performance through on-road testing, is time consuming and expensive. In recent years it has become common practice to use 'model-based design' techniques to design controllers in a simulation environment, which then require considerably less testing and refinement. Model-based design practices have also made possible the study of strictly optimal controllers for which a 'cost function' is defined. The optimum strategy is that from the set of all possible strategies which minimises this cost.

In practise it is common that the control problem definition for an HEV is in fact a multi-objective one. For example, as well as reducing fuel consumption, which is

18

usually the primary goal, other objectives may include reducing pollutant emissions and improving drivability; these secondary objectives often oppose the primary and so a trade-off is necessary. In order to manage the compromise between multiple objectives a cost function may be formulated in which all of the costs/objectives are represented and weighted to varying degrees depending on their relative importance; the optimal all-round solution is then found by minimising the total cost incurred.

The distinction between 'global' and 'real-time' optimisation (Figure 2-1) stems from whether the control problem is formed with the assumption that the drive cycle is known entirely in advance or not. As discussed previously the truly optimal control strategy would require information about the future in order to decide how best to manage the stored energy. Given this information a 'globally optimal' strategy would control the vehicle powertrain so as to minimise the cost incurred over an entire drive cycle. An important quality of such a system is that the control decision at any time is not only a function of the vehicle's past and present states, but also its future. As a result the strategy's response to any vehicle state in the present time is not fully defined without the knowledge of the future and so, for two identical vehicle states, the strategy may make different control decisions because of differences in the projected futures. For this reason globally optimal control strategies may be labelled 'time variant'.

Real-time optimal controllers are those which do not assume any knowledge of the future, but retain some degree of optimality. Given that in the vast majority of cases the drive cycle is not known in advance this set is concerned with finding the most optimal controller which is practically implementable, i.e. 'time invariant'. In the following sections the three most prominent optimisation-based control solutions will be presented and more thoroughly reviewed.

2.1.2 Dynamic Programming

Dynamic Programming (DP) is perhaps the most established global optimisation technique, however there also exists some potential for its adaptation to real-time implementation. The method is based on Bellman's Optimality Equation which is used to formulate a recursive equation describing the minimum possible cost that can

be incurred during operation of the system under the optimal control policy and starting from its present state:

$$J_k^*(x_k) = \min[c(x_k, u_k) + J_{k+1}^*(x_{k+1})] \tag{2-1}$$

where k is a discrete time index, x_k is the vehicle state vector at time k and u_k is the control decision at time k. The instantaneous cost arising from the present vehicle state and control decision is $c(x_k, u_k)$, however it is expected the present control decision will affect the vehicle state at the next timestep and so $J_{k+1}(x_{k+1})$ is the total summative cost incurred in the next timestep and all timesteps in the future thereafter. The superscript * denotes optimality, and so Equation (2-1) represents the statement that *the minimum total cost to progress from state x_k is achieved by the optimal control policy which minimises the sum of the instantaneous cost and all future costs*.

At face value this statement appears rather trivial, however the power of DP lies in the method which may be applied to solve for the optimal control policy. Before explaining this however, the single biggest limitation of Equation (2-1) should be highlighted: since the equation is recursive, containing total cost J on both sides, an attempt to solve the equation without a finite end time would result in an infinite regress. As a result the method is only applicable to situations where the drive cycle is known entirely in advance and so a finite chain of causal events may be considered. It is for this reason that DP is a global optimisation technique. The equation is solved backwards in time, starting from $k=N$ (the end of the drive cycle) and moving towards $k=0$, to solve for the absolute optimum system performance over a known drive cycle subject to constraints.

DP is helpful when the control decision at the present timestep directly affects the system state at the next timestep; for example in the problem of hybrid vehicle control the decision made in the present timestep will affect the subsequent SOC. This is an important notion because the instantaneous fuel consumption will always be reduced when tractive force is provided by the EM rather than by the combustion engine; however for a system with a limited reserve of electrical energy this use of electricity necessarily means that less will be available in the future, and so the future

cost may be higher as a result. The optimality problem could be solved by brute force, by directly enumerating every possible control trajectory; however this is extremely wasteful as many of the trajectories will be very similar. For example, starting from state x_k a set of simulations may follow the same control trajectory over the entire problem up until the very last timestep, at which point each simulation makes a different decision in order to see which trajectory was best. The result is a set of control trajectories that are extremely similar, only different at the final timestep, but each is simulated entirely. If N_x is the number of discrete states, N_u the number of possible control decisions in each state, and N_k the number of timesteps in the drive cycle, this direct enumeration would therefore require the simulation of $N_x \times N_u^{N_k}$ state transitions.

Instead of enumerating every possibility DP considers the problem starting from the final timestep, working backwards in time towards the beginning of the problem, quickly rejecting sub-optimal solutions. It may be conceptually helpful to consider the algorithm's benefits in a more trivial case: The common shortest path problem may require the shortest path from B→C to be found. If a new problem were then posed to find the shortest route from A→B→C (i.e. A to C, via B) only the solution to A→B must now be found, because the solution to B→C was found previously and is still optimal. In the same way DP breaks down top level problem into a series of single stage problems, each solved in turn, to build up the full solution.

The procedure for DP is to first consider the one-stage (instantaneous) cost for being in the final state, for example at *k=N* in Figure 2-3. Moving backward in time one step to *k=N*-1 a control decision in each state will define the state at the next timestep (*k=N*), and so the one-stage cost of being in each state is added to the optimal cost from the resulting state onward. By following this procedure a table can be constructed containing the optimal cost of moving from any state *x* at time *k* to the end of the problem, where non-optimal control trajectories are ignored completely. By removing the repeated simulation of trajectories that are known to be sub-optimal the total number of state transitions which need to be simulated to establish the table of costs is $(N_x \times N_u \times N_k)$, which is likely to be considerably lower than $N_x \times N_u^{N_k}$.

Having populated a table of costs the optimal control decision from any starting state may be read from the table, moving to the subsequent state defined by the present control decision, and continuing to follow the optimal path from there.

These descriptions of the procedure are an extremely brief summary as a full explanation would be considerably longer and not proportional to the method's pertinence to the present work; an excellent practical guide to the method with worked examples is provided by Larson and Casti [18].

As well as being non-causal in time the DP result is also a *time-variant* full-feedback controller; that is to say that the controller adapts with time to suit future power demands and so a given system state (perhaps defined by vehicle speed, power demand, and SOC for example) may result in a different control output at two different times. As a result the controller's response cannot be reproduced in real-time without complete knowledge of the driving cycle and so is not implementable in the majority of applications. Having said that, by examining the average behaviour of the DP controller it is possible to derive a deterministic controller designed and tuned to emulate the average behaviour of the DP control, as is demonstrated by Lin et al. [19, 20]. The resulting controller is a near-optimal approximation of the DP control when operating on the drive cycle which the DP controller was developed on,

$x_k \in \{x_1 \dots x_6\}$

	0	1	2	3	4	5	…	…	N-2	N-1	N
x_6					1.3	1.75	1.21	1.46	0.91	1.47	0.81	0.81
x_5					1.98	1.81	1.43	1.36	0.89	0.83	1.03	0.24
x_4					**1.69**	1.52	1.07	0.9	1.36	1.31	0.55	0.93
x_3					2.15	1.15	1.5	1.23	1.61	1.3	0.67	0.35
x_2					1.23	1.29	1.16	1.05	1.77	0.93	0.78	0.2
x_1					1.59	1.98	1.55	1.68	0.96	1.12	0.75	0.25

Timestep, k
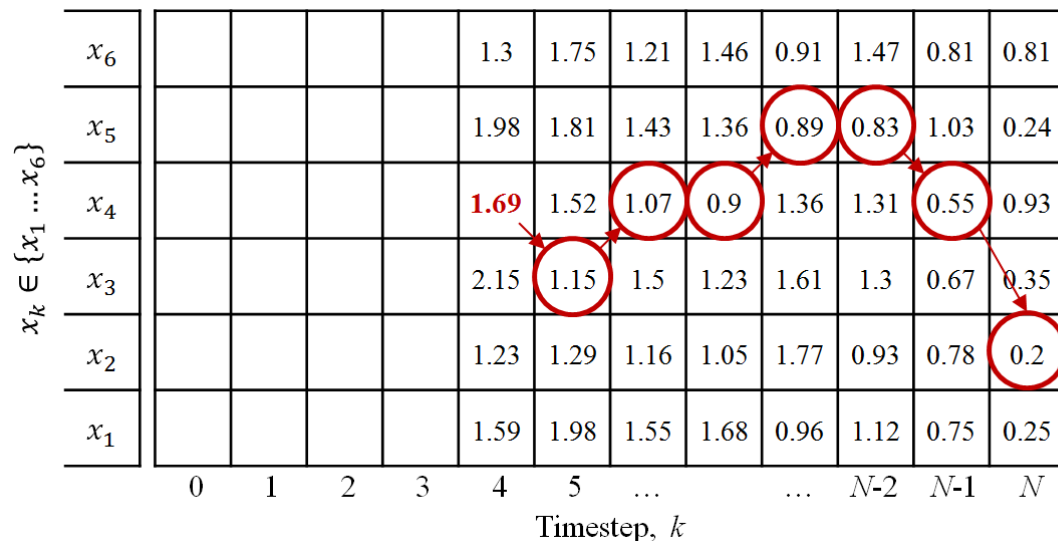
Figure 2-3: Diagrammatic representation of a Dynamic Programming cost table. The number in each grid space is the minimum cost which can be incurred in progressing from that state to the problem end. Starting from state x_4 at time $k=4$ the minimum is 1.69. The instantaneous cost of the optimal trajectory at time $k=4$ is 0.54, and so the minimum cost from the next timestep ($k=5$) to the end is 1.15.

however its performance over other drive cycles (that the DP controller was not developed for) may not be at all optimal, and these controllers are therefore accused of being 'cycle beating'.

Despite its limitations the DP approach is very useful in gauging the theoretical performance limit of a system if it were controlled perfectly. Furthermore there remains significant potential for DP when integrated with in-vehicle Global Positioning Systems (GPS) to predict the future vehicle path and optimise the HEV accordingly [21-23], though this is by no means a simple task.

2.1.3 Stochastic Dynamic Programming

To overcome the limitations of DP with a view to implementing an optimal control strategy which is causal Stochastic Dynamic Programming (SDP) has been proposed, the application of which to HEV control is usually credited to Lin et al. [24]. In SDP the precise future expected cost $J_{k+1}^*(x_{k+1})$ in the DP algorithm is instead replaced with a stochastic representation of what this cost is *likely* to be, based on historical data, thereby avoiding the requirement for *a priori* knowledge of the exact drive cycle.

SDP requires a library of historic driving data which can be statistically analysed to generate a state transition probability matrix, describing the probability of transitioning from any vehicle state to any other. Given the present vehicle state this stochastic representation of state transitions can therefore be used to find the statistical distribution of likely states at the next timestep, and this distribution used the find the statistical distribution at the following timestep, and so on; the statistical distribution of vehicle states may therefore be forecast far into the future with decreasing certainty. For a particular control policy which is a function of vehicle state the cost associated with each state (and the resulting control decision) can be explicitly calculated, and so it becomes possible to create a *future expected cost* function, which takes into account the likelihood of future vehicle states as well as the cost of a particular controller's response to those states. With respect to the original DP formulation presented in Equation (2-1) the stochastic version replaces the second term (future cost) with an equivalent *expectation* of the future cost, denoted by *E*:

$$J_{k+1}(x_{k+1}) \sim \lim_{N \to \infty} E \left\{ \sum_{k=1}^{k=N-1} \lambda^k c\left(x_k, \pi(x_k)\right) \right\}. \qquad (2\text{-}2)$$

This expected future cost has an infinite horizon, that is there is no limit as to how far into the future the expected cost could be predicted, but over an infinite time horizon the expected cost would clearly be infinite. For this reason the *discount factor*, $0 < \lambda < 1$, is essential in ensuring that the sum converges to a finite limit. Since the discount factor is raised to the power of the time index the value of λ^k decays exponentially such that costs in the distant future carry far less weighting than costs in the immediate future. Although there is no threshold or fixed time window being considered (the discounting is exponential) the value of the discount factor effectively determines how far into the future the controller considers by defining how rapidly future costs are discounted. For this reason the value of the discount factor has a considerable effect on the resulting control strategy; however this effect, and how a final decision on its value should be reached, is not adequately explored in existing literature. In two worked examples Kolmanovsky et al. [25] state that they used discount factors of 0.95 and 0.98, though without further explanation. Johannesson [22] used 0.995, while Lin et al [24] used 0.95, whose example was followed by Tate el al [26] and Moura et al [27], with Moura noting that the choice of discount factor is a matter open for further research.

Beyond the choice of discount factor the existing literature in which SDP is used generally provides very limited guidance on how the algorithm should be applied successfully; there is no well documented method summarizing good practice implementation. For example, in order to find the optimal control policy SDP requires an iterative procedure in which the first step is to find the expected cost function for an arbitrary controller. Based on this cost function an improved control strategy is found, for which the cost function may be calculated, and so on. Much iteration of successive policies is generally necessary, but how to determine when iterations may be truncated is not explored. Yet another issue not satisfactorily resolved in previous works on the subject is that the resulting policy is a multi-dimensional state lookup table which is practically impossible to examine manually,

and how the charge sustenance properties of the resulting policy may be examined or adjusted is therefore not clear.

The potential of SDP is extremely appealing because it allows a time-invariant full-feedback optimal controller to be generated. Although the controller may not be optimal among the set of all possible controllers, it is guaranteed to be optimal among the subset of controllers which are 'stationary' (time-invariant) policies, with the proviso that the historic stochastic driving data continues to be an accurate representation of the vehicle's future use. The poor understanding of how the algorithm should be implemented robustly in real vehicles is likely due to the fact that in the vast majority of cases its evaluation has been limited to simulation; only one example of the algorithm's application to control of a real HEV is known of [28], and though the results looked promising the study had limited applicability because no real world driving data was available from which to form a stochastic model.

2.1.4 Equivalent Consumption Minimisation

Moving away from global optimisation altogether some degree of optimality may still be retained by subtly changing the objective to the minimisation of instantaneous equivalent cost. When first introduced by Paganelli et al. [29] this approach, named the Equivalent Consumption Minimisation Strategy (ECMS), marked a significant step forward in the approach to HEV control beyond deterministic logic because for the first time it allowed the controller to be optimal to some degree. The term 'equivalent consumption' refers to the equivalence drawn between fuel energy and electrical energy, which is used to ensure that the ESS maintains a suitable SOC. It is defined by Salmasi [12] as the "extra fuel consumption that will be required for the battery recharge in the near future" if the electric motor is used to supply some of the driver's power demand instead of the ICE. Since the use of electrical energy always reduces fuel consumption this equivalent cost is essential to ensure that the value of the electricity is considered; the instantaneous control decision is always the one which minimises the *combined* cost of fuel and electricity, and in the same manner as previously described other objectives may also be represented in the cost function with equivalence factors used to weight their relative importance.

Paganelli et al. [29] applied ECMS to a parallel torque addition HEV where the EM added torque to the drivetrain between the gearbox and final drive of the vehicle, and was therefore directly coupled to the wheel speed. The ICE was replaced with a downsized engine, and the EM accounted for approximately one third of the power capability of the new powertrain. In this example a known torque demand could be fulfilled by the ICE and EM in different proportions, and so in order to find the optimum torque split ratio the ECMS considered the sum of ICE fuel consumption and the 'equivalent fuel' consumption of the EM for every possible ratio. The torque split resulting in the absolute minimum equivalent consumption was then selected. In order to ensure that SOC constraints were adhered to a nonlinear correction factor was applied to the EM demand torque after the optimum operating condition had been determined. To compensate for this adjustment the ICE torque demand was corrected by an equivalent amount in the opposite direction to ensure the driver torque demand was satisfied exactly.

Clearly for ECMS controllers the price of electricity is extremely important, and was defined by Paganelli et al. [29] using the average cost of the fuel used to generate the electricity. In reality however the cost to generate electricity may vary depending on operating conditions or drive cycle, so a price with some degree of variability may be an improvement. Also, by adding the SOC correction after the optimum operation point has been determined and adjusting the torque split accordingly, the controller does not capitalise on its potential to optimise the system holistically.

Paganelli et al. went on to show how the ECMS could be used to control a fuel cell hybrid vehicle [30], demonstrating the versatility of the approach to different architectures. A significant advance in this paper was the adjustment of the SOC correction so that it acted as a multiplier on the cost of electricity rather than simply modifying the torque demand of the EM. In this way it was more integrated in the decision making process rather than a retrospective adjustment. This scenario is still not ideal in that SOC correction is enforced immediately, and the controller is allowed very little flexibility in when it is applied. It can be seen from some of the simulation results that at times in the drive cycle the equivalent fuel consumption curve is relatively insensitive to changes in power split, and therefore a large change in power split may sometimes only deviate from optimal operation very slightly.

These scenarios would be ideal for correcting SOC, and so a further improvement to the strategy may be to alter the SOC correction so that it is another control objective handled in a similar way to emissions in [29], allowing the controller greater flexibility in when it corrects the SOC.

Liu and Peng [31] compared SDP and ECMS strategies in the control of the Toyota Prius power-split system. They also compared the effectiveness of their controllers against the DP theoretical benchmark, unlike others [16, 29, 30], which allows for much better evaluation of their performance against a theoretical limit rather than against other heuristic controllers which are unfamiliar to the reader. They showed that both SDP and ECMS outperform traditional deterministic logic, and give relatively similar fuel consumption results despite delivering them in very different ways, though the SDP approach was marginally better in their simulations. For control using ECMS a tendency was observed for the ICE power demand to alternate between a high value and zero (engine off). This is logical because the controller tries to maximise instantaneous efficiency and, since the most fuel efficient operating condition for the engine is in the high load region, the controller exploits this. In contrast the SDP controller is much smoother in its control of the ICE as it also considers the future to some degree (statistically modelled), and so produces a control signal closer to the DP result.

2.1.5 Real-World Considerations

Optimal control design techniques, such as those discussed, rely heavily on model-based-design; they use computer simulations to assess and optimise the performance of a control strategy. This approach is enormously powerful, however computer models always require simplifications and assumptions to be made since including every physical dependence would make the model impossibly complicated and computationally expensive. While many simplifications are relatively inconsequential some risk introducing a significant disparity between the simulated and physical environments. As an example, it is common practice in all of the research discussed thus far that the electrical powertrain is modelled with a fixed maximum power limit which it may be operated at as much or as little as desired. In fact this is far from the truth, since electric motors can often deliver considerably

more than their 'continuous rated' power in short bursts, usually limited by temperature. Similarly, batteries may be operated at high charge/discharge rates in short bursts, but at the expense of increased heat generation and very often the Battery Management System (BMS) will struggle to maintain even temperature and SOC distributions between cells. In addition to these short-term problems batteries' usable life span may suffer as a result of consistently heavy use.

Since the great majority of the work in this area is based on simulation, with minimal real-world validation or implementation, the significance of ignoring battery cell balancing issues, battery thermal management and motor temperature has been understated, though these concerns are familiar in industry and to those working with hardware [32]. Paganelli et al. demonstrated real world control implementation [29] and Dubarry et al. [33] demonstrated how battery models built on laboratory data could increase understanding of real-life performance, but it is certain that more real world field trials are needed.

It has been noted by Plett [34] that the HEV environment is particularly harsh for batteries as it is frequently desirable to draw and return energy at extremely high rates. This makes sense because batteries are heavy and expensive and so it is essential that maximum benefit is obtained in exchange for carrying them. In contrast, full Battery Electric Vehicles require much larger battery packs to achieve satisfactory vehicle range, and this added capacity also tends to yield a higher power capability (even if the batteries in the two vehicles have the same energy and power densities), meaning that the powertrain is rarely operating at its power limit.

Cell Aging

Unfortunately high power use of battery cells is a well-known stress factor, accelerating capacity fade and reducing the useable life of the battery [35, 36]. For this reason it is usually preferable to operate batteries more consistently at a low power, rather than have a 'peaky' duty cycle with high rates of charge draw/acceptance. In order to smooth out the duty cycle batteries are exposed to and reduce battery stress load-levelling control strategies have been proposed which intelligently control the electrical load of auxiliaries [37, 38]. Hybrid energy storage systems have also been proposed, in which super capacitors (or mechanical

flywheels) are used during high power events to smooth out the peaks in battery power demand, while batteries remain the primary means of energy storage [39, 40]. It is widely supposed that such load levelling would also increase the service life of batteries, though good quality data to quantify this is not readily available.

The problem of high power operation is made considerably more pertinent in HEVs because it is normal to use battery packs made up of numerous cells. Ohmic losses cause heat generation in each individual battery cell proportional to its internal resistance. Manufacturing tolerances mean cells in a pack will have a spread of nominal internal resistance and as a result some cells generate more ohmic heating than others, which can cause thermal gradients to develop within the pack during normal use. Since thermal cycling plays a role in cell aging and affects internal resistance [41] this thermal gradient can cause cells to age differently and their internal resistances to deviate further. Uneven heat generation in battery packs is also problematic from a safety viewpoint because in extreme cases it introduces a risk of thermal runaway, requiring the battery power to be temporarily limited; for this reason battery pack cooling has become important [42].

Battery Management

On a shorter timescale the spread of internal losses between cells can also be very problematic because, for an identical current cycle, cells which start with identical SOC will begin to develop a disparity [43]. Since battery cell internal resistance is typically much higher at low SOC [44] this SOC disparity can become self-perpetuating if not properly managed by a BMS, since cells with low SOC become even less efficient.

The consequences of SOC imbalances are considerable because battery cells usually have specified safe working limits for terminal voltage. Nominal cell voltage at rest varies between cell types, but as an example a lithium ion cell may have a nominal voltage of 3.2 V which varies slightly between 3.1 V at low SOC and 3.3 V at high SOC. Drawing current from a cell causes a voltage drop proportional to the current draw and the cell internal resistance ($\Delta V = I \cdot R_{int}$), so cells with a higher internal resistance will experience a greater voltage drop for the same current. Should the terminal voltage fall below 2.7 V permanent damage may be done to the cell, and the

situation is mirrored during charging where the maximum safe terminal voltage may be in the order of 3.7 V. One of the roles of a BMS is therefore to implement over- and under-voltage protection logic such that the terminal voltages of all cells are kept within the safe range 3.2 ± 0.5 V. When the battery pack is considered as a whole the problem becomes apparent: in order to avoid damage to any cells the working limits, or operational window, of the pack must be limited by the highest and lowest cell voltages [45]. In a battery pack in which the majority of cells have the same SOC but one cell is lower the permissible charge/discharge rates of the pack may be severely limited because of the higher internal resistance of the one cell with a low SOC. The same scenario is true of cells having different thermal cycling histories.

The net result of all of these cell balancing issues is that at various times the BMS may be forced to reduce the power demand on the battery in order to ensure no cells incur any permanent damage, and to try to re-balance the pack [43].

Motor Considerations

Operating the electric powertrain at high powers also has consequences for the EM because ohmic losses in the motor windings, which are proportional to the square of current, cause the EM to become hot. If the EM is consistently operated at high powers and the rate of cooling is not sufficient the control system may be forced into thermal cutback, reducing the availability of the electric powertrain.

Relevance to Optimal Control

Since research in the area of HEV control tends not to consider the physical effects of high power operation an 'optimal controller' in simulation may actually prove highly sub-optimal in practice if it demands too much of the electric powertrain, which then goes into a self-protection mode. Conversely if researchers choose to stay within the confines of the continuous power limits then the efficacy of the system is not maximised.

Most HEV control strategies being developed employ a cost function to define when a control strategy is optimal [22, 31]; the optimal strategy is the one which minimizes the total cost over a drive cycle. In its most simple form this cost function is probably equal to fuel consumption, however efforts have been made to include several other parameters to achieve multi-objective optimisation; examples include equivalent cost

of electricity used [23, 27, 29, 46], NO_X and PM emissions [20, 24] and gear shift busyness [19]. Salmasi [12] suggested that further work should add durability extensions (or extension of the life of hybrid powertrain components) to controller cost functions. In light of the real-world problems experienced with hybrid systems, as discussed, it certainly seems there might be benefits to also considering short-term State of Health (SOH), or a measure of electrical powertrain stress.

There is no uniformly agreed measure of battery stress, however Moura [47] proposed two cost functions aimed at improving long-term battery health, one based on minimizing anode side solid electrolyte interface (SEI) film growth (which is widely regarded as a key cause of capacity fade) and the other designed to reduce battery current throughput. Of these two cost functions the former requires a relatively detailed model of the battery cell chemistry which may not be appropriate for higher level holistic system optimisation – the relative mathematical simplicity of throughput models makes them better candidates for control optimisation purposes. Both functions are aimed primarily at reducing long-term capacity fade with a view to increasing battery service life. Ebbesen et al [48] and Serrao et al [49] also demonstrated controllers incorporating throughput aging models. In both examples the rate of aging is acknowledged to increase at high battery currents, and so the cost function amounts to a non-linear function of instantaneous current, while [49] also accounts for temperature and SOC.

An alternative metric which has been proposed is the mean square of the battery current ($\overline{I^2}$) [50], with the reasoning that ohmic losses, and therefore heat generation, are proportional to the square of current. This is perhaps more appropriate for management of short-term system SOH, with the advantage that its physical meaning can be extended to the EM as well as the battery, therefore offering some indication of the stress of the entire electrical powertrain.

2.1.6 Drive Cycles and Driver Behaviour

Vehicles will achieve different levels of efficiency depending on the driving conditions in which they are used, and for this reason it is important when testing a vehicle to use a defined 'drive cycle' which specifies the vehicle's speed as a function of time. This issue is particularly pertinent for HEVs because the

effectiveness of the hybrid powertrain is highly dependent on the drive cycle it is used over. A hybridised vehicle may achieve considerable fuel savings over its standard counterpart on some drive cycles, but none at all on others. For driving cycles that involve a lot of accelerating and braking, such as urban conditions, there is considerable scope for kinetic energy recovery during braking which can be used to assist the engine when required. In contrast, for haulage vehicles which spend the majority of their time cruising at steady speed on a motorway no opportunity exists for kinetic energy recovery, and so the only energy available would be 'traded' by first using the ICE to generate the electricity. This generally does not result in worthwhile fuel savings because of the roundtrip efficiency losses during energy conversion. For these reasons it is important when developing and testing an HEV that the test cycle is a realistic representation of the real-world vehicle use.

For legislative purposes pre-defined drive cycles are used to test and report vehicle fuel consumption in a consistent manner, with tests varying between different countries worldwide. However, whether driving cycles accurately reproduce the real-world conditions that they intend to emulate is questionable [51-53]. Adornato et al. [53] compared legislative drive cycles with real-world data collected by the authors using average measures such as mean speed and specific energy (energy consumed per unit distance). It was decided that none of the standard drive cycles represented the real-world data collected by the authors, and so HEV simulations were conducted using the real-world data instead. The disadvantage of this is that simulations are not reproducible by others; for this reason where possible simulations and laboratory experiments should be based on widely used drive cycles, and then compared and contrasted later with the results of real-world implementation. Daniel et al. [52] used a more sophisticated approach of constructing cumulative probability plots of various parameters such as vehicle speed and aggressivity for each of the legislative drive cycles. The level of agreement between these plots and similar plots based on observed data was then established using the coefficient of determination (R^2). Although this approach gives more insight than simply comparing average values it is not ideal because it assumes that individual variables such as speed, acceleration and power can be decoupled and considered independently from one another. In fact

rates of acceleration are likely to be highly dependent on the vehicle speed, for example.

It is clear that for a fair comparison to be made between a standard vehicle and a hybridised vehicle both should ordinarily be tested over identical drive cycles. In the case of a retrofit hybrid system however the situation must be more carefully considered. By adding an electric powertrain alongside the vehicle's standard ICE, without modifying the Engine Control Unit (ECU) to reduce the ICE power demand, the vehicle is effectively made *more powerful*. Since in the case of the Ashwoods Automotive system the ICE is unaware that the electric powertrain exists at all the driver has use of the full power of the ICE, augmented by the power of the EM. One may hope that in response to this additional power the driver would reduce the accelerator pedal position, thereby reducing engine load and saving fuel. For the driver the perceived effect of adding the hybrid system would be a minor re-mapping of the accelerator pedal response, so that the vehicle delivers slightly more power at each pedal position, and so the pedal needs to be activated slightly less. In reality however there is a risk that the driver would instead continue to use the ICE at a similar operating point, benefiting from the additional power provided by the EM for increased acceleration.

Use of the retrofit electric powertrain to increase rates of acceleration would completely undermine the addition of the hybrid system and not result in any fuel saving at all; it is therefore important that drivers are encouraged to limit their acceleration. One way in which driver behaviour and acceleration can be controlled is with a Driver Advisory System (DAS) which can give drivers feedback on their performance and help them to limit rates of acceleration; such devices will be discussed in the following section.

2.2 Driver Behaviour Improvement

It has been known for some time that driver behaviour has a significant effect on fuel consumption, which is reflected in the growing popularity of 'eco-driving' courses as fuel economy has become an increasingly important issue. As discussed, the retrofit

hybrid electric system studied in this work would likely require some monitoring of driver behaviour to achieve its potential. In view of this it seems only logical that such monitoring should also be carefully designed and its effects investigated with a view to achieving the maximum fuel savings that might be deliverable through this route, alongside the hybrid electric technology.

Studies have shown that suitable driver training can reduce fuel consumption by 10% on average [54, 55]. However, it has also been suggested that the long-term effects of such courses are less significant. Beusen et al. [56] followed a set of car drivers for 5 months before and 5 months after such a course and noted that the long-term effects varied between drivers, with around 20% relapsing to old habits. The authors of the study acknowledged that since the drivers volunteered for the course there was also likely to be some bias in the mentality of the drivers (many showed an increase in fuel efficiency in the months leading up to the course, prior to any training). It seems likely that relapse amongst an accurate sample of the population would be higher. A further problem with eco-driving training when applied to drivers of light commercial vehicles is that usually the driver does not pay the fuel bill, and as such may have significantly less motivation to save fuel. A similar study following bus drivers for a period of years found that 12 months after an eco-driving course fuel consumption was reduced by just 2% [57]. There is an apparent need to give continuous real-time advice to drivers to ensure they do not forget what they have learnt.

The idea of using a real-time driver feedback device to try to improve fuel economy is not new. Van der Voort et al. [58] conducted experiments where such a device encouraged drivers to keep the engine near its point of optimal efficiency, and demonstrated fuel savings of up to 23% in urban driving. However this was demonstrated on a driving simulator and based on a relatively small number of driving hours. Furthermore, despite stating that using specialised sensors which must be added to the vehicle should be avoided, the system used inputs such as steering angle and headway (gap to vehicle in front) which are not readily accessible on most vehicles.

More recently Wu et al. [59] showed that fuel savings during acceleration events (not over a drive cycle) of up to 31% could be achieved by encouraging drivers to follow an optimal acceleration profile. However these were shown using a driving simulator, and once again the system was quite complex, requiring information about headway as well as the state of traffic lights being approached. Furthermore, the human machine interface consisted of a colour bar representing good/bad levels of acceleration overlaid with a black line corresponding to the current rate of acceleration. Drivers were expected to alter their acceleration, moving the black line until it rested in the optimal region. Whilst this is fine in a simulator environment, the safety implications of having a driver concentrate on a moving display during transient events in the real world are questionable. It is likely the algorithms developed here are better suited to autonomous vehicle applications. This tension between safe driving and 'green' driving, and the volume of information made available to a driver was highlighted by Young et al. [55].

As part of the European FP7 project "ecoDriver" Nouveliere et al. [60] presented a more fully developed system with similarities to that developed by Wu et al. [59]. In this implementation Dynamic Programming was used to calculate the optimal future speed trace over a short time horizon, and the driver was encouraged to meet this. However, in order to derive the optimal speed trace access to GPS data, road speed limits and headway were all required. Information was fed back to the driver by means of a large ring around the speedometer (in the instrument cluster) which would change colour depending on whether the driver was above or below the calculated optimal speed. It is conceivable that this mechanism of feedback could be safely adopted as it is likely the driver may be generally aware of the colour in their peripheral vision without having to deliberately look at it frequently. As well as fuel consumption reduction this device placed considerable emphasis on potential for safety improvement, which was delivered primarily through a reduction of speeding events (time spent above the legal speed limit). Trials undertaken on-road with 6 drivers showed substantial average fuel savings of 7.5%, though this was achieved at the expense of journey times being increased by 8% which may not always be tolerated.

Van Driel et al. [61] also set out some guidelines and lessons learnt from development of such a device, suggesting amongst other things that integrating the system with the vehicle CAN-bus would eliminate the need for dedicated sensors, reducing complexity and cost.

Larsson and Ericsson [62] developed an acceleration advisory tool with a novel implementation, in that it provided feedback to the driver by adding resistance to the throttle pedal. Therefore if the system deemed that the driver was accelerating unnecessarily harshly it would make the throttle pedal more difficult to press. The results showed a significant reduction in throttle depression but no significant reduction in fuel consumption, and it was concluded that rate of acceleration is not the only parameter affecting fuel consumption.

The findings of Larsson and Ericsson [62] highlight the fact that in order to reduce fuel consumption by modifying driver behaviour it is important to first understand what behaviours affect fuel consumption and to define quantifiable metrics. This in itself is no simple task as there are many facets of driver behaviour, some of which will vary depending on driving conditions and drive cycle, and not all of which the driver will necessarily be willing to change. Ericsson [63] defined 26 parameters to characterise driving patterns, divided into level measures (for example average speed and average acceleration), oscillation measures (which describe 'jerkiness') and distribution measures (proportions of time spent at various operating points). Three oscillation measures were defined:

(1) *Frequency of maximum and minimum values*: This is calculated by finding the time between peaks and troughs in the vehicle speed trace, where the minimum speed difference between a peak and a trough is defined (for example 10 mph);

(2) *Integral of the square of the acceleration*: This is defined as $\frac{1}{n} \int a^2 \, dt$ where a is the vehicle acceleration and n is the number of time steps;

(3) *Relative Positive Acceleration (RPA)*: This is defined as $\frac{1}{x} \int v a^+ \, dt$ where v is the vehicle speed, x is the total distance, and a^+ is positive acceleration only.

Whilst level measures and distribution measures are good for quantifying behaviour over a drive cycle they rely on collecting a sample of data over a period of time and then reviewing it, as does oscillation measure #1. For this reason they are difficult to use as instantaneous measures of driver performance. A measure which can be calculated instantaneously is likely to be of more use in modifying driver behaviour real-time for two reasons: firstly, the driver is immediately aware of actions which negatively impact fuel economy rather than trying to relate statistics to their driving style retrospectively; secondly, the driver is not required to set aside time to analyse their feedback. A study conducted by Tulusan et al. [64] confirmed that corporate car drivers prefer instantaneous real-time feedback from an eco-driving device and find this more helpful in reducing their fuel consumption.

Taking the terms inside the integrals only, the square of the acceleration and the RPA may both be calculated instantaneously. The RPA has an additional strength over the square of the acceleration in that small accelerations at high speed may produce values equally as large as rapid accelerations at low speed. This dependence on speed compensates for the fact that vehicles are capable of less severe acceleration when at high speed because the engine has a power limit and mechanical power is proportional to speed. RPA was shown to have a strong positive correlation with fuel consumption [65].

Fomunung et al. [66] defined the same quantity (speed times acceleration) as the Inertial Power Surrogate (IPS), also defining a Drag Power Surrogate (acceleration times velocity squared). The IPS was shown to have a positive correlation with NO_X emissions. In an effort to quantify driver aggressiveness Ford Motor Co. later used a similar approach [67] to define a Power Factor, $Pf = 2va$. Power Factor was identified as a loose measure of inertial load, or change in kinetic energy, and the driver's total 'aggressivity' was defined as the root mean square of *Pf* over a journey.

2.3 Chapter Conclusions

Energy management in a hybrid electric vehicle is a complex control task which ideally requires knowledge of the future. Most of the literature in the area assumes

the vehicle possesses a level of hybridisation that is consistent with a system originally conceived as a hybrid powertrain, not with a retrofit system. For example, the EM rated power consistently accounts for between a quarter and a third of the total powertrain rated power and the ICE is commonly downsized to reflect the extra power available from the EM [17, 19, 20, 24, 29]. The exception to this is [16] where the EM accounted for approximately 15% of the total powertrain rated power, but was integrated pre-transmission and so is equivalent to a larger EM mounted post-transmission. For retrofit systems where the electrical powertrain is likely to be much smaller some of the functionality assumed, for example the ability to turn the ICE off and operate in 'electric only' mode, is not available. While the underlying concepts remain common this observation suggests there is a gap in the present research to examine good control of retrofit HEVs.

HEV applications represent a particularly intensive duty cycle for the electric powertrain because it is likely to be operated near its maximum power capacity a great deal of the time. Much of the previous research in the area is based entirely in simulation and so rarely acknowledges the practical issues and limitations arising from battery thermal management, cell SOC imbalances and motor thermal management, for example. For this reason the addition of a factor representing powertrain stress to the cost function may prove beneficial when applying the control strategy to a real vehicle, and improve the optimality of controllers developed in simulation when applied to real vehicles.

With regard to state-of-the-art control algorithms the approach achieving the greatest theoretical optimality while retaining the feasibility to be implemented in real-time is SDP, though in practice this may not be a great deal more effective than ECMS in reducing fuel consumption. Evaluation of the performance of SDP in real hybrid vehicles has thus far been extremely limited, and so there exists a research opportunity to progress this; however, to do justice to the algorithm a representative library of historic driving data would be required and the vehicle should be tested on a driving cycle consistent with these data.

Adding a retrofit hybrid electric system to a standard vehicle effectively makes the vehicle more powerful. To ensure that in the real world deployment of the system the

electrical power capacity is used instead of, and not as well as, the ICE a driver advisory system would most likely be required. The development and trial of such a system would also provide the library of historic driving data necessary to design a SDP control strategy. With this in mind this research will aim to further the current body of knowledge by examining optimal control design for a retrofit HEV, conscious of stress on the electric powertrain, and where the driver does not have the explicit goal of saving fuel.

2.4 Research Aims

In view of the opportunities for further research highlighted in the Literature Review the scope of the work is broken down into the following aims.

Aim 1: Design and test a driver advisory system to encourage eco-driving and restrict use of vehicles' maximum available power. Use the real-world driving data collected to characterise vehicles' normal use and select a representative drive cycle for testing HEV control strategies.

Aim 2: Develop an optimal SDP control strategy for the retrofit HEV which allows stress exerted on the electric powertrain to be controlled, and evaluate this controller's performance on-vehicle.

Chapter 3   Driver Behaviour

This section of the work may be regarded as a standalone contribution of its own merit, as well as an essential component in the development of the HEV control strategies. Development of a driver assistance system designed to reduce fuel consumption is described, and fleet trials of the system demonstrated fleet fuel savings of 7.6%. Driving data collected during the trials is statistically compared against common legislative cycles, with the conclusion that the LA92 cycle is the closest match to the vehicles' typical in-service use. Recorded data is then used to generate a gear shift schedule for use with the LA92, which will together form the test cycle for dynamometer testing of the HEV.

Parts of this chapter have been published in the following separate works:

C. Vagg, C. J. Brace, D. Hari, S. Akehurst, J. Poxon and L. Ash, "Development and field trial of a driver assistance system to encourage eco-driving in light commercial vehicle fleets," *IEEE Transactions on Intelligent Transportation Systems,* 14 (2), pp. 796-805, 2013.

C. Vagg, C. J. Brace, D. Hari, S. Akehurst, and L. Ash, "A driver advisory tool to reduce fuel consumption : SAE Technical Paper 2012-01-2087." *In: 5th International Environmentally Friendly Vehicle Conference,* 2012-09-09 - 2012-09-10, Baltimore, Maryland, USA.

In the scheme of this work this chapter will characterise the typical usage of the vehicle to be hybridised, which is crucial to achieving an optimal control strategy. The aim of this is to collect data which may be used to build a stochastic model of typical driving patterns, which will in turn form the basis of the optimal hybrid controller. As well as this the data collected will be compared against common legislative driving cycles to determine which is most representative of the vehicle's real-world usage. However, it is also noted that driver behaviour has a considerable effect on fuel economy, and that it may be possible reduce the vehicles' fuel consumption by encouraging drivers to adopt a less aggressive driving style, incorporating some techniques which are taught as 'eco-driving'. This is particularly important when considered alongside the fact that the retrofit hybridisation being considered effectively increases the installed power of the vehicle, and the desired effect would not be achieved if this additional power were simply exploited, and wasted, in more aggressive driving.

In light of these observations this chapter will describe the development of a retrofit real-time DAS designed to reduce fuel consumption by encouraging eco-driving in drivers of light commercial vehicles. Field trials of the system will then be described, with analysis of the driving patterns observed. Finally, the recorded data will be compared against legislative driving cycles to determine the most representative test cycle for the hybrid vehicle and controller to be evaluated over.

3.1 Driver Assistance System (Lightfoot)

With the intention that the DAS developed here should be viable as a commercial product, and potentially independent of the retrofit hybrid electric system, it must be designed such that it is:

a) Cheap: requiring a minimum of dedicated sensors, therefore also allowing quick installation;

b) Simple: such that its functionality is transparent to the driver and perceived as fair, and to avoid the need for calibration on different vehicle models;

c) Safe: demanding minimal active attention from the driver such that they are not distracted from the road conditions.

Focussing on commercial vehicle drivers offers greater opportunity in some respects, because a driver's performance can be fed back to their employers and so the system can assume some authority. Whereas drivers in their own vehicles may choose to ignore driving advice the subjects of this system may be obliged to improve their driving behaviour. On the other hand commercial drivers present additional challenges because they do not typically pay for the fuel they use, and so may be less motivated to reduce their consumption. Furthermore, the driver's obligation to obey the system increases the necessity for good design, and since its installation is most likely the decision of the employer and not the driver there is a risk of resentment if the driver perceives the advice of the system to be unfair.

3.1.1 System Design

The system implemented was based on the method set out by Fomunung et al. [66] to calculate the IPS real-time. As discussed in Chapter 2 this metric is one of relatively few which may be determined instantaneously, and has been shown by others to have a clear link to fuel consumption.

IPS is calculated instantaneously by multiplying the speed and acceleration, as shown in Figure 3-1, where the acceleration is derived from change in speed over
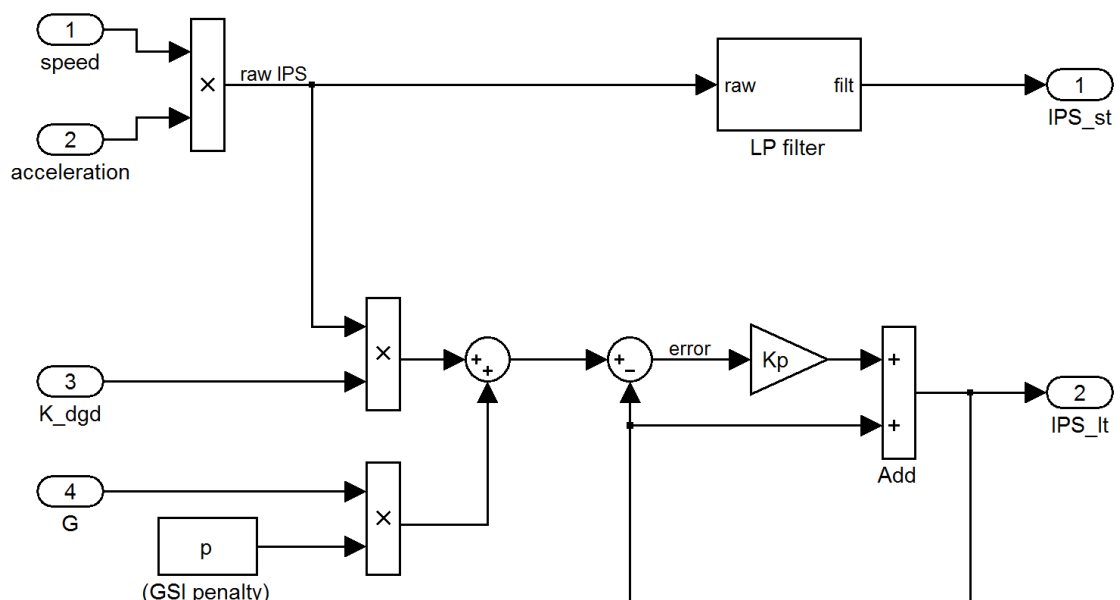

Figure 3-1: Schematic showing the calculation of short-term and long-term IPS, on which the core functionality of the DAS algorithm is based. Vehicle speed is the only essential input.

0.5s. The instantaneous IPS is passed through a low-pass filter with a time constant of 0.5s to smooth it; this is referred to as the Short-Term IPS (IPS_{ST}) and is displayed to the driver using a strip of 9 LEDs mounted inside the vehicle instrument cluster. Of the LEDs the first 4 are green, followed by 3 amber and then 2 red (Figure 3-2), to convey at a glance how economical the instantaneous driving behaviour is.

Monitoring a driver's behaviour directly with IPS_{ST} was found to be problematic as all drivers occasionally need to accelerate harshly. Some method of monitoring IPS over a longer period of time was therefore required, and this was achieved with the implementation of the Long-Term IPS (IPS_{LT}). As shown in Figure 3-1 the value of IPS_{LT} is calculated in a manner analogous to a P-control loop, where a modified value of the instantaneous IPS is the 'reference', IPS_{LT} is the process variable, and K_p the proportional gain.

The effect of this formulation is that IPS_{LT} tracks the instantaneous IPS, continually moving towards it. Calculation of IPS_{LT} is in fact directly analogous to passing IPS through a low-pass filter, however the implementation shown was preferred because of the ease with which signals could be intercepted and modified. The value of IPS_{LT} was the basis for determining whether a driver's behaviour was acceptable or not, and its inherent properties meant that brief episodes of extremely aggressive driving could be identified as well as moderately bad driving sustained over a longer period.

In calculation of the IPS_{LT} two notable modifications were made to the instantaneous value of IPS to refine the algorithm's behaviour: use of a degrade factor and inclusion of gear shift indicator (GSI) advice, such that the discrete-time transfer function of the IPS_{LT} logic may be written as

$$\text{IPS}_{LT[k]} = \text{IPS}_{LT[k-1]} + K_P \left(K_{dgd}\big(v_{[k]} \cdot a_{[k]}\big) + \big(p \cdot G_{[k]}\big) - \text{IPS}_{LT[k-1]} \right) \qquad (3\text{-}1)$$

where k is the discrete time index, K_{dgd} the degrade factor, p the value of the penalty for ignoring the GSI indicator, and G a Boolean flag indicating the application, or not, of this penalty. This calculation is performed at a 10 Hz task rate.

By multiplying the raw IPS signal by a degrade factor, $0 < K_{dgd} < 1$, the IPS could be artificially reduced during moderate accelerations in situations where it might be

unfair to penalise the driver. Specifically, two cases were screened for: long accelerations of moderate magnitude which might indicate accelerating up to a safe speed on a slip road, and accelerations at very low throttle which would suggest a steep negative road slope.

It was decided to incorporate gear shift advice since it is well know that reducing engine speed is one of the most straightforward ways of maximising the efficiency of modern automotive engines [68] by reducing air pumping losses. The GSI was integrated with the rest of the logic so that its advice could be enforced. Many modern vehicles are equipped with GSIs as standard – mainly for legislative reasons, and these often go unnoticed by the driver – and so the vehicle's built-in gear shift signal was used where this could be detected on the CAN-bus. For the case where OEM-designed gear shift advice was not available a simple logic was implemented which advised upshifts at 2200 rpm. In order to avoid inappropriate gear shift advice this shift flag was suppressed at high throttle positions where a steep gradient or an overtaking manoeuver was suspected.

The gear shift flag was conveyed to the driver with a green light marked "Shift Up" in the instrument cluster (Figure 3-2), however since it is not expected that drivers

Figure 3-2: Photograph showing a vehicle instrument cluster with Lightfoot during key-on initialisation. The Lightfoot add-on can be seen in the centre; LEDs indicating IPS_{ST} are below the standard vehicle computer screen, and those indicating IPS_{LT} and the GSI are positioned above the screen.

routinely watch the instrument display during driving the light was reinforced with a short beep to notify the driver when it illuminated. Upon illumination of the GSI drivers were given 1 s in which to change gear, after which a penalty p of 200 km^2/h^2s was added as an offset to the raw IPS (equation (3-1) and Figure 3-1), in turn causing the value of IPS_{LT} to rapidly climb.

In a similar fashion to IPS_{ST} the value of IPS_{LT} is displayed to the driver, this time using a 5 bar LED display following the green-amber-red convention (Figure 3-2). As already noted it is not expected that the driver will ordinarily look at this visual display and so the driver is also informed audibly when their driving style is deemed unnecessarily aggressive, as determined by the level of IPS_{LT}. Three thresholds are used such that the crossing of the first causes "Warning 1" to be issued audibly, the second causes "Warning 2" to be issued, and a failure of the driver to effect a reduction in IPS_{LT} causes a "Violation" to be issued if the third threshold is reached. These three thresholds were set at 100, 150 and 185 km^2/h^2s respectively, and the value of IPS_{LT} was saturated at the Violation threshold such that this could not be exceeded. The time taken for a warning to be issued based on GSI logic would depend on the value of IPS_{LT} when the GSI was activated, however because the penalty offset was higher than the Violation threshold ignoring the GSI advice would always ultimately result in a Violation, usually after around 25 s.

Although IPS inherently and logically depends on speed, and so for the same acceleration is larger at higher speeds, in practice this intrinsic adjustment was not found to result in a subjectively consistent feel to the logic across the speed range. This was understood to be because at low speeds, which tend to suggest urban driving, accelerations tend to be relatively brief. Conversely at higher speeds, which are indicative of rural driving, accelerations tend to be sustained over longer periods of time resulting in greater accumulation of IPS_{LT}. For this reason it was deemed necessary to vary the value of K_p with speed, using a value of 0.11 below 50 km/h and 0.07 above 75 km/h, with linear interpolation in the range 50-75 km/h; this logic resulted in better subjective fairness. It may also be obvious that K_p controls the rate at which IPS_{LT} falls as well as rises, and it was generally felt that the natural fall rate was too fast. In order to address this a rate limit was placed on IPS_{LT} such that the fall rate was limited to -3 km^2/h^2s^2.

3.1.2 Implementation of Embedded Code

The logic described in Section 3.1.1 was implemented in SIMULINK and the Real-Time Workshop add-on was then used to auto-generate C code. This source code was integrated with other low-level functionality by engineers at Ashwoods Automotive, to create a project structure which could be compiled and run as embedded code on a target microprocessor. In this work a Microchip PIC32 microcontroller was used to run the driver behaviour logic as well as the hybrid control strategy described in Chapter 4.

3.1.3 Field Testing

For assessment of the advisory system's efficacy in helping drivers to reduce their fuel consumption field trials were undertaken with 7 companies in the UK, with the device fitted into a total of 15 vehicles. Companies used in the trials represented a range of business sectors from mail delivery to environmental site surveyance, however the majority were providers of delivery or technical support services primarily operating in urban environments. During the trial baseline data were recorded from each vehicle for a period of approximately 2 weeks, of which the drivers were unaware and so would have driven naturally. Data collection was achieved by installing the electronic hardware of the DAS equipped with data logging capabilities, but without fitting the display and with audio feedback disabled. Following this period the display was fitted, audio feedback was enabled and the trial run for approximately 2 further weeks with the system active. Throughout the trials essential data were recorded from the vehicle CAN-bus at 10 Hz via the On-Board Diagnostics (OBD) port and stored to a SD card. These data included the most essential inputs and outputs of the ECU such as vehicle speed, accelerator pedal position, engine speed, engine load, engine coolant temperature, and fuel injection rate, allowing insight into the behaviour of both the vehicle and driver.

Details of the companies involved in the trial can be seen in Table 3-1 as well as the total duration of each phase of the trial. The number of active days shown are the number of days during which the vehicles were in active use, since they tended not to be used every day of the week; these numbers are approximate as usage may have varied slightly between vehicles within each company.

Table 3-1: Details of companies and vehicles involved in the trial.

Company	Nº Vehicles	Vehicle Use	Baseline duration days (active days[a])	Live duration Days (active days[a])
A	3	Technical call-out service	14 (13)	14 (13)
B	3	Retail parts delivery	15 (13)	14 (12)
C	2	Fresh produce delivery	11 (9)	14 (12)
D	2	Technical call-out service	14 (10)	22 (16)
E	2	Site visits	16 (10)	22 (13)
F	2	Technical call-out service	14 (9)	18 (12)
G	1	Support service	14 (12)	16 (14)

[a]Active days are the days during each phase of the trial where the vehicle(s) were in active use.

The total duration of each trial was about 4 weeks; this was considered to be long enough to collect representative data from each vehicle which is desensitised to daily or weekly fluctuations in usage pattern, cargo load, traffic or weather, while short enough not to be affected by seasonal changes; thus the effects of these confounding factors were minimised. Each vehicle was normally paired with only one driver and so a comparison between vehicles is also a comparison between drivers. All of the vehicles involved in the trial were Ford Transit vans of Euro IV emissions stage specification built from 2008-2011, further details of which can be found in Table 3-2.

For the calculation of vehicle fuel consumption a reliable measure of fuel consumed is clearly essential, and for this purpose the logged ECU engine fuelling rate was used. This measure was regarded as precise but not accurate, that is each ECU may have a small calibration offset, but this should be constant. With this in mind it is reasonable to calculate percentage fuel savings of each vehicle and to compare these, but the comparison of absolute fuel consumption figures between vehicles should be cautioned against.

3.1.4 Results

During the trials a total of 1,107 hours of real world driving data were recorded, covering 39,300 km and comprising 5,587 individual trips. In the greater context of this work these data offer a wealth of information pertaining to typical vehicle usage patterns and driver behaviour which is essential knowledge for good design of hybrid vehicle control strategies; however it is also worthwhile examining the fuel savings achieved by the system.

Table 3-2: Vans involved in the trial and the idle-corrected fuel savings achieved by each.

Company	Van	Vehicle Type	Total Distance Covered (km)	Baseline Fuel Consumption (L/100km)	Live Fuel Consumption (L/100km)	Fuel Saving (%)
A	1	N/A	2375.2	8.20	7.44	9.16
	2	260S 2.2L 5-speed	1635.0	9.12	9.09	0.43
	3	260S 2.2L 5-speed	1402.7	9.08	7.99	12.03
B	4	350E 2.4L 6-speed	7666.1	10.12	9.00	11.08
	5	280S 2.2L 6-speed	6604.2	7.99	7.20	9.89
	6	350L 2.4L 6-speed	2613.4	9.74	9.57	1.75
C	7	350L 2.4L 6-speed	3727.9	11.00	10.01	8.97
	8	300M 2.2L 6-speed	2918.3	10.43	9.20	11.82
D	9	350M 2.2L 6-speed	1740.7	9.92	9.08	8.48
	10	350M 2.2L 6-speed	1737.5	9.18	8.86	3.44
E	11	350M 2.4L 6-speed	2065.3	11.04	10.98	0.49
	12	350L 2.4L 6-speed	2355.3	11.14	10.83	2.77
F	13	300S 2.2L 5-speed	917.1	9.05	8.86	2.08
	14	300S 2.2L 5-speed	507.3	10.77	10.31	4.29
G	15	280S 2.2L 5-speed	1036.0	9.74	8.97	7.91
			Average fuel saving (weighted by distance travelled per vehicle)			*7.61*

Fuel consumption results for each vehicle are shown in Table 3-2 for each vehicle, as well as the total fleet savings calculated as the average of each vehicle's savings weighted by distance covered. Overall it can be seen that the fleet fuel consumption was reduced by 7.61% - an extremely encouraging figure carrying considerable environmental and monetary value. It is also interesting to note that there is a large range in savings, from 0.43% to 12.03%. This range is presumably because of a difference in the drivers' aggressivity to begin with meaning that some had more potential for improvement than others, though possibly also because the efficacy of the device depends somewhat on the drive cycle a vehicle is used over.

In order for the reported fuel savings to be meaningful it is important that the vehicles usage was broadly similar during each phase of the trial. Simple analysis using cumulative probability plots showed that on the whole the usage was comparable, however it became apparent that the data for some vehicles showed inconsistencies in the amount of time spent stationary with the engine idling. Further analysis showed that whilst the great majority of idling events were of less than 60 seconds, as might be expected, a small subset were far longer, with the longest being 2.5 hours. This was problematic for the analysis because although relatively little fuel is consumed at idle this can become significant over a long period of time, and since no distance is covered this can cause skew in the fuel consumption per unit
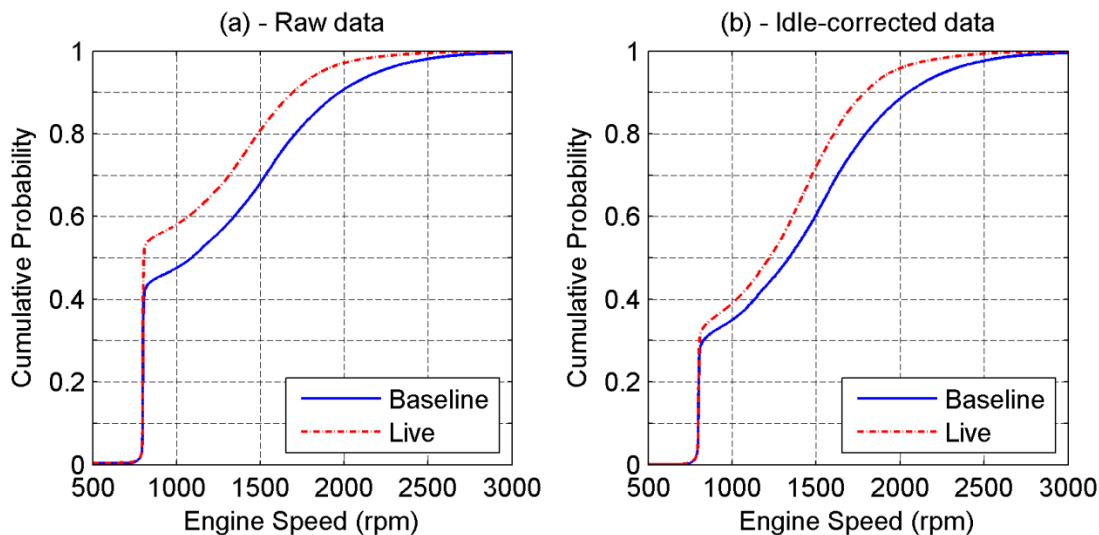
Figure 3-3: Cumulative probability distribution of engine speed for Van 14 before removal of exceptional idling instances (a) and after (b). The 10% discrepancy in idling time between Baseline and Live phases of the trial could have skewed the results, and is reduced to 3% in the idle-corrected data set.

distance (L/100km for example). For this reason all idling events longer than 90 seconds, representing the 97[th] percentile, were regarded as outliers and removed from the analysis; the data presented in Table 3-2 have been adjusted for idling time in this way.

Correction of the data to remove exceptionally long idling instances resulted in satisfactory parity between the Baseline and the Live data sets. The most acute example of idling time discrepancy in the raw data is shown in Figure 3-3(a), along with the corrected data in Figure 3-3(b). This figure illustrates the perhaps surprising proportion of time for which light commercial vehicles are at idle, and suggests that the introduction of start-stop systems which are now common in new vehicles may have great potential for reducing fuel consumption in real world driving. Figure 3-3 also indicates a significant shift towards lower engine speeds during the Live trial when compared against the Baseline for this vehicle, a phenomenon which will be further discussed with respect to the complete fleet later (Figure 3-7).

To illustrate the functioning of the algorithm and the effect of the device on driver behaviour Figure 3-4 and Figure 3-5 are provided as exemplary excerpts of urban driving before and after the device was activated. These data are recorded by the same driver completing similar journeys with roughly the same average speed. As explained in Section 3.1.1 periods of acceleration cause peaks in IPS_{ST}, generally

causing IPS_{LT} to rise. With the system active the driver has been encouraged to reduce the magnitude of acceleration events with respect to the baseline condition, and the IPS_{LT} is correspondingly reduced.
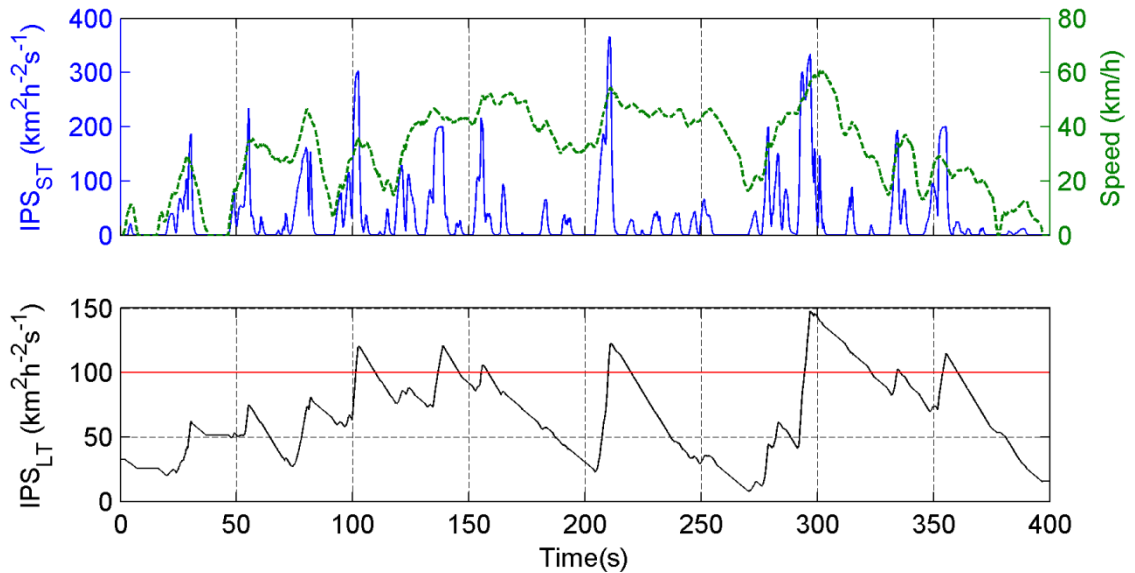
Figure 3-4: Example of urban driving during the Baseline phase of the trial, and the accompanying response of IPS_{ST} and IPS_{LT}. The solid red line at 100 $km^2h^{-2}s^{-1}$ on the IPS_{LT} graph indicates the level of the Warning 1 threshold; this driving would therefore have warranted several warnings, though the driver would have been unaware.
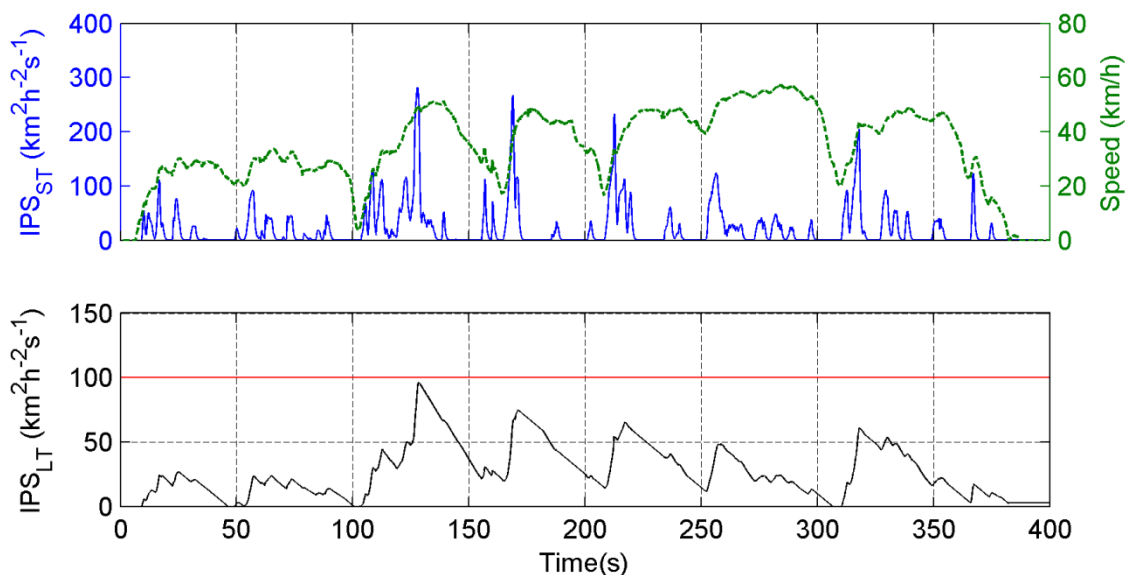
Figure 3-5: Example of data collected from the same van as in Figure 3-4, during the Live phase with the system active. Note the smoother driving style has resulted in lower values of IPS_{ST} and correspondingly lower values of IPS_{LT}.

Table 3-3: Changes in average operating parameters with *Lightfoot* active.

	Fuel Consumption (L/100km)	Engine Load (%)	Engine Speed (rpm)	Accelerator Position (%)
Baseline	9.68	36.65	1575	17.45
Live	8.94	39.63	1412	15.54
Change (%)	*-7.61*	*+8.14*	*-10.36*	*-10.93*

Analysis of the instantaneous algorithm response is helpful in understanding the interaction of the system with the driver, however to gain insight into the change in driving style effected by the system it is more helpful to consider the higher-level statistics. Table 3-3 shows the changes in average fuel consumption, engine load, engine speed, and accelerator pedal position during the Live phase of the trial, which give a good indication of how the drivers' behaviour was altered. In calculating these figures the data from all vehicles was regarded as one continuous data set, thus the figures represent the 'fleet average', with vehicles having recorded the most usage contributing more to the overall average; this is not the same as the average which any one driver, selected at random, might expect to see.

The observed drop in average accelerator pedal position by 10.93% is an indication of the less aggressive driving style being adopted, and the reduction in average engine speed is the result of gear shift advice. It is interesting as an aside to note that the average engine load was actually increased, most likely because the shift of engine operating point toward lower speeds requires an increase in torque to deliver comparable tractive power. The observed reduction in fuel consumption is understood to be a result of both smoother driving and reduced engine speed. A reduction in the average positive IPS value corresponds to a decrease in the tractive work consumed per kilometre driven. In contrast a reduction in average engine speed would imply a shift towards higher torque operation for the same net power output, moving the engine operating point towards the region of the speed-torque envelope where automotive engines are typically most efficient. There are therefore at least two mechanisms by which fuel could be saved – one pertaining to how much energy is required, and the other how efficiently this energy is supplied. Further insight into the two mechanisms of fuel saving is offered by consideration of Figure 3-6 and Figure 3-7, which compare the probability distributions of raw IPS and engine speed during the Baseline and Live phases of the trial.
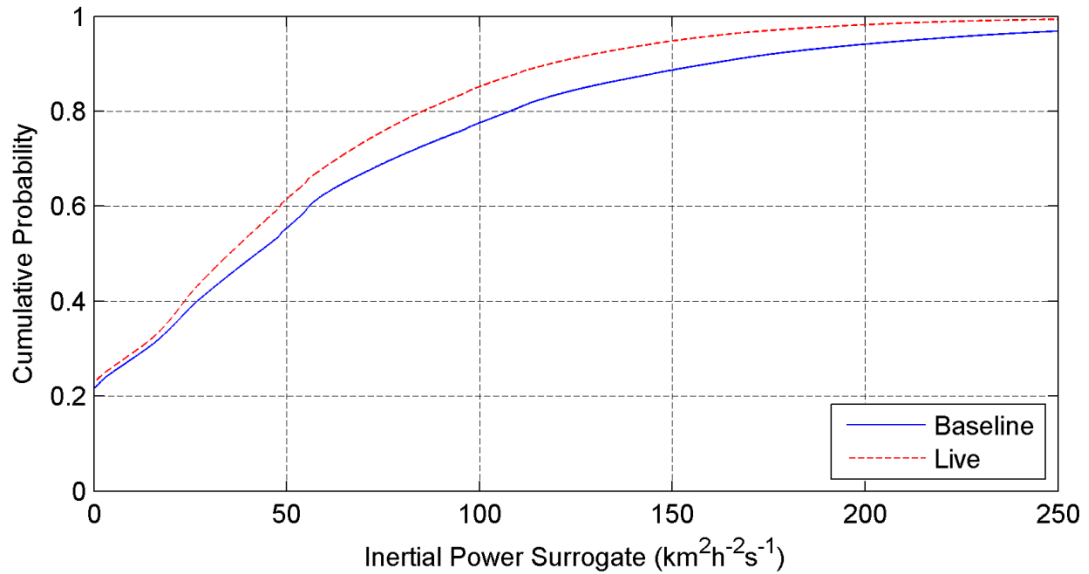
Figure 3-6: Cumulative Distribution Functions for positive IPS during the Baseline and Live phases of the trial. Low values of IPS are more common in the Live phase and the frequency of IPS>200 km²/h²s is halved.

Probability data in Figure 3-6 ignore periods where the vehicle is stationary, as well as when the value of IPS is negative. The data therefore represent the distribution of positive IPS while the vehicle is moving, and the y-axis intercept represents the proportion of time spent at cruise. The values of IPS used in generating the plot are not those which would have been calculated by the algorithm real-time, but are instead calculated off-line over the recorded speed trace. The acceleration across each time sample is therefore calculated using the adjacent speeds (at the previous and subsequent timesteps), $a(t) = \frac{v(t+1)-v(t-1)}{2 \cdot \Delta t}$, which gives a more precise and less noisy value for instantaneous acceleration than is possible to calculate in real-time.

Examination of Figure 3-6 reveals that the IPS distribution has shifted considerably away from high values and towards the mid-range. For example, the frequency of IPS exceeding 200 km²/h²s, which corresponds to the threshold of the 7th LED (3rd amber), has been reduced by half (from 11.0% to 5.4%), while moderate values of IPS have become more frequent. In contrast the frequency of low values of IPS below 50 km²/h²s is relatively unchanged, and the proportion of time spent at cruise (indicated by the y-axis intercept) is almost identical. This is the clearest indication that the IPS logic achieved the design intent of discouraging aggressive acceleration in favour of more moderate accelerations.

To study the effect of the GSI logic Figure 3-7 presents histograms of engine speed during the trial and shows that drivers' choice of engine operating point changed substantially. During the Live trial the advice of the GSI resulted in drivers upshifting earlier, greatly reducing the proportion of time spent above 1700 rpm and encouraging use of the engine around 1450 rpm. In both phases of the trial an anomaly is evident in the region 2100-2200 rpm where the frequency density is very high; this engine speed typically corresponds to 96-98 km/h in top gear, and many of the vans were electronically limited to 100 km/h, so it is likely that this discontinuity results from high-speed cruising.

In the context of determining the success of *Lightfoot* it is important to remember that in the commercial sector the financial gains achieved through fuel savings may be negated if journey times were increased at all. It is perhaps tempting to suppose that reduced rates of acceleration result in slower speeds, and therefore longer journey times. In fact this logic is flawed because during normal driving it is unusual for speed to be unrestricted and so journey time is more often determined by traffic conditions than by how quickly one can accelerate; the opposite is of course true on a race track. This logic is supported by the observation that the average speed of all vehicles during the Baseline phase of the trial was 38.59 km/h, which actually *increased* slightly to 38.75 km/h during the live phase, despite reduced rates of
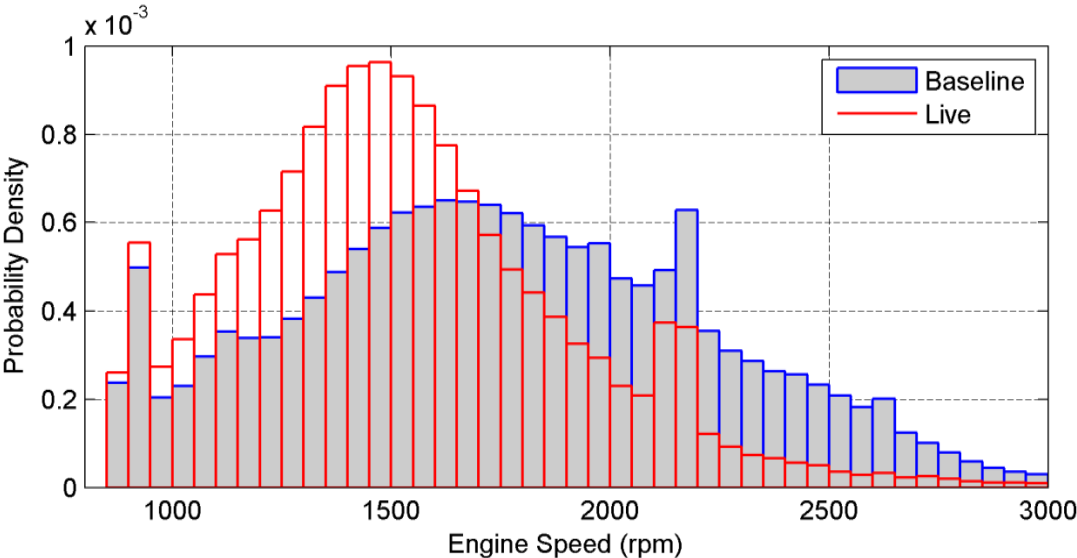

Figure 3-7: Engine speed probability density during the Baseline and Live phases of the trial. A considerable shift towards lower engine speeds is evident during the Live phase. Bars showing engine speed below 850 rpm have been omitted for clarity.

54

acceleration. It is difficult to comment on the statistical significance of such a small difference, however it suffices to say that the introduction of *Lightfoot* appeared not to have adverse effects on average speeds, and therefore journey times, and so the system shows strong potential for commercial application.

3.2 Drive Cycle Selection

The aim of this work is to develop a hybrid vehicle control strategy which is optimal for the usage patterns to which the vehicles are likely to be exposed. In order to evaluate the performance of the eventual control strategy it is therefore essential that the test cycle over which it is assessed is representative of this same usage, otherwise the evaluation phase will not give an accurate impression of the controller's behaviour under real world conditions. In view of this an important part of this work is to define a representative test cycle over which to evaluate the performance of the hybrid electric system and its controller.

Whilst the huge quantity of real world drive cycle data collected would have made it possible to design a bespoke cycle from scratch this would have been an enormous undertaking in its own right, and would also make the results less transparent to others not familiar with this new drive cycle. For these reasons the preferred method was to compare the real world data against a broad range of widely used legislative test cycles, and to select the cycle most similar.

Analysis of driving patterns is a complex activity, since there are a great many facets and parameters which can be used characterise a cycle, and since it is unlikely to be possible to perfectly align all of these it is necessary to focus on matching some while compromising on others. The most basic of approaches to this problem is to consider a table of average or extrema parameters – such as vehicle speed, acceleration, and power – and to try to match these between the collected data and a drive cycle. This may serve as a very approximate tool, but is far from ideal as an average gives only extremely limited insight. A more advanced approach which has been used by some [52] is to consider the cumulative frequency distributions of these same parameters, with the aim of matching the curve shapes. Although a considerable improvement this still leaves open substantial margin for inadvertent erroneous matching, since considering the frequency distribution of each parameter

in isolation is an unfair simplification. For example, an acceleration of 2 m/s^2 at low speed (e.g. pullaway) may be regarded as quite modest; however the same acceleration at motorway speeds would imply a much larger power use and be regarded as extremely aggressive driving. Decoupling speed and acceleration and considering each separately does not therefore give sufficient insight into the aggressivity of the driving cycle.

3.2.1 SAFD Analysis

The approach taken here was to define a two dimensional discrete state space consisting of speed and acceleration. By placing each sample point into its appropriate bin a two dimensional histogram, or surface, was constructed representing the speed-acceleration frequency distribution (SAFD) of the cycle. The surface represents the probability distribution of the driving cycle in the speed-acceleration plane, and so the volume under the surface has a sum of exactly one. The shape of a SAFD surface gives detailed insight into driving patterns, where the speed and acceleration distributions are not artificially decoupled, but instead may be considered simultaneously to build a better picture of the duty cycle and its aggressivity. Though it is not often employed the benefits of this method of analysis have been understood for some time [69].

SAFD data were calculated for each vehicle involved in the trial as well as for a range of common legislative driving cycles, and examples of the resulting surfaces are presented in Figure 3-8 and Figure 3-9. In the probability calculations a linear interpolation regime was used so that data samples not lying exactly on a grid node of the state space could be apportioned appropriately between the surrounding nodes, rather than simply rounded to the nearest. Interpolation was particularly important in the analysis of the set of legislative tests to ensure the resulting surfaces were not unduly jagged because of the relatively few data points in any one drive cycle. The idle condition, where speed and acceleration were exactly zero, was not included in the analysis as it otherwise dominates. Since the hybrid system is inactive at idle it is not a particularly important condition to consider in the selection of a representative driving cycle.
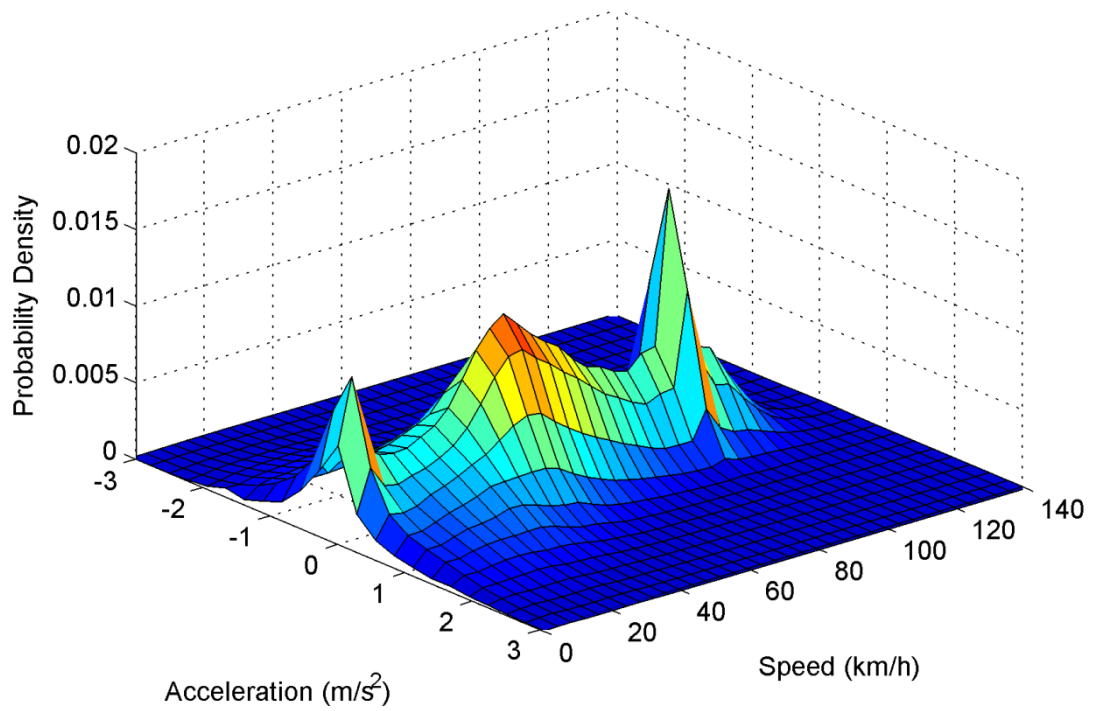
Figure 3-8: SAFD plot for the combined data of vans 4, 6 and 7 with *Lightfoot* active.
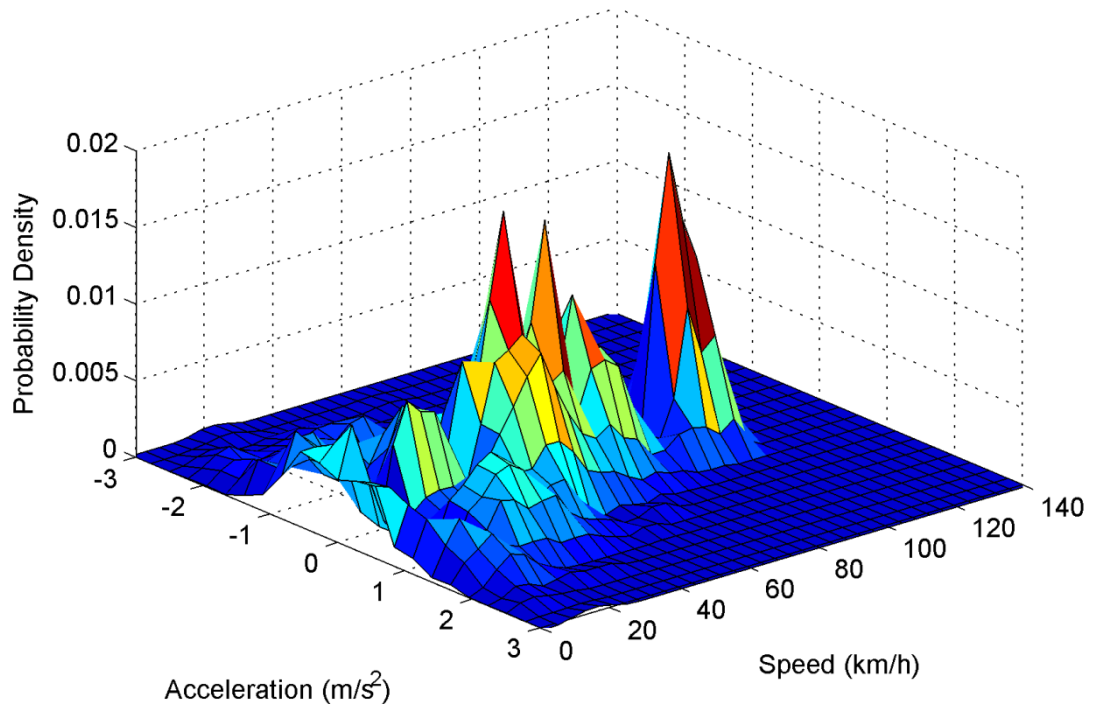

Figure 3-9: SAFD plot for the LA92 legislative drive cycle, sampled at 10 Hz.

Having obtained SAFD data for each legislative drive cycle these were compared against the real world data collected from each van, both with and without *Lightfoot* fitted. Using equation (3-2) a *percentage agreement* was obtained for each combination in an effort to determine which test cycle was the closest statistical match to the real world usage patterns observed.

$$\% \text{ agreement} = 100 \times \left(\frac{2 - \sum|SAFD_1 - SAFD_2|}{2} \right) \qquad (3\text{-}2)$$

In the calculation $SAFD_1$ and $SAFD_2$ are the two frequency distributions to be compared. The difference between their elements is summated and subtracted from 100%. Note that a factor of two is necessary since two SAFDs with no overlap at all would have a summative difference of two. An alternative calculation can be done using a chi-square (χ^2) test, however this process squares the differences and it was not felt there was any particular justification for this in the analysis. The percentage agreement between each of the legislative drive cycles and the vans is shown in Table 3-4 for the case without *Lightfoot*, and in Table 3-5 and Figure 3-10 for the case with *Lightfoot*.
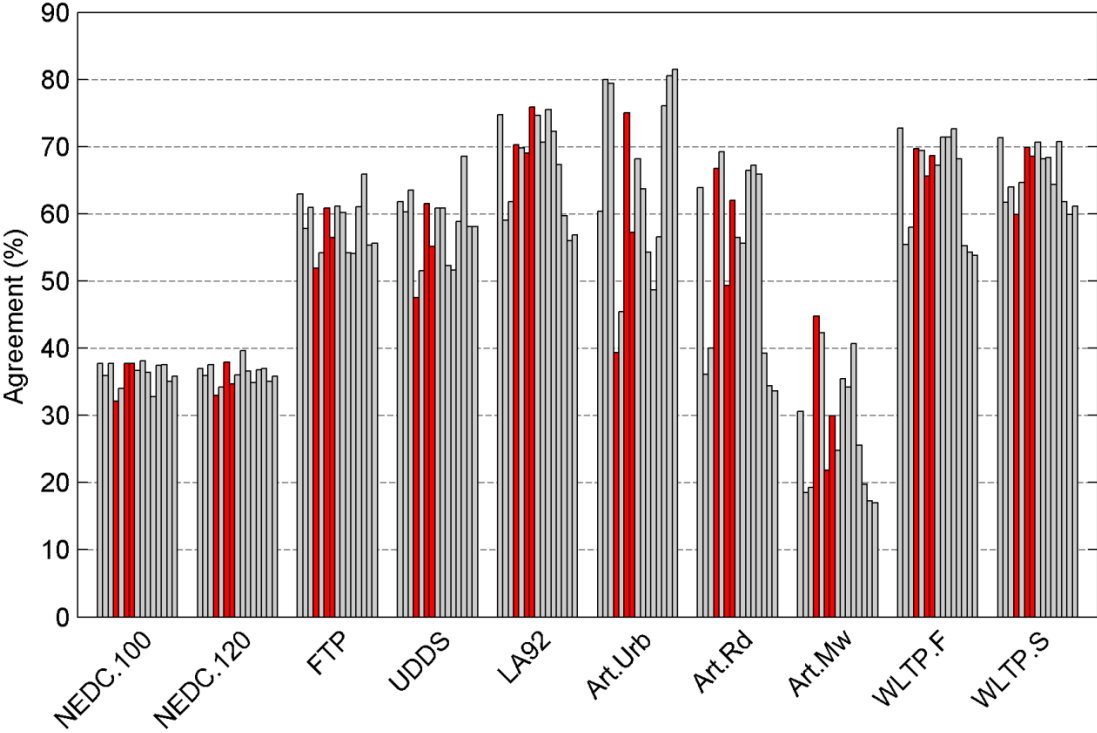

Figure 3-10: Percentage agreement between common legislative cycles and each vehicle data set with Lightfoot active. Data corresponding to van numbers 4, 6 and 7 are highlighted in red since these have the gear ratios most representative of the test vehicle.

Table 3-4: Percentage agreement between common legislative driving cycles and each vehicle data set without Lightfoot. The best agreement in each row is highlighted in bold.

Van #	NEDC$_{100}$	NEDC$_{120}$	FTP	UDDS	LA92	Art.Urb	Art.Rd	Art.Mw	WLTP.F	WLTP.S
1	38.1	38.0	64.5	64.4	74.2	66.7	59.1	29.0	70.6	70.5
2	37.3	37.2	58.1	59.9	64.2	77.2	42.6	23.9	60.1	63.2
3	35.3	36.1	57.7	59.1	64.8	76.7	44.0	25.8	61.0	61.9
4*	28.4	30.7	42.5	40.7	62.6	34.2	58.3	56.6	**62.9**	50.5
5	32.9	34.1	53.4	51.7	69.7	48.9	66.7	41.7	68.5	62.2
6*	36.9	37.3	59.9	60.5	70.4	**73.0**	51.0	24.0	66.7	67.9
7*	35.7	33.5	53.4	52.4	**74.6**	57.0	60.1	30.7	67.4	65.7
8	34.8	33.8	55.1	54.9	72.8	64.4	53.5	27.1	64.5	64.4
9	38.2	39.2	61.1	62.4	67.6	75.7	48.3	25.4	65.2	67.2
10	33.5	33.9	47.2	47.1	68.5	53.4	51.9	38.3	63.8	58.4
11	32.9	34.5	52.8	50.2	71.7	45.9	66.9	43.4	72.0	62.8
12	38.3	37.6	63.3	61.8	65.7	58.2	64.7	25.3	67.8	71.0
13	36.5	36.5	64.8	67.3	58.1	77.2	37.8	17.6	53.7	60.4
14	35.0	35.0	57.5	60.3	57.4	81.2	34.7	17.5	53.0	59.0
15	34.7	34.7	52.1	54.9	55.8	81.2	32.4	17.8	53.4	59.1
Mean- 3 vans (*)	33.7	33.8	51.9	51.2	**69.2**	54.7	56.5	37.1	65.6	61.4
Mean- All vans	35.2	35.5	56.2	56.5	**66.5**	64.7	51.5	29.6	63.4	62.9

Table 3-5: Percentage agreement between common legislative driving cycles and each vehicle data set with Lightfoot active. The best agreement in each row is highlighted in bold.

Van #	NEDC$_{100}$	NEDC$_{120}$	FTP	UDDS	LA92	Art.Urb	Art.Rd	Art.Mw	WLTP.F	WLTP.S
1	37.7	37.0	62.9	61.8	74.7	60.4	63.9	30.6	72.7	71.3
2	35.9	35.9	57.8	60.3	59.0	80.0	36.1	18.5	55.4	61.7
3	37.7	37.5	60.9	63.5	61.8	79.4	40.0	19.3	58.0	64.0
4*	32.1	33.0	51.9	47.5	**70.3**	39.3	66.7	44.8	69.7	59.9
5	34.0	34.2	54.2	51.5	69.8	45.4	69.2	42.3	69.4	64.6
6*	37.7	37.9	60.8	61.5	69.0	**75.0**	49.3	21.8	65.6	69.9
7*	37.7	34.7	56.5	55.1	**75.9**	57.2	62.0	29.9	68.6	68.5
8	36.7	36.0	61.1	60.8	74.6	68.2	56.5	24.8	67.2	70.6
9	38.1	39.6	60.2	60.8	70.6	63.7	55.6	35.4	71.4	68.2
10	36.4	36.6	54.2	52.3	75.5	54.3	66.5	34.2	71.4	68.4
11	32.8	34.9	54.1	51.6	72.3	48.7	67.2	40.7	72.6	64.4
12	37.4	36.8	61.0	58.8	67.3	56.6	65.9	25.5	68.2	70.7
13	37.5	37.0	65.9	68.5	59.7	76.1	39.2	19.7	55.2	61.8
14	35.1	35.1	55.3	58.1	56.0	80.5	34.4	17.3	54.3	59.9
15	35.8	35.8	55.6	58.1	56.8	81.5	33.6	17.0	53.8	61.1
Mean- 3 vans (*)	35.8	35.2	56.4	54.7	**71.7**	57.1	59.3	32.2	68.0	66.1
Mean- All vans	36.2	36.1	58.2	58.0	**67.6**	64.4	53.7	28.1	64.9	65.7

Key to non-standard abbreviations:

NEDC$_{120}$	Standard NEDC	Art.Urb	Artemis Urban
NEDC$_{100}$	NEDC limited to 100 km/h	Art.Rd	Artemis Road
WLTP.F	WLTP "fast" version	Art.Mw	Artemis Motorway
WLTP.S	WLTP "slow" version		

Comparison of driving patterns using SAFDs considers only the distributions of speed and acceleration. When constructing the stochastic driving distributions which the SDP hybrid vehicle controller will be based on (Chapter 5.2) it will be extremely important these incorporate representative and consistent use of vehicle gears, since for a fixed driving cycle it is the gear selection which determines the engine speed-load operating point. Considering the data from all of the vehicles involved in the trial would not result in consistent use of the gears, because when purchasing a LCV customers often have a choice of gearing, achieved through different final drive ratios, allowing the same vehicle model to be tailored for different purposes. Use of data from all of the vehicles may therefore introduce a large variance in gear shift speeds. A further consideration for implementation of the SDP hybrid vehicle controller is that the vehicles used to construct the stochastic gear shift data ought to have similar overall gear ratios to the test vehicle on which the control strategy is physically implemented and tested; otherwise the speed-load operating points of the engine will be fundamentally shifted.

Of the 15 vans analysed three have identical gear ratios to one another, and are extremely closely matched to the test vehicle. For this reason these three vans - numbers 4, 6 and 7 - are of particular interest. By selecting data from these vans to build the stochastic drive cycle model it is ensured that the test vehicle will be exercised in a manner most similar to real world conditions.

Whether considering the data from all of the vans or from the three highlighted as of particular interest, and whether with or without Lightfoot, the LA92 drive cycle is unanimously the best fit. This was therefore selected as the drive cycle used in the development and evaluation of hybrid control strategies both in simulation and for dynamometer testing.

3.2.2 Gear Shift Schedule Design

Following the selection of the LA92 speed trace, a final step was required to fully define the drive cycle. Unlike the NEDC, Artemis cycles and others, the LA92 does not have a standard time-based gear shift schedule which can be used with manual transmissions. Instead, the cycle requires that a custom gear shift schedule be

designed according to the same regulatory framework set out for the FTP drive cycle in CFR Title 40 §86.128-00 [70]. This framework does allow for gear shifts to be effected as a function of road speed or engine speed, however the speeds defined are now extremely unrepresentative of modern vehicles with more than 5 gears and so the alternative method using an empirical "Shift Speed Survey" described in Advisory Circular (A/C) No. 72A [71] was used instead. The procedure set out in A/C 72A considers the road speed at which gear shifts occur during acceleration events and in cruise conditions during normal driving. For acceleration events the mean shift speed is used; the calculation of shift speeds during cruise conditions is a little more elaborate however:

> "For any particular gear, the cruise speed should be the speed at which the number of cruise shift data points … in lower gears at higher speeds is exactly offset by the number in that and higher gears at lower speeds. For example, the fourth gear cruise speed can be determined by adding the number of points in fourth and fifth gears starting at the lowest speed, adding the number of points in first through third gears starting at the highest speed, and selecting the speed where the two sums become equal. This speed can be graphically determined by plotting the cumulative points in fourth (and higher gears) starting at the lowest speed and the cumulative points in third (and lower gears) starting at the highest speed. The speed at which the two plots intersect is the speed at which the cumulative points offset one another."

– A/C 72A, pp10.

An applied example of this method is shown in Figure 3-11. The only deviation from the method set out was that the three vans were not weighted in any way to account for differences in the number of shift points recorded by each. Instead the three data sets were treated as though they were one continuous data set, therefore reflecting the relative use of each vehicle.

The shift speeds derived using these legislative procedures are shown in Table 3-6 and these were applied to the LA92 speed trace to complete the drive cycle. Having applied the procedure two further filtering operations were deemed necessary in order to remove anomalous gear shift points and ensure the resulting gear shift schedule was realistic:

1. Where a gear was selected, and then the previous gear re-selected only one second later, the gear change events were regarded as noise and ignored.

2. During harsh accelerations it is possible that the acceleration shift speed values dictate an upshift every second. This behaviour was not regarded as realistic, and the shortest time in which the driver could perform an upshift while also delivering tractive power to accelerate the vehicle was felt to be two seconds. In these cases, therefore, upshifts were delayed as necessary to ensure that each gear was engaged for at least two seconds.

The final time-based gear shift schedule used is made available in Appendix 1.

Table 3-6: Gear shift speeds determined from road survey, according to the procedure set out in Advisory Circular No. 72A.

	Gear 2	Gear 3	Gear 4	Gear 5	Gear 6
Acceleration Shift Speed (km/h)	14.1	28.7	39.7	51.3	62.4
Cruise Shift Speed (km/h)	12.0[†]	25.0	36.9	48.0	58.9

[†] The calculated cruise shift speed here was 8.0 km/h, which would result in engine speeds less than 700 rpm. This was therefore replaced with 12.0 km/h, corresponding to an engine speed of 1000 rpm.
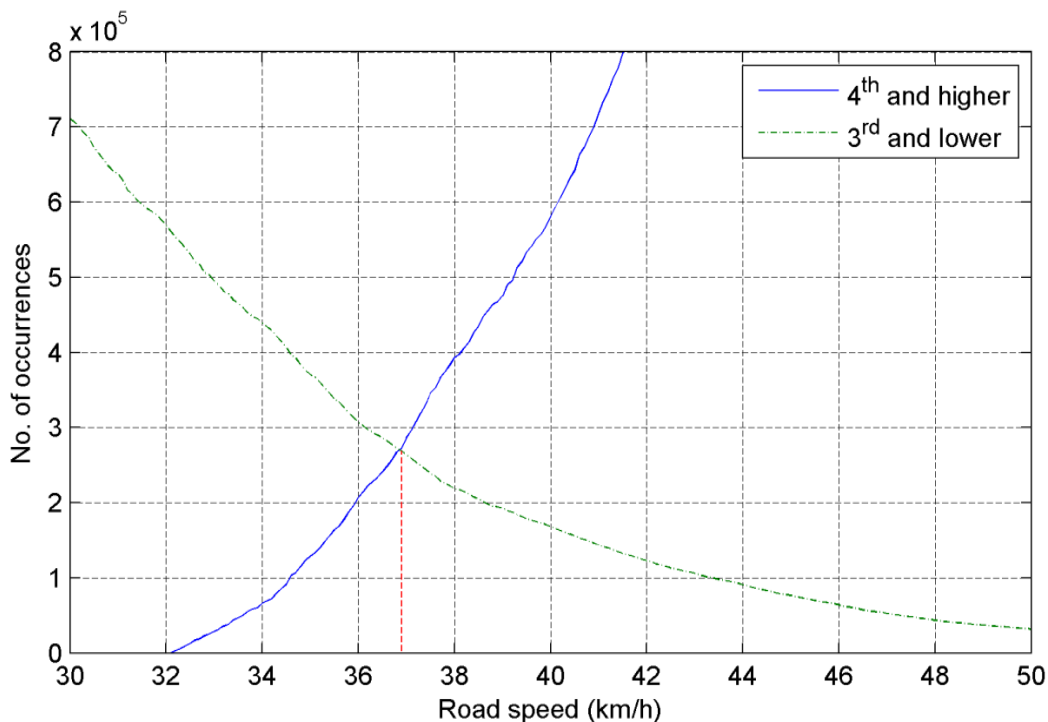

Figure 3-11: Determination of cruise shift speed for gear 4. The number of shifts into gears 4 and higher exceeds those into gears 3 and lower at 36.9 km/h. This is therefore the cruise shift speed used.

3.3 Chapter Conclusions

In this chapter the development and field trials of a driver assistance system have been presented. During fleet trials of the system it was shown to reduce fleet fuel consumption by 7.6% through reduced accelerations and encouragement of early gear upshifting resulting in considerably lower engine speeds.

Naturalistic driving data collected during the trials has been compared against legislative drive cycles, resulting in the LA92 cycle being identified as the most representative of the vehicles' typical in-service usage. For this reason the LA92 is selected as the reference drive cycle for testing and evaluation of the hybrid vehicle energy management strategies developed in Chapter 5. Use of a standardised, well defined and publically available drive cycle was felt preferable over the development of a bespoke drive cycle for reasons of reproducibility and transferability of the work. Recorded driving data was used to generate a gear shift schedule for use with the LA92 in accordance with relevant legislative procedures, and the drive cycle and gear shift schedule are carried forward to the development and evaluation work carried out in Chapters 4 – 6.

In addition to the selection of a reference drive cycle the driving data obtained here are also processed to obtain the probability data required for the development of the SDP control strategy in Chapter 5.

Chapter 4   Hybrid Vehicle Modelling

Computational modelling of the HEV is of key importance to the control design techniques described later, and this chapter describes in detail the experimental data collection and empirical component sub-models which form the whole HEV model.

A bespoke model was developed in the SIMULINK environment with a focus on characterising components at the highest possible level. In particular, the conventional powertrain components (ICE, transmission, differential, tyres) were mapped as a single unit on a chassis dynamometer to obtain a direct link between tractive force and fuel consumption. This modelling approach is a novel contribution and is presented in the following separately published work:

C. Vagg, C. J. Brace, R. Wijetunge, S. Akehurst and L. Ash, "Development of a new method to assess fuel saving using gear shift indicators," *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering,* 226 (12), pp. 1630-1639, 2012.

Optimal control of a hybrid powertrain is highly dependent on the drive cycle to which the vehicle is subjected. Having gathered a considerable amount of data relating to the expected typical usage of Ashwoods' hybrids, and designed a driver feedback device with the aim of minimising the energy intensity of this drive cycle, the focus of this work will now shift towards optimal control of the hybrid vehicles.

All of the optimal control algorithms presently proposed in literature rely heavily on Model-Based Design (MBD) techniques already used widely in industry for control system design. In this approach a computational model of the physical system is built in which the effects of any control signal can be simulated, enabling the best controller behaviour to be found. Clearly the optimality of the final control strategy therefore depends heavily on the fidelity of the system model – an optimal controller designed with the aid of a poor simulation is unlikely to be optimal when implemented on the physical system. The bespoke powertrain model developed in this work for the purposes of control strategy development is described in this chapter.

4.1 Hybrid Vehicle Description

At the most basic level Ashwoods' hybrid system consists of an EM and a battery pack which enable kinetic energy usually dissipated as heat during braking to instead be captured through regenerative braking and stored in the batteries. This energy may then be used to assist at other times, supplying tractive force and therefore reducing the load on the engine, reducing the vehicle's fuel consumption and emissions.

In the common classification of HEVs the configuration adopted here is categorized as a parallel torque-assist hybrid vehicle, and the architecture is illustrated in Figure 4-1. The EM adds tractive power through a belt and pulley, with the pulley being sandwiched between the propeller shaft and the final drive (differential). Since the propeller shaft incorporates some telescopic travel to accommodate suspension movement the pulley can be inserted without any modification to the standard components. In this arrangement the speed of the EM is directly proportional to the vehicle road speed as it is integrated downstream of the clutch and transmission; this

has the advantage that regenerative braking is possible regardless of whether the clutch is engaged, but carries the disadvantage that the EM must be able to operate over a wide speed range.

The hybrid system adds tractive force, and the mechanism of fuel saving is through the driver's assumed reduction in accelerator pedal position – there is no direct intervention in the accelerator pedal position reported to the ECU or the engine control signals. It is possible that the driver would not reduce the accelerator pedal position and would simply make use of the additional power, however in most cases drivers are not operating at the limit of the power already available to them, but rather are limited by traffic or road conditions. Furthermore any driver consistently making full use of the vehicle's installed power would likely fall foul of the DAS presented in the previous chapter, and it is the commercial intention that these two systems be sold together.

The battery pack is able to store around 2.16 MJ (0.6 kWh) of energy at a nominal voltage of 78 V and the system has a peak power of about 6 kW. The system is a 'mild hybrid' because the electrical powertrain is much less powerful than the standard ICE. Some of the functionality common in full hybrid vehicles, such as the
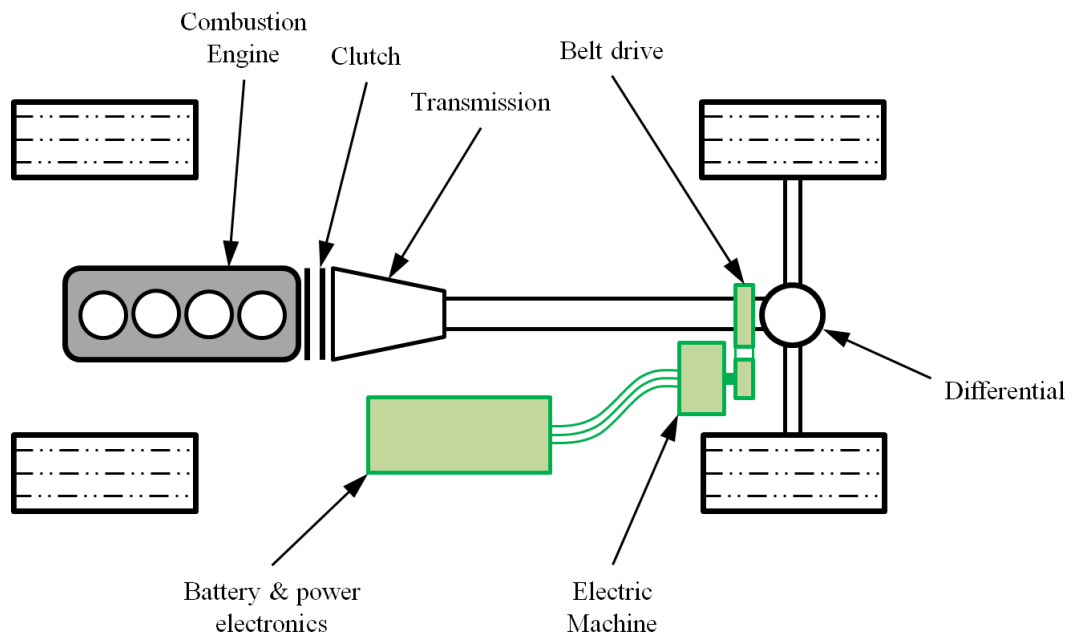

Figure 4-1: Schematic of the Ashwoods Automotive hybrid electric conversion system, showing the parallel torque-assist architecture. Stock components are in black, while retrofit components are highlighted in green.

ability to turn off the ICE and operate in 'electric only' mode, is not available.

The stock vehicle which was hybridised in this work was a Ford Transit 330S, details of which are given in Table 4-1. A model of this vehicle with the retrofit hybrid electric system fitted was developed in the MATLAB/SIMULINK environment, and the following sections describe the model's architecture, implementation, and validation.

Table 4-1: Characteristics of the baseline vehicle.

Symbol	Quantity	Value
Vehicle		2011 Ford Transit 330S
Engine	Model	2.2 L 100 PS diesel
	Peak power	74 kW @ 3500 rpm
	Peak torque	310 Nm @ 1300-2100 rpm
Transmission	Model	MT82 – 6 speed manual
	Gear 1	5.441
	Gear 2	2.839
	Gear 3	1.721
	Gear 4	1.223
	Gear 5	1.000
	Gear 6	0.794
	Final Drive	3.909

4.1.1 Model Architecture

When considering the modelling of a vehicle two main approaches to the model architecture exist, referred to as 'Forwards' and 'Backwards'. A forwards model represents the physical vehicle and the chain of causality observed in real life: a desired vehicle speed (reference signal) is defined, and the difference between this and the actual vehicle speed is processed by the driver. The driver's response is to manipulate the vehicle controls (accelerator, brake and clutch pedals, and the gear selection) so as to reduce the speed error, and the vehicle's response to the control inputs is calculated.

In contrast, the backwards approach to vehicle modelling assumes that the drive cycle is completely described in advance and is followed perfectly by the vehicle. Given the vehicle speed trace the tractive force required to have achieved this is calculated, and the fuel consumed in the process of delivering this is found. In this case the effect of the hybrid control strategy is actually to determine the proportion of the required tractive force which is delivered by the electric powertrain, while the remainder is assumed to have been provided by the conventional powertrain.

Forwards models are in principle the more thorough approach to modelling the system, which includes the driver. Nevertheless they do not always represent the more appropriate choice for a particular application. The inclusion of a driver model means that the ideal speed trace defined by the drive cycle is not accurately followed and the actual speed trace completed may be affected by changes to the vehicle, for example to the hybrid controller. Furthermore accurate modelling of driver behaviour is actually extremely difficult since human drivers often respond non-linearly to differences between intended and actual speed, as well using foresight to anticipate accelerations and decelerations. If undertaken to a high degree of fidelity forwards models may allow interactions between driver and powertrain to be modelled and enable some study into drivability; however if done only to a basic level the approach adds little value.

A further significant advantage of the forwards modelling approach is that the model can be used to simulate the non-ideal response of the vehicle in cases where it is not capable of completing the drive cycle perfectly. Specifically, for a low power vehicle situations may arise where its performance limits mean it is unable to accelerate at the rate defined by the speed schedule. In this case a forwards model allows the vehicle to accelerate at its maximum possible rate until it achieves the ideal speed, in line with the likely response of a human driver, and so the model can provide some insight. In contrast a backwards model in the same scenario may identify that the speed schedule is outside of the vehicle performance limits and therefore return an error, but it would be unable to provide any further insight as to what the non-ideal response would look like.

Present models used in state-of-the-art controller design are typically based on the backwards architecture. Since the study of drivability is outside of the scope of this work this is also the class of model implemented here. A top-level block diagram of the model is illustrated in Figure 4-2 which shows the various subsystems in their order of execution, and the inputs and outputs of each. A brief description of the function of each subsystem is given below, followed by detail on the individual components and their modelling in Section 4.1.2.
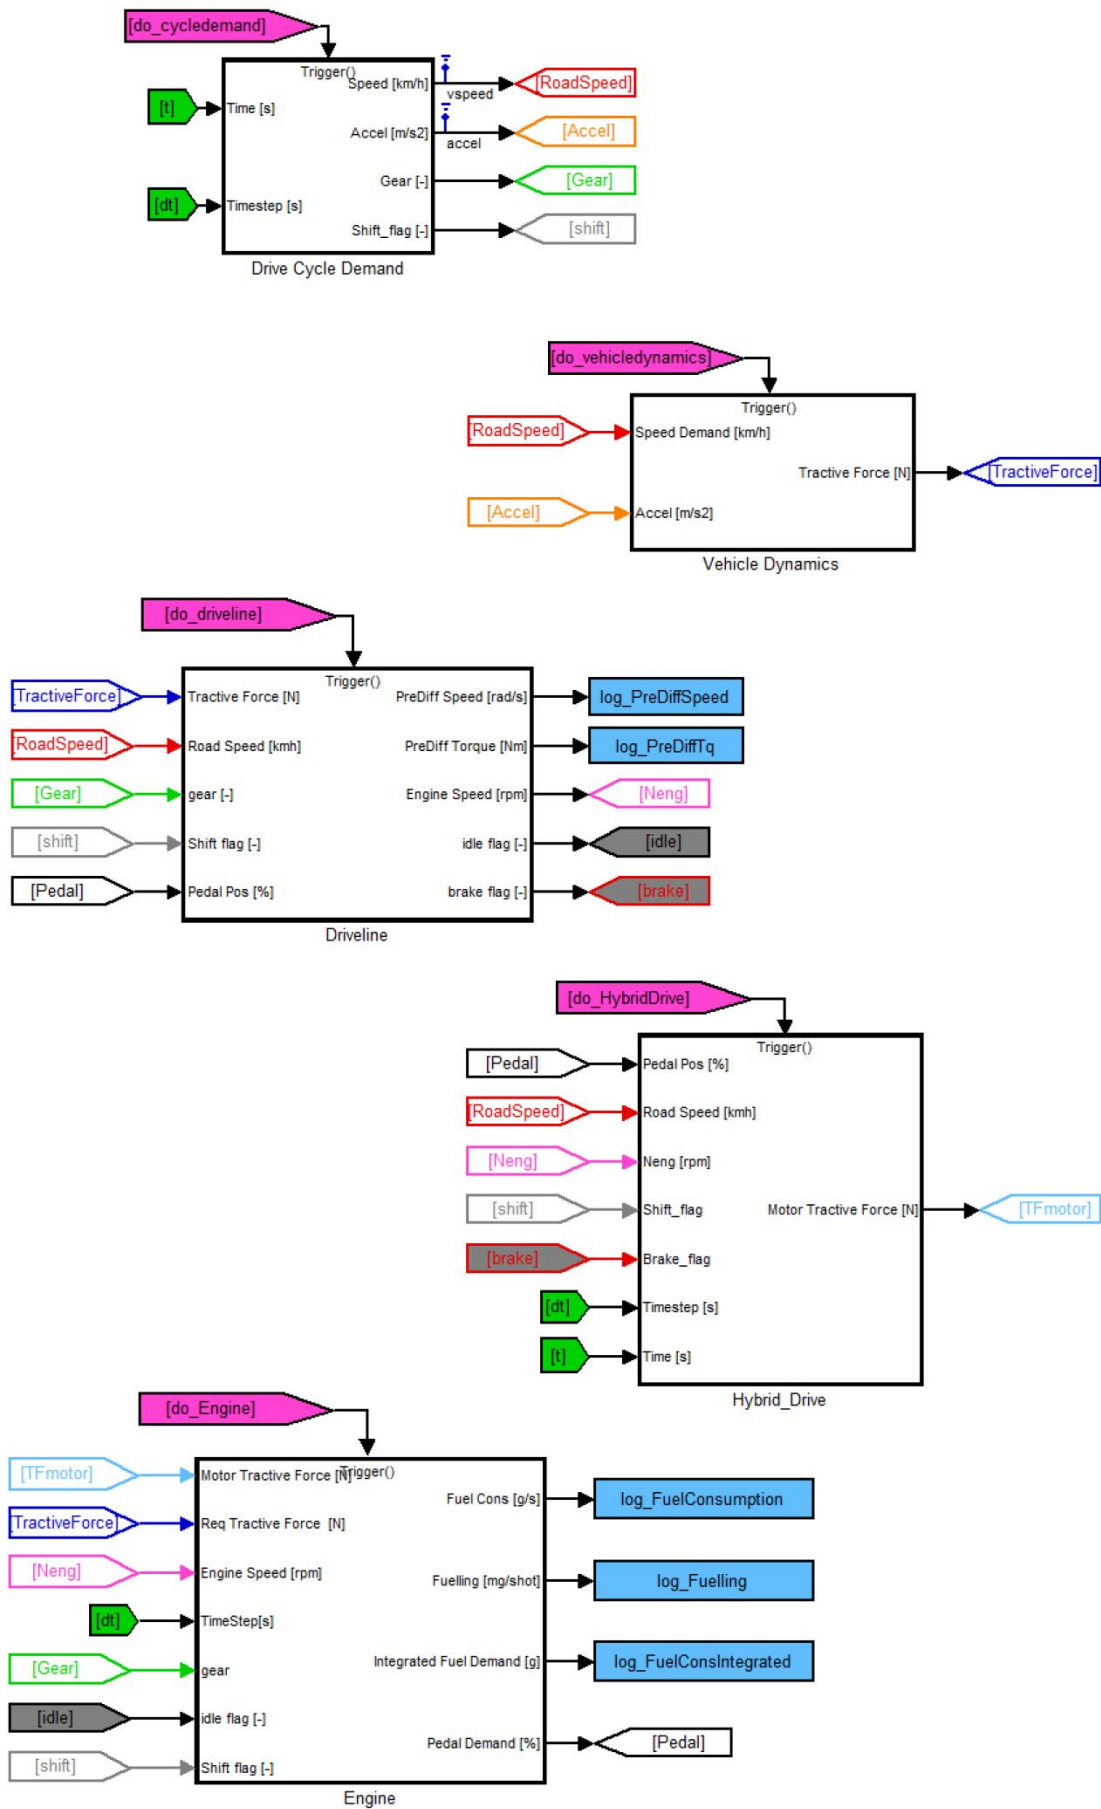
Figure 4-2: Top-level block diagram of the vehicle powertrain model in SIMULINK.

70

At the most fundamental level the model receives the drive cycle speed, acceleration, and gear as inputs, deduces the response of the hybrid controller in this state and the resulting tractive force from the electric powertrain, assumes that all of the remaining tractive force requirement is delivered by the diesel engine, and returns the fuel consumed in achieving this. To this end each block in the model is called by a function call generator in the following order, and performs the functions described accordingly:

(1) *Drive Cycle Demand*: The present simulation time is fed into a lookup table, which returns the present vehicle speed according to the prescribed drive cycle. The change in vehicle speed since the previous timestep is used to calculate the vehicle acceleration, which is passed through a low-pass filter. A second lookup table is accessed to retrieve the present gear selection from the gear schedule, and a change in the gear since the previous timestep causes the 'shift flag' to be triggered.

(2) *Vehicle Dynamics*: Given the vehicle speed at the present timestep the drag force due to aerodynamics and rolling resistance can be calculated. Present acceleration is multiplied by the vehicle mass to find the inertial load. These two components are then summated to find the total tractive force required at the wheels to achieve the present speed and acceleration.

(3) *Driveline*: Present vehicle speed and gear selection are combined to find the engine speed; the tractive force and gear selection are similarly combined to find the engine torque. The 'shift flag' signal allows an alternative engine speed to be specified during gear shift events. The accelerator pedal position (calculated during the previous execution of the 'Engine' subsystem) is used to detect when complete lift-off occurs, and it is assumed that the brake pedal is applied in this case.

(4) *Hybrid Drive*: This block contains three further subsystems: 'Hybrid Controller', 'Battery', and 'Motor'. The Hybrid Controller subsystem simulates the controller's response to the present vehicle state; this is a unitless number in the range ±255 which specifies the battery current demand as a proportion of the maximum available at the present vehicle speed. The demand is interpreted and scaled by the Battery subsystem (as per the BMS

on the physical vehicle), and the present DC battery current, voltage, and SOC are output. Finally, in the Motor subsystem the motor speed is calculated from the vehicle speed, and this is combined with the DC power to find the present motor torque and tractive force at the wheels.

(5) *Engine*: Tractive force supplied by the engine is found by subtracting that supplied by the motor from that required at the wheels. The tractive force provided by the engine, engine speed, and gear selection are used to infer instantaneous fuel consumption. The same three variables are used to find the present accelerator pedal position (used in the 'Driveline' subsystem). The 'idle flag' is used to swap the fuel consumption to a known engine fuelling rate during idling; the 'shift flag' may be used to select zero fuelling rate during gear shift events, though this was found to introduce inaccuracies during ideal drive cycle simulation and was only used when the input to the simulation was recorded data from a real driver following the ideal drive cycle.

4.1.2 Component Characterisation and Modelling

Detailed empirical models of each powertrain component were constructed. Empirical component modelling based on measured performance data was preferred over physical modelling, since for these purposes accurate representation of the macroscopic component behaviours is more important than an understanding of the underlying physical causes. For each component the data collection and modelling approaches are described below.

Modelling of Vehicle Dynamics

Forces acting on the vehicle are assumed to include only aerodynamic drag, rolling resistance and inertia. At this stage for the purposes of simulation and dynamometer testing forces arising due to road gradient are not considered, as the drive cycle data does not contain information on this.

Rather than the classical approach of modelling aerodynamic drag from first principles the steady state road forces were modelled as a second order polynomial of vehicle speed, as is common in dynamometers. The quadratic coefficients are commonly found either from a 'coast-down test' on a straight level road, or from

legislative tables of standard values, as were used here. Inertial forces are proportional to the vehicle acceleration, with the greatest by far being due to the vehicle mass. Other minor contributions to vehicle inertia are made by the rotational inertias of wheels and driveline components, and these were also represented in the simulation using approximate equivalent masses suggested by Miller [72]. Combining all these forces,

$$F_{TR} = F_2 \cdot v^2 + F_1 + \frac{dv}{dt}(m_v + m_w + m_d).$$

(4-1)

In this approach F_{TR} is the total tractive force acting at the road-tyre interface, v the vehicle speed, F_2 a constant resulting from the vehicle aerodynamic properties, F_1 a constant rolling resistance force, m_v the vehicle mass and m_w and m_d the equivalent inertial masses of the wheels and driveline components respectively. Values used for these constants are given in Table 4-2.

Table 4-2: Constants used in modelling the vehicle dynamics.

Symbol	Quantity	Value
m_v	Vehicle mass	1930 kg
m_w	Equivalent mass of wheels	40.9 kg
m_d	Equivalent mass of drivetrain	39.9 kg
F_1	Tractive force term	11.05 N
F_2	Tractive force term	0.9733 N·s^2/m^2

Modelling of the Diesel Engine

In order to obtain fuel consumption data for the engine a mapping exercise was undertaken on a chassis dynamometer. The objective of this exercise was to obtain an accurate map of fuel consumption as a function of engine speed, tractive force, and gear selection. Performing this mapping exercise on a chassis dynamometer rather than an engine dynamometer was advantageous because the result directly linked fuel consumption to tractive force, accounting for all losses that occur between the engine crankshaft and the road surface such as transmission and tyre losses. Engine fuel consumption should be identical for any engine speed and torque regardless of gear, however it was felt worthwhile repeating the exercise for each gear so that any differences in gearbox or tyre transmission efficiencies at different speeds were captured. Engine load and accelerator pedal position were also logged

throughout the exercise and the availability of this information was important for other facets of the control strategy development; it was therefore also worthwhile completing the exercise in each gear to obtain these data.

The test facility is temperature controlled and was maintained at 25°C throughout all tests. Regular checks of vehicle battery charge, engine oil level and tyre pressures were performed to ensure repeatable data collection and so that these variables could be similarly monitored for parity during the subsequent vehicle testing described in Section 6.2.2.

Fuel consumption was determined by analysing the components of the vehicle exhaust gasses and performing a carbon balance calculation. Exhaust gas composition was measured by two Horiba MEXA 7000 series emissions analysers, and knowing the carbon composition of the fuel the mass of fuel burned (*f*) was calculated by

$$ f = \frac{1}{d_{CWF}} \left(d_{CWF} \cdot m_{THC} + 0.428 \cdot m_{CO} + 0.273 \cdot m_{CO_2} \right) \qquad (4\text{-}2) $$

where d_{CWF} is the carbon weight factor of the diesel fuel (the mass ratio of carbon in the fuel, here 0.867), $m_{THC}, m_{CO}, m_{CO_2}$ are the measured masses of THC, CO and CO_2 emissions, and 0.428 and 0.273 represent the ratio of the atomic weight of carbon to the molecular weight of carbon monoxide and carbon dioxide, respectively [73].

Two forms of emissions readings are available from the dynamometer systems: bag analysis values, and modal emissions. Bag analysis is the industry standard method for measuring total cumulative emissions over a test, and is achieved by taking a proportional sample of exhaust gasses throughout the test, storing this in one or more bags, and analysing the contents of the bag after completion of the test to infer the total mass of emissions during the test, in accordance with UN/ECE Regulation No 83 [4] and ISO 16183:2002 [74]. This approach is extremely accurate but does not allow any insight into the instantaneous emissions production; only the total emissions over the test are reported. In modal emissions measurement a continuous

sample of tailpipe gas is passed through emissions analysers real-time and the results recorded at 10 Hz.

During engine mapping the dynamometer was used in closed-loop speed control mode and manually adjusted to achieve a range of engine speeds. At each speed the accelerator pedal was progressively applied in steps to achieve the full range of tractive force, while the dynamometer supplied the necessary reactive torque to maintain the set speed. Each speed-force point was held for 5 seconds to allow gas concentrations to stabilise and during this period modal emissions measurements were used to sample the rate of fuel consumption. This procedure of progressive application of the accelerator pedal in steps and then holding each sample point represents a significant amount of learning from a pilot study published by the author elsewhere [68]. During the pilot study an engine was mapped in a similar way but applying a continuous ramp to the accelerator pedal. As a result of gas mixing in the connecting pipes the sampled emissions concentrations were found not to respond instantly to a step change and so during the application of a continuous accelerator pedal ramp a time lag is introduced, which caused hysteresis in the fuel consumption loop. The stepped ramp procedure used during this mapping was found to yield better results.

Each pedal ramp was applied over 2 minutes and then released over 2 minutes, however the pedal step size was increased during the second half of the ramp because the engine is far less responsive to changes in accelerator pedal position near
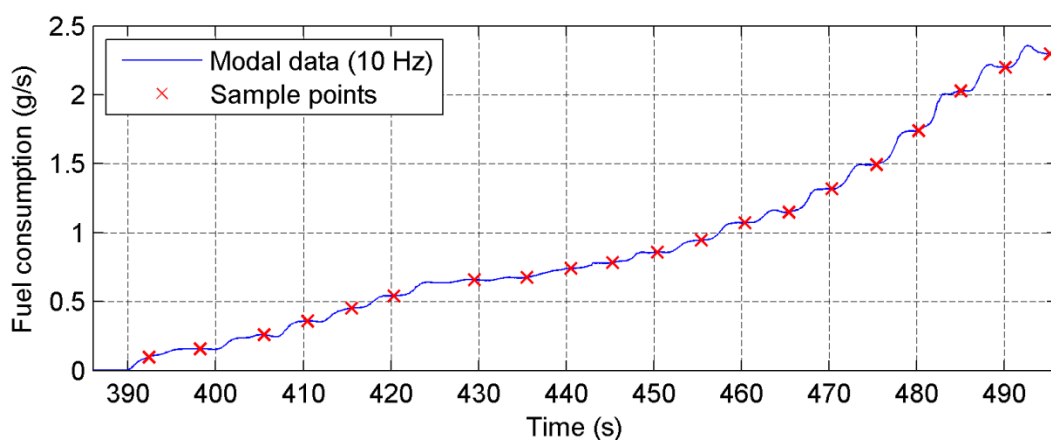

Figure 4-3: Sample points of engine fuel consumption during a stepped ramp application of the accelerator pedal on the chassis dynamometer. These data were recorded at 2500 rpm in gear 4 (69.5 km/h).

full load, as was also observed during the pilot study. A MATLAB script was written to identify points during the ramps at which CO_2 emissions were stable for periods longer than 1.5 seconds and the tractive wheel force was also stable, and these points were used as steady state samples of the engine operation as shown in Figure 4-3.

As well as showing how sample points were selected Figure 4-3 also hints at the non-linear relationship observed between accelerator pedal position and engine response. Throughout the time window shown the steps in accelerator pedal position between each sample were approximately constant, but it can be seen that the fuel consumption (which tracks engine torque) increases linearly at first, then is almost unresponsive at around 435 seconds, before becoming increasingly sensitive.

Collating the steady state samples of engine operation from ramps at different engine speeds allowed a complete picture to be built of the engine fuel consumption as a function of engine speed and tractive force in each gear, an example of which can be seen in Figure 4-4. It is perhaps worth noting that fuel is burned even when some negative tractive force is being generated, as the engine will consume fuel in overcoming its own internal losses; it is for this reason important to find the tractive force at zero load (i.e. maximum engine braking) as this defines the accelerator pedal "lift-off" point described in Section 4.1.1.
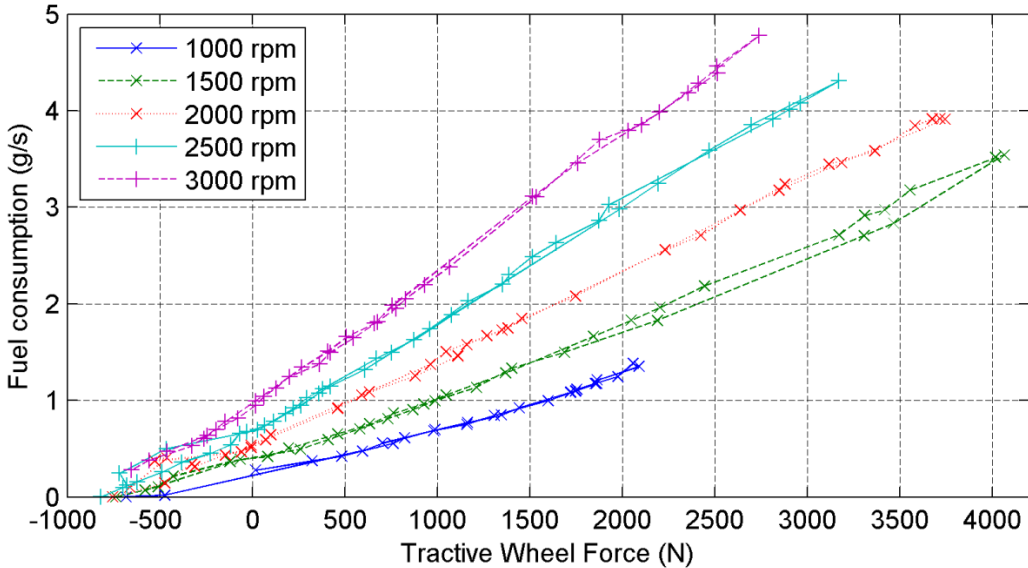

Figure 4-4: Fuel consumption map obtained in gear 4 using the accelerator pedal stepped ramp technique. Data at intermediary speeds were also obtained but are omitted for clarity.

The data sets for each gear were parameterised as third-order polynomial surfaces using the MATLAB Model-Based Calibration toolbox, and the polynomial coefficients used to calculate fuel consumption in the vehicle powertrain model. This type of engine model, based on steady-state measurements, is often called *quasi-static*. As previously described the models implemented directly linked tractive force and engine speed to fuel consumption, and an example of the gear 5 polynomial surface is shown in Figure 4-5. Data for the quality of fit of the engine model in each gear are given in Table 4-3.

Table 4-3: Quality of fit statistics for fuel consumption models in each gear.

Symbol	R^2	PRESS R^2
Gear 1	0.991	0.987
Gear 2	0.996	0.995
Gear 3	0.998	0.998
Gear 4	0.998	0.998
Gear 5	0.999	0.998
Gear 6	0.993	0.993

If desired, the data collected may also be manipulated using the known gear ratios to present the results in more conventional forms, such as the Brake Specific Fuel Consumption (BSFC) map shown in Figure 4-6, however it must be noted that in this
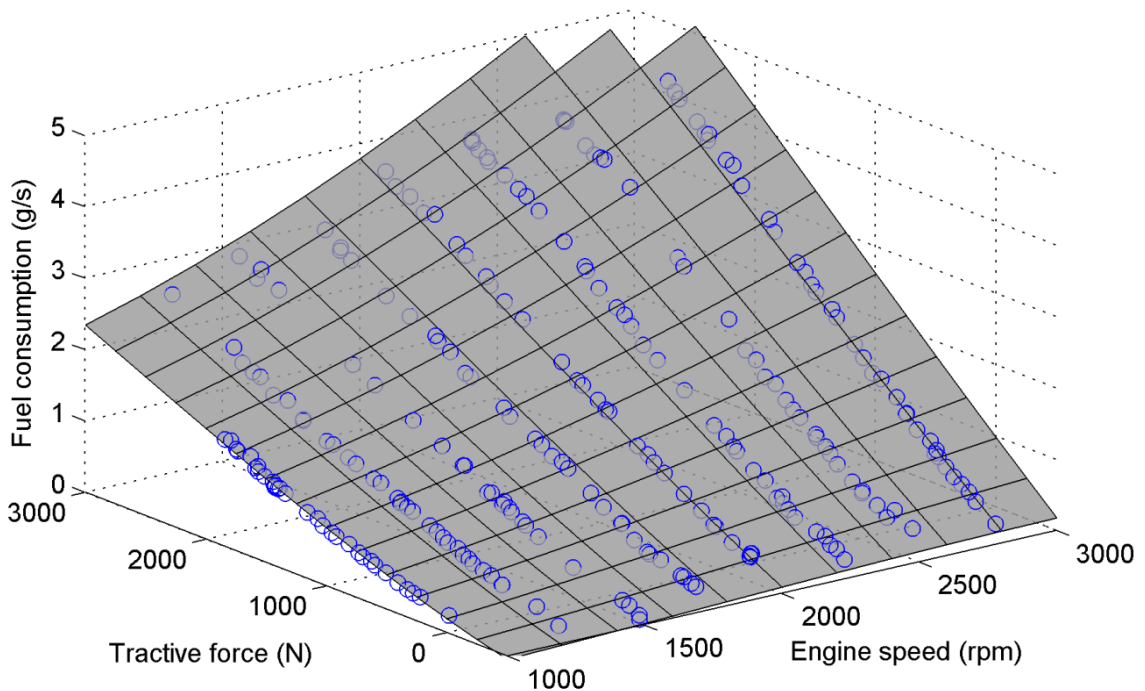

Figure 4-5: Polynomial model of fuel consumption in gear 5 as a function of engine speed and tractive force.
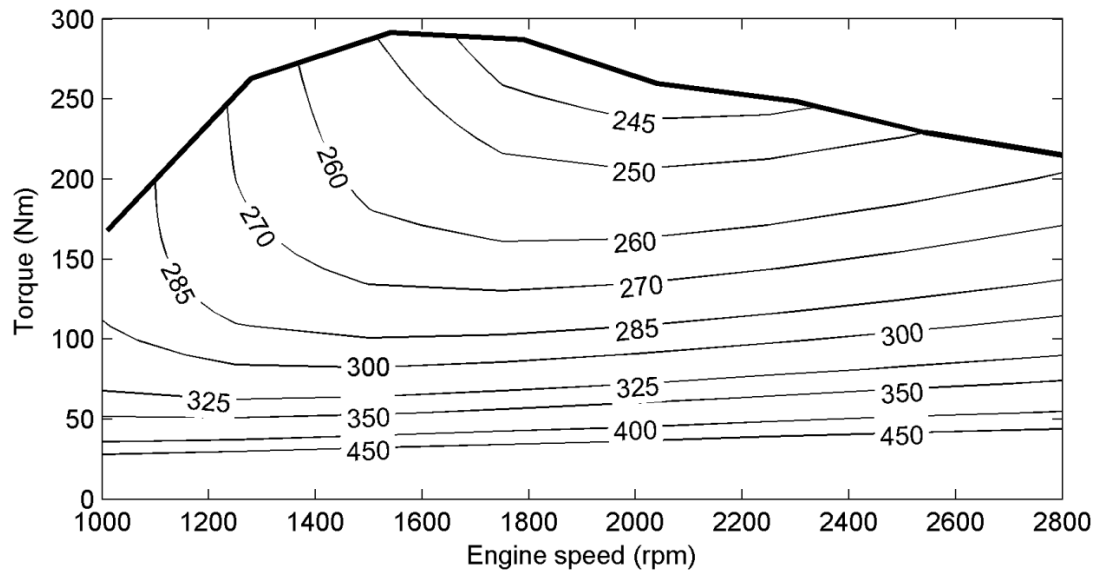
Figure 4-6: BSFC (g/kW·h) of the engine, constructed from the data collected in gear 5. Note that the engine torque is calculated from measurements taken at the tyre, not crankshaft, and so include transmission losses.

case the results represent the brake work measured at the tyre (after transmission losses) and not at the crankshaft as would be usual.

Modelling of Motor and Power Electronics

The electric motor used in the retrofit hybrid electric system is a bespoke unit manufactured by Ashwoods Automotive; it is a three-phase axial flux permanent magnet machine rated to 6000 rpm and producing a maximum torque of 65 Nm. Key specifications of the motor can be found in Table 4-4, and a photograph of the unit in Figure 4-7.

Table 4-4: Specifications of the Ashwoods Automotive motor.

Quantity	Value
Maximum torque	65 Nm @ 2200 rpm
Maximum power	14.8 kW @ 3250 rpm
Continuous torque	11.5 Nm @ 2500 rpm
Continuous power	3.0 kW @ 2500 rpm
Operating speed	0-6000 rpm
Peak efficiency	>95%
Cooling	Air
Weight	15.0 kg
Dimensions	ø237 mm × 130 mm
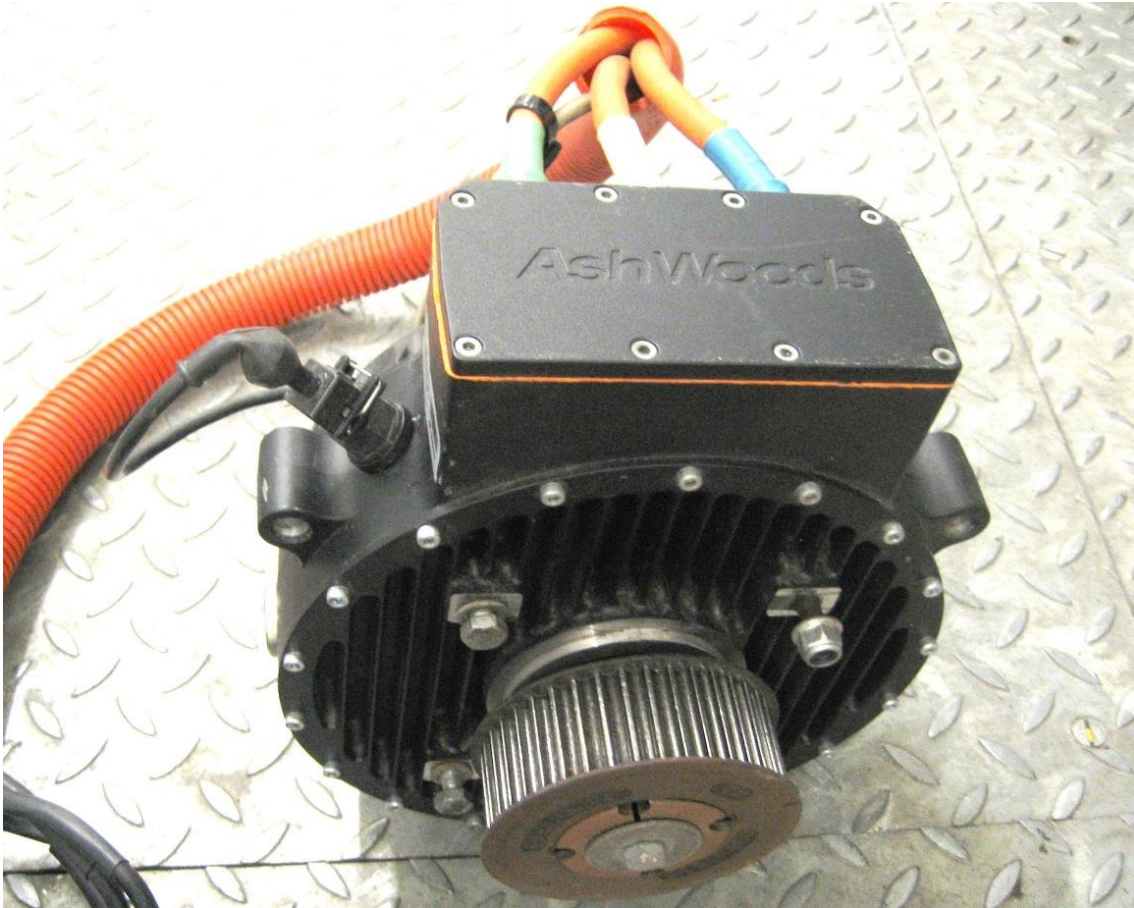
Figure 4-7: Ashwoods Automotive axial flux permanent magnet motor, as fitted to the hybrid vehicle.

Using a similar method to that described for the engine mapping a model of the power electronics and electric motor was developed from data collected running both on a motor test facility. The tests and data collection were performed by other researchers and so no credit is claimed for this. The data were processed and then implemented in the powertrain model as a 2D look-up table, using linear interpolation between measurements. Motor speed and DC power consumption define the operating point, and the torque produced is calculated. This representation of the data, as shown in Figure 4-8, is not the typical way in which motor performance is reported; it is perhaps more normal to have efficiency as a function of a speed and torque, however because the hybrid controller specifies a DC current demand it makes more sense in this case to find the torque output resulting from this.

As with the engine mapping exercise the characterisation of the motor and power electronics as a discrete unit is a simplification, but because the measurements
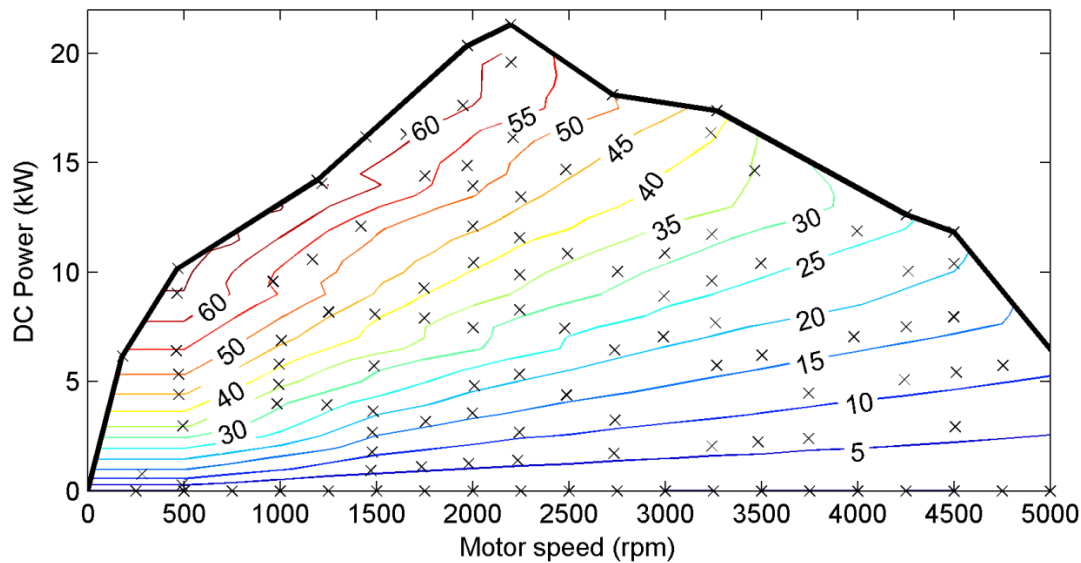
Figure 4-8: Motor map defining the torque output (Nm) for any combination of motor speed and DC power within the operating envelope. Points marked (x) are data points obtained experimentally.

account for any losses or component interactions it is possible to achieve greater accuracy in the model without added complication.

Although the peak power of the motor is 14.8 kW in fact the peak system power is limited to 5.8 kW by the battery current, as will be discussed in the following section. Furthermore the temperature of the motor rapidly increases during operation at high power, and so for continuous operation the motor power is limited to 3.0 kW. The power electronic inverter (motor controller) used during motor characterisation was a SEVCON Gen4 600 amp unit capable of supplying the full power of the motor; however since the available system power on-vehicle is less than this a smaller variant in the same range – the SEVCON Gen 4 350 amp – is used. The use of a different inverter during motor characterisation to on-vehicle may mean that the model is not entirely representative of the vehicle system, however it is assumed that since the two variants are from the same manufacturer and part of the same range any differences are likely to be relatively small.

On-vehicle the electric motor was geared such that at the vehicle's maximum speed of 100 km/h the motor was also at its maximum speed of 6000 rpm. At this speed it is not possible to demand any torque from the motor because the back emf approaches the DC battery voltage.

Modelling of Batteries

Lithium iron phosphate (LiFePO$_4$) battery cells are used as the electrical energy store for the hybrid system. Each cell is a Headway HW38120HP, and the battery pack consists of 24 cells arranged in series; key parameters describing the cells and the pack are shown in Table 4-5, and a photograph of the assembled pack is shown in Figure 4-9.

Table 4-5: Specification of the battery cells and pack.

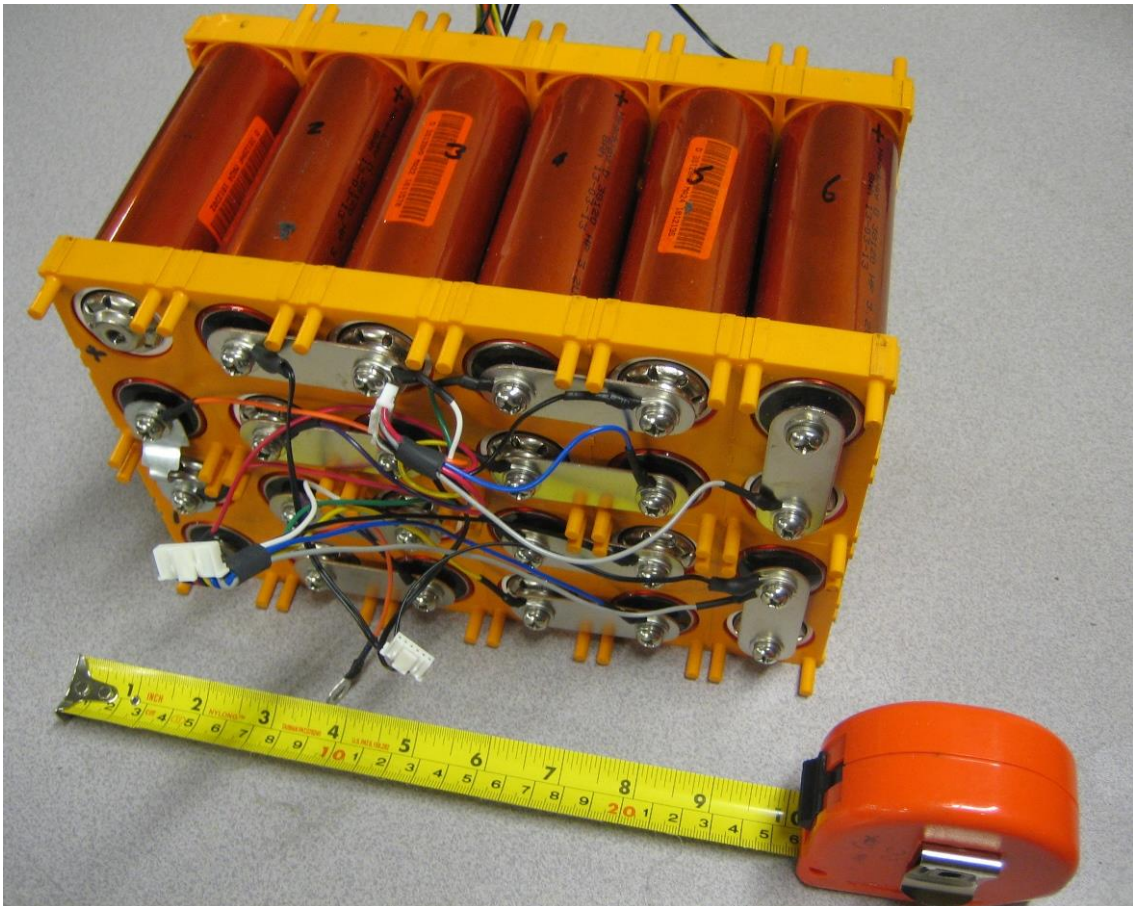| Quantity	Cell Value	Pack Value
Manufacturer	Headway	
Model	HW38120HP	
Nominal voltage	3.2-3.25 V	76.8-78 V
Internal resistance	8 mΩ	192 mΩ
Nominal capacity	8.0 Ah	
Max charge current	80 A (10 C)	
Max continuous discharge	80 A (10 C)	
Max pulse discharge current	200 A (25 C)	

Figure 4-9: Battery pack used in the retrofit hybrid electric system, comprising 24 lithium iron phosphate cells in series. The wires shown allow the BMS boards to monitor individual cell voltages.
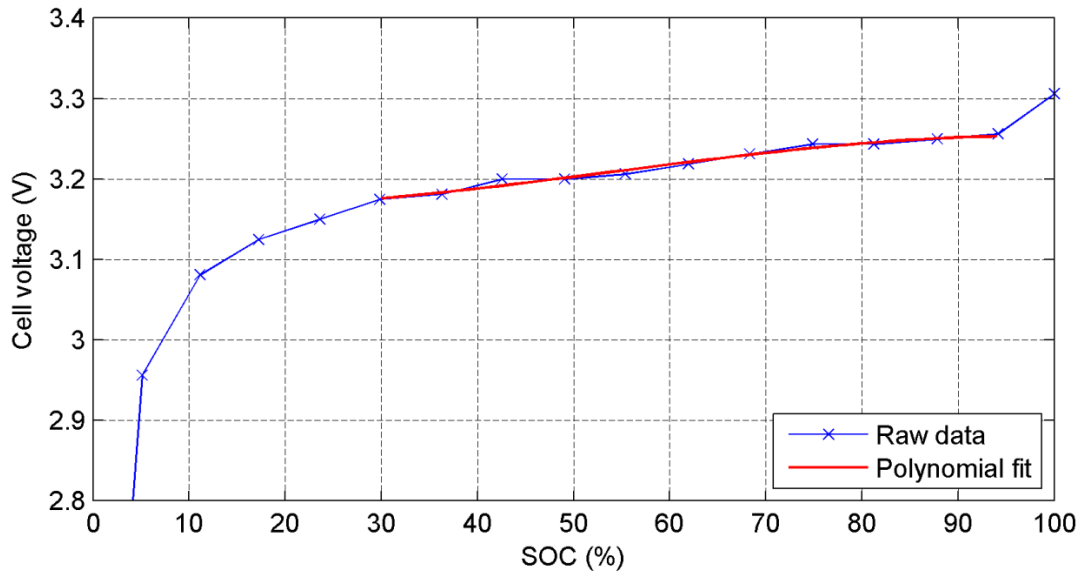
Figure 4-10: By processing and rearticulating the 1 C time-voltage discharge plot provided in the cell data sheet a SOC-voltage plot was obtained. A polynomial function was used to model open circuit cell voltage in the range 30-90% SOC.

In the absence of battery testing facilities the battery pack was modelled using data supplied on the manufacturer data sheet. Since the hybrid control strategy specifies a battery current the primary purpose of the model is to estimate the change in SOC as a result of the current; for this purpose the prediction of battery terminal voltage is not of major importance and so a relatively simplistic battery model was deemed adequate, with an equivalent circuit comprising only of an internal resistance, and not accounting for any time-dependent capacitive effects. Data sheets for the battery cells provide time-voltage discharge curves at various discharge rates. For estimation of the open circuit cell voltage the curve for a 1 C discharge rate was digitized using MATLAB and the energy dispensed during a complete discharge evaluated, thus allowing the data to be rearticulated as a SOC-voltage plot as shown in Figure 4-10. It was found that the entire curve could not be represented well by a polynomial function, however the range 30-90% SOC could be well represented. Since 30-90% SOC covers the desirable operational state of the cell a polynomial was used to approximate the open circuit cell voltage within this range, $V_{oc}(SOC)$.

Of greater importance to SOC estimation is that the energy losses incurred during charge/discharge are accounted for, which was achieved by calculating an *equivalent current*, I_{eff}.

82

Equivalent current is the hypothetical current which we may imagine has been drawn from or absorbed by each cell to account for its charge/discharge inefficiencies. During discharge the equivalent current will be higher than the DC current, in recognition of the fact that some energy is consumed by internal cell losses; similarly during charge the equivalent current will be lower than the DC current, to represent the fact that not all of the DC current is converted into stored chemical energy.

During charge events the cell efficiencies were assumed to be dominated by ohmic (I^2R) losses, and therefore quantified by using the cell internal resistance. For discharge events the time-voltage graphs from the cell data sheets were employed again, and the curves for a variety of discharge rates ranging from 1 C to 20 C were digitized to calculate the energy delivered during a complete discharge at each C-rate. At higher C-rates the total energy delivered is less than at low C-rates, allowing an effective current to be calculated at each C-rate. To characterise the reduction in available energy capacity of a battery when operating at higher C-rates Peukert's Law is often applied. When considered in terms of effective current Peukert's Law may be expressed as stating that the effective current is proportional to the discharge current raised to the power of the *Peukert constant*, K_P, which is an exponent that may be experimentally determined for any battery. Figure 4-11 shows how the Peukert constant was determined as 1.0526 using a trendline through the calculated
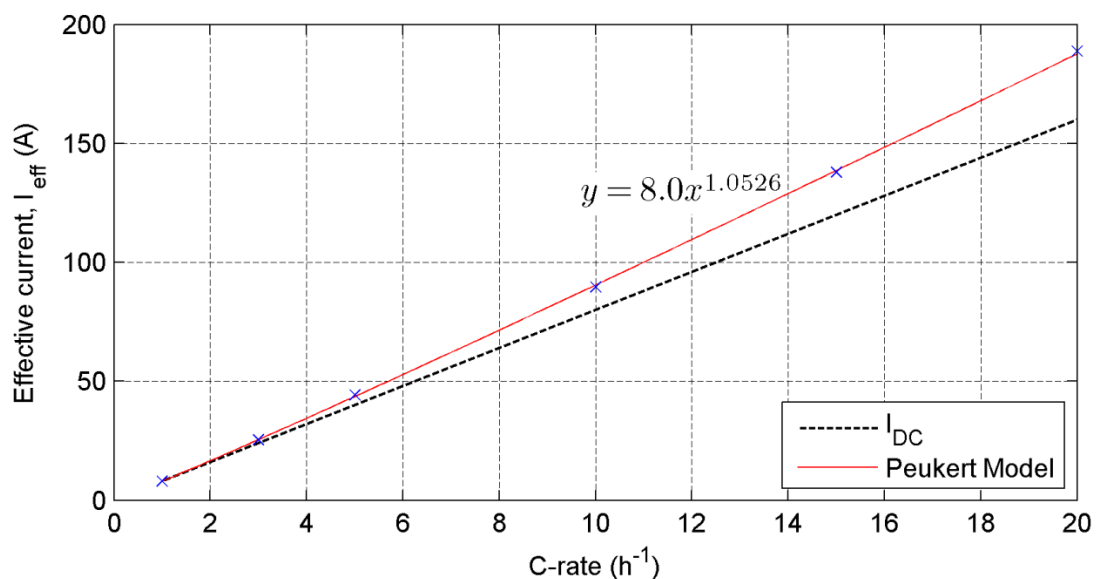

$$y = 8.0x^{1.0526}$$

Figure 4-11: Effective currents for six discharge rates were calculated based on effective capacity determined from manufacturer discharge curves. At high discharge rates effective current is substantially higher than the DC current due to internal cell losses. The Peukert coefficient best matching the trend is 1.0526.

effective currents, and also conveys graphically the difference between the measured DC current and the effective current at higher discharge rates. The equations describing the battery pack model are therefore:

$$V_{bat} = 24 \times (V_{oc}(SOC) - I_{DC} \cdot R_{int})$$ (4-3)

$$SOC_{k+1} = \frac{(SOC_k \times Q \times 3600) - (I_{eff} \cdot \Delta t)}{Q \times 3600}$$ (4-4)

where R_{int} is the internal resistance of each battery cell, I_{DC} is the DC bus current flowing out of the battery, V_{bat} is the terminal voltage of the battery pack, subscript k is a discrete time index, and Q is the nominal battery capacity in amp-hours. The effective current is defined by

$$I_{eff} = \begin{cases} I_{DC} + \dfrac{I_{DC}^{2}(24 \cdot R_{int})}{V_{bat}}, & I_{DC} < 0 \\ \\ Q \cdot C^{K_P}, & I_{DC} \geq 0 \end{cases}$$ (4-5)

where

$$C = \frac{I_{DC}}{Q}$$ (4-6)

is the battery C-rate: a fraction representing the instantaneous current draw as a proportion of the rated battery capacity Q. A continuous current draw of 1 C therefore flattens the battery completely in 1 hour.

Finally, it should be noted that although the specifications from the battery data sheet, reproduced in Table 4-5, suggest that charge/discharge rates of 80 A/200 A respectively are possible this is a simplification. In practice there are many other factors that affect maximum or minimum charge/discharge rates, with maximum/minimum allowable terminal voltage being one of the most significant. In order to the ensure cell voltages were reliably operated within their specified terminal voltage range the battery management system limited current to 63 A (7.9 C) in charge and 54 A (6.8 C) in discharge. These were therefore regarded as the absolute maximum limits within which the hybrid control strategy could operate.

4.1.3 Engine Model Validation

Since the engine model is quasi-static, based solely on steady-state measurements, there was concern that it might be a less accurate predictor of highly transient behaviour. Since the NEDC is a relatively gentle drive cycle with much time spent in steady state and very few aggressive transients any difficulties predicting transient fuel consumption may not be detected through validation on a NEDC. For this reason the engine model's behaviour was also compared against recorded data on a LA92 cycle, which is highly transient; this cycle was selected due to its strong correlation with observed real-world use, as already described.
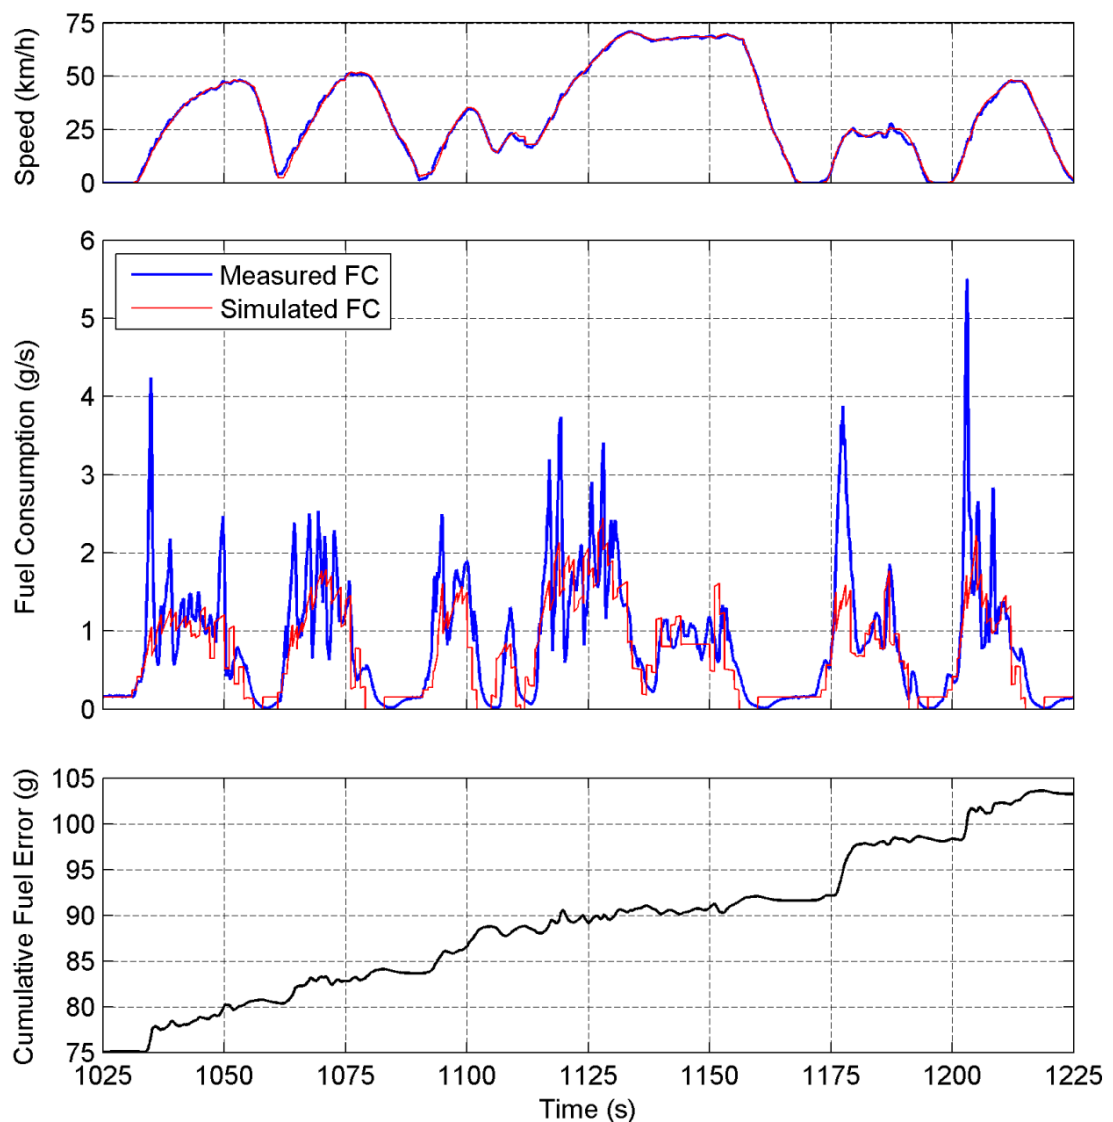

Figure 4-12: Correlation between measured and simulated instantaneous fuel consumption over an excerpt of the LA92. This correlation is considered very good considering the engine model is quasi-static and the drive cycle is extremely transient.

Comparison of the predicted and recorded instantaneous fuel consumption over a LA92 is shown in Figure 4-12. In order to gain insight into the ability of the model to predict fuel consumption even when the speed trace is not precisely followed the modelled fuel consumption is based entirely on the ideal speed trace, not the recorded speed trace. It can be seen that the model generally predicts the instantaneous fuel consumption extremely well despite the highly transient drive cycle. Since the ideal speed trace is defined at 1 Hz there are abrupt changes in acceleration every 1 second, which gives rise to the step-like appearance of the modelled fuel consumption.

The greatest discrepancies between the simulation and measured data are during pull-away events, as shown by rapid increases in the cumulative fuel error. In these cases the simulation under-predicts fuel consumption because clutch slip was not modelled in detail. Overall the simulation under-predicts the fuel consumption by about 4.9% over a LA92; this was deemed extremely encouraging as a worst-case scenario, using a quasi-static engine map to predict fuel consumption over a highly transient drive cycle, and where the measured drive cycle completed may differ slightly from the ideal speed trace.

4.2 Chapter Conclusions

This chapter has presented the development of a bespoke vehicle model in SIMULINK representing the vehicle hybridised in subsequent chapters. Powertrain components were characterised at the highest possible level, based on experimental data where possible and physics-based approaches otherwise.

The component models developed in this chapter are used in Chapter 5 to calculate the consumptions of fuel and electrical energy in any state, which are fundamental to the model-based control design approaches pursued. In Chapter 6 the complete vehicle model will be employed to simulate the performance of the HEV when running the control strategies developed in Chapter 5.

Chapter 5   Hybrid Control Development

In this chapter the development of HEV controllers is described in detail. The primary focus of effort is on the design of SDP controllers using the real-world data already recorded as the basis for a stochastic drive cycle model (Markov chain). ECMS controllers are also developed, primarily as a benchmark. For both control approaches a set of controllers are developed which exert stress on the electric powertrain to different extents.

The development process used in the design of the SDP controllers is extremely thorough and believed to be the most concerted effort to date to produce a controller, based on recorded driving data, which is suitable for real-world implementation. Elements of the process, notably the use of *probabilistic charge sustenance plots*, are believed to be entirely novel. Further novelty lies in the design of the cost function proposed which allows fuel saving to be traded-off against electric powertrain stress. Parts of this chapter have been published in the following separate works:

C. Vagg, C. J. Brace, S. Akehurst and L. Ash, "Model-based Optimal Control of a Hybrid Electric Vehicle Using Stochastic Dynamic Programming." *In: 6th Conference on Simulation and Testing for Automotive Electronics,* 2014-05-15 - 2014-05-16, Berlin, Germany.

C. Vagg, C. J. Brace, S. Akehurst and L. Ash, "Minimizing Battery Stress during Hybrid Electric Vehicle Control Design." *In: 9th IEEE Vehicle Power and Propulsion Conference,* 2013-10-15 - 2013-10-18, Beijing, China.

5.1 Defining Cost

As outlined in the aims one of the primary goals of this work is to integrate powertrain stress considerations, notably those of the battery, into the cost function. Moura [47] proposed two functions which could be incorporated into a cost function: the first aimed at accounting for SEI film growth at the anode of the battery, however this required a relatively advanced electrochemical model of the battery cells. The second function proposed by Moura simply accounts for battery current throughput, i.e. minimising the battery use, and similar cost functions were proposed by Ebbesen et al [48] and Serrao et al [49]. These functions were designed primarily to enhance the longevity of battery cells in the long term, rather than shorter term effects which this thesis aims to address, such as cell voltage and temperature imbalances within the battery pack, and motor temperature.

Long-term and short-term system healths certainly share similarities and the functions previously discussed would most likely have a positive effect on short-term health, however for the specific goals of this thesis an alternative is proposed. It is noted that the majority of electrical losses in battery cells and the motor are due to ohmic heating, which is proportional to the square of current (I^2R losses). Avoiding high power peaks in the use of the system and introducing a partiality towards more constant low power use would minimise these ohmic losses, with potentially beneficial effects on the system performance in the real world. It is not entirely straightforward to do this because the DC battery current and AC motor current are in fact decoupled from one another by the power electronic inverter, however the DC current is also a reliable indicator of system power, as the battery voltage is relatively steady. For these reasons the cost function

$$c = f + \alpha \cdot I_{DC}{}^2 \tag{5-1}$$

is proposed, where c is the instantaneous cost, f is the instantaneous fuel flow (g/s), α is an equivalence factor and I_{DC} is the DC battery current. Since the acceptable battery current is very much dependent on the size of a battery it is more helpful to work with the C-rate. In order to generalise the cost function and make the results more transferable it is therefore rephrased as

$$c = f + \alpha \cdot C^2 \qquad\qquad (5\text{-}2)$$

where C is the instantaneous C-rate (h^{-1}) at which the battery is being operated, and α therefore has units of $g \cdot h^2$. This cost function provides a mechanism to deter high power system operation, in turn reducing heat generation in the battery cells and EM.

Apart from defining a metric of optimality an important secondary purpose of the cost function was to implement any constraints which the control strategies must observe. This was done by assigning certain operating conditions either with a null value, for example 'NaN' (Not a Number), or with an extremely high cost, for example 'bitmax' (the largest double precision number in MATLAB), such that these operating points were never selected. An example of such a constraint is that when the vehicle is at 100 km/h the motor will be at 6000 rpm – its maximum speed – and so it is not permissible to demand any power from it.

5.2 SDP Controller Development

The SDP controller concept seeks to find the stationary policy which minimises the future cost incurred over an infinite time horizon, when starting from any vehicle state. For this aim a statistical model of the future is required, and the control policy is iteratively refined to minimise the expected future cost. This policy will not necessarily be optimal for any one drive cycle in particular, but rather is the single time-invariant control policy that would return the lowest cost when run for an infinite time period on a drive cycle with statistical distributions matching the design set. Over a finite drive cycle, say an hour in duration, it may be possible that an alternative policy would have better minimised the cost, or that there would be a more globally optimal non-stationary policy (i.e. changing over time) that could adapt to suit specific conditions or different drivers. However, assuming that over a long period of time the drive cycle which the vehicle is exposed to approaches that which was used in the design of the SDP controller, this policy is the single time-invariant policy which would yield the lowest cost.

The problem is formally defined and the algorithm formulation explained from a high level conceptual perspective in the next paragraphs, followed by a more detailed explanation as to how this was implemented.

We require a control decision to be made based on a discrete set of states,

$$u = \pi(x_k) \tag{5-3}$$

where π is a stationary control policy which may be interrogated for any state x_k to return the instantaneous control decision in that state, u. As described in Chapter 4 the drive cycle is defined entirely by three variables: vehicle speed (v), vehicle acceleration (a) and gear selection (g). To describe the state of the hybrid system the present SOC of the battery is also required and so, in our case, the full state vector on which control decisions are made includes four state variables.

$$x_k = \{v_k, a_k, g_k, soc_k\} \tag{5-4}$$

It is worth noting that there is no mathematical reason why the state vectors should have constant grid spacing, and so in fact the grid used for vehicle speed was sparser at higher speeds. The grid spacing for each state vector is given in MATLAB notation in Table 5-1.

Table 5-1: State vectors used in control development

Symbol	Quantity (Unit)	Parameter Value
v	speed state space (km/h)	[0:2:80 85:5:120]
a	acceleration state space (km/h·s)	[-14:14]
g	gear state space	[1:6]
soc	SOC state space (%)	[55:85]
u	control state space	[-255:17:255]

Any vehicle state combined with a control decision will result in some instantaneous cost, $c\big(x_k, \pi(x_k)\big)$, and by summing the instantaneous costs over an infinite time horizon we may arrive at an estimate of the future cost expected when the policy π is used starting from the present state x_0:

$$J_\pi(x_0) = \lim_{N \to \infty} E\left\{ \sum_{k=0}^{k=N-1} \lambda^k c\big(x_k, \pi(x_k)\big) \right\}. \tag{5-5}$$

Since the actual cost over an infinite horizon would clearly be infinite, the discount factor, $0 < \lambda < 1$, is introduced to ensure that the sum converges. The effect of λ is not to define a discrete window or time horizon for the sum, but to exponentially reduce the relative weighting of samples that extend into the future, such that in the far distant future the relative weighting approaches zero. Having said that, it is also true to say that λ defines how quickly the importance of the future is degraded. For example in the extreme case when $\lambda=0$ only the instantaneous cost is considered and the future is disregarded entirely; meanwhile $\lambda=1$ would be a true infinite sum, and therefore not converge at all.

Since the future states are not known these are instead modelled as a probability distribution and so Equation (5-5) is broken into two terms: the instantaneous cost and the future expected cost from the next state onward. Again, the impact of a control decision is twofold: firstly the decision will affect the instantaneous cost, for example we would expect the decision to assist the engine result in reduced instantaneous fuel consumption; secondly the decision will affect the system state in the next timestep, for example an assist will deplete the battery, and the SOC will consequently be lower at the next timestep. This effect on the future state is critical, since any depletion of charge must necessarily be made up in the future, and therefore the future cost expected when at a low SOC will be higher than the future cost from a high SOC. The reduction in instantaneous cost must therefore be considered with respect to the increase in the future cost which is expected as a result of depleting the battery SOC. This is expressed in Equation (5-6) where, on the right hand side of the equation, the first term represents the immediate cost and the second represents the future expected cost of every state, multiplied by the probability of being in each state at the next timestep as a result of the present control decision.

$$J_\pi(x_k) = c(x_k, u) + \lambda \sum_{x_{k+1} \in X} \mathbb{P}(X_{k+1} = x_{k+1} | x_k, u) \cdot J_\pi(x_{k+1}) \tag{5-6}$$

where

$$\mathbb{P}(X_{k+1} = x_{k+1} \mid x_k, u)$$

defines the probability of progressing into every vehicle state at the next time step, given knowledge of the present vehicle state and the present control policy decision resulting from this vehicle state. This matrix therefore has 8 dimensions in this case.

The objective is to determine the optimal control policy, π^*, which minimises the expected cost when starting from any state x_0. Put formally,

$$\pi^*(x_0) = \arg\min_u \left\{ c(x_0, u) + \lambda \sum_{x_1 \in X} \mathbb{P}(X_1 = x_1 | x_0, u) \cdot J_{\pi^*}(x_1) \right\}. \qquad (5\text{-}7)$$

Equations (5-6) and (5-7) introduce the idea that the vehicle state at the next timestep may be defined as a probability distribution, based only on the vehicle state in the present timestep. This mathematical description is known as a Markov chain, and posits that the vehicle states may be defined in a discrete state space where the transition probabilities between states at each timestep are time-invariant. An important property of the Markov chain is that it is memoryless – the transition probabilities depend only on the present state, and not on the history of the system. As an illustrative example Figure 5-1 portrays an extract of a basic Markov chain where the state variable is simply vehicle speed.
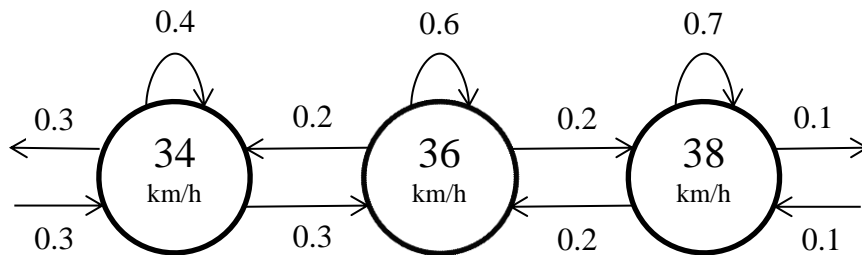
Figure 5-1: A portion of a Markov chain showing three states and the transition probabilities between them. In this simplistic example the state is defined solely by vehicle speed.

The real world driving data described in Chapter 3 contains a wealth of information recorded at 10 Hz, which was used to construct a Markov chain describing the vehicles' real-world use. This therefore enabled the design of a SDP control policy optimised to the recorded driving data. As an example of real data Figure 5-2
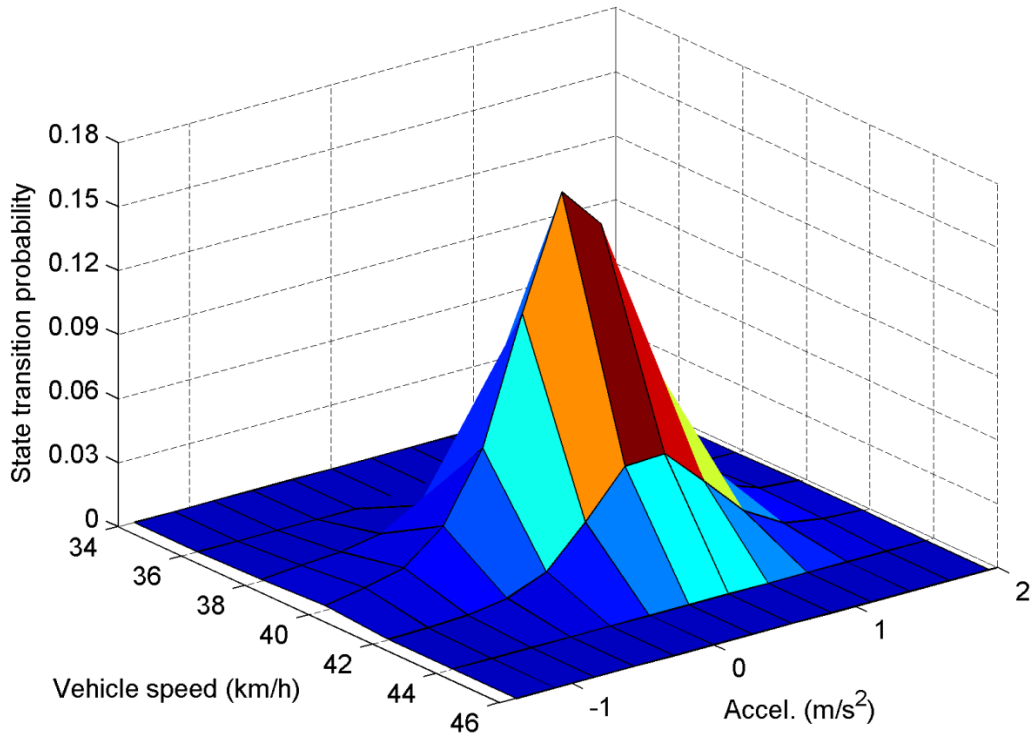
Figure 5-2: Probability distribution of the next vehicle state given the present vehicle speed and acceleration are 40 km/h and 0.56 m/s^2 in 4th gear. The surface resembles a two dimensional normal distribution centred on the present vehicle state.

illustrates the probability distribution of speed and acceleration states at the subsequent timestep given the present state of 40 km/h, 0.56 m/s^2 and 4th gear. The resulting surface broadly resembles a two dimensional normal distribution centred on the present vehicle state. Vehicle speed must be continuous in real-time and so cannot deviate greatly from the present speed. In contrast acceleration may in theory be discontinuous and therefore change significantly from one timestep to the next, but in reality rarely does so. Since the acceleration at the present timestep (calculated over the previous 0.1 s) is positive the probability distribution of speed at the next timestep has a slight skew towards speeds over 40 km/h; nevertheless the acceleration over the previous 0.1 s is no guarantee of the acceleration over the next 0.1 s, and so speeds less than 40 km/h are also possible. Finally, the surface illustrated accounts only for the state transitions in which the gear selection is unchanged, and the sum of probabilities is 0.9934; there is therefore a 0.0066 chance at each state transition that a different gear may be selected.

It should be noted that the probability distribution described in Equations (5-6) and (5-7) may actually be regarded as two probability distributions which can be

decoupled. The progression of v, a, and g depend exclusively on these same parameters at the previous timestep – the drive cycle probabilities – and not on the control decision. Conversely the progression of SOC depends only on SOC, v and the control decision at the previous timestep, and not a or g. We may therefore regard the overall state transition probability set not as one 8-D matrix, but rather as one 6-D matrix and one 4-D matrix. In summary, in order to implement this algorithm and find the optimal control policy we require three fundamental pieces of information:

(1) The instantaneous cost resulting from the combination of being in any vehicle state and making any control decision. Since the instantaneous cost is not actually dependent on the battery SOC, only on the electrical power, this cost matrix has four dimensions.

$$c_k(v_k, a_k, g_k, u_k) \tag{5-8}$$

(2) The drive cycle probability matrix. This is based on the real-world recorded drive cycle data, and defines the probability of progressing from any speed-acceleration-gear combination into another.

$$\mathbb{P}(\{V, A, G\}_{k+1} = \{v, a, g\}_{k+1} \mid \{v, a, g\}_k) \tag{5-9}$$

(3) The SOC progression matrix. This defines what the battery SOC will be in the subsequent timestep, given the SOC, vehicle speed, and the control decision made in the present timestep.

$$\mathbb{P}(SOC_{k+1} = soc_{k+1} \mid \{soc, u, v\}_k) \tag{5-10}$$

The SOC progression matrix is a probability distribution because it is extremely unlikely that the charge depletion or gain over a single timestep be enough to move the SOC exactly onto the next grid point. In practice the change will only be a tiny fraction of one grid point. Therefore the subsequent state is represented by a probability distribution split between two SOC grid points. Similarly, the subsequent driving state will depend on the model described in Equation (5-9), which will also be some probability distribution of $\{v, a, g\}$ centred around the present $\{v, a, g\}$. These two probability distributions may then be combined to obtain the likelihood of transitioning to each $\{v, a, g, soc\}$ given a control decision.

The matrices described by Equations (5-8) and (5-10) contain values which must be obtained from the vehicle model. It is worth noting that SIMULINK contains functionality to create such outputs relatively easily, without needing to simulate each point in the matrix individually. By saving the components of the vehicle model as blocks in a "Block Library" they may be re-used in a stand-alone model with the specific purpose of producing the required matrices. The blocks can be connected up within a "For Each" subsystem which accepts vectorized inputs, and so produces the matrices in one call. This approach ensures the model components used to create the SDP matrices are exactly those which form the vehicle model, and is both less time consuming and more robust that writing equivalent MATLAB code for the models.

In order to converge on the optimal control policy the algorithm is divided into two steps which are repeated alternately many times. Eventually the policy π will stop changing, though in practice it is necessary to implement some threshold at which to stop iterating. The steps are:

(1) *Policy Evaluation*: This is the evaluation of equation (5-6), which returns the cost of the present control policy, J_{π_i}.

(2) *Policy Improvement*: This is the evaluation of equation (5-7), which returns a refined control policy π_{i+1}, optimised for the latest estimate of the cost matrix calculated in the Policy Evaluation step.

5.2.1 Practical Implementation: Step-by-Step

In practice the procedure described above is not straightforward to implement because probability matrices soon become enormous. For example the full state transition probability matrix $\mathbb{P}(X_{k+1} = x_{k+1} \mid x_k)$ would be an 8-D matrix, which if stored only in single precision would occupy 326 GB of memory – far in excess of the practical RAM capacities of present day computers. Even with terminal servers, where memory limitations are less of a concern, the element-wise dot product of two matrices of these proportions would take an inordinately long period of time, and must be performed several thousand times over. By rationalising the problem a solution may be found, and the actual implementation described below makes extensive use of the most highly optimised elements of MATLAB such as sparse

matrices, matrix operations, matrix cross-multiplication, and one dimensional linear interpolation. The gradual introduction of these functionalities as the code was refined reduced computational times by several orders of magnitude, however the price is that the solution becomes more abstract and it becomes increasingly difficult to visualise the placement of dimensions.

Before proceeding the concept of sparse matrices should be briefly introduced, since this is used extensively. For an array that is mostly composed of zeroes, as the majority of the probability matrices are, it is hugely inefficient to store every element in the array. A variable stored as "sparse" does not store a complete list of every element value in the array. Rather, only the non-zero elements of the array are stored, along with the indices in the array which each element occupies. For arrays which are very space this is a much more memory efficient way of storing the variable, and so for the probability matrices used here this is an attractive option. Having the matrices stored as sparse arrays also greatly reduces the computational burden of multiplying each element in one matrix by each in another, as there are no $0 \cdot 0$ calculations taking place: the only calculations taking place are those where both elements are non-zero. For any matrix multiplication the product is at least as sparse as the factors. The final word about sparse matrices is that MATLAB restricts this class to a maximum of two dimensions for technical reasons. It is however possible to circumvent this limitation by nesting dimensions inside one another, such that a multi-dimensional matrix is represented as an extremely large 2-D matrix. This approach proved invaluable, but does reduce the transparency of the operations and means that great care is required to ensure matrices retain their integrity while having their dimensions collapsed, expanded, and permuted.

Policy Evaluation

The purpose of this step is to calculate the future cost which may be expected if the present control policy is used. For every $\{v, a, g, soc\}$ the algorithm combines the two probability matrices defined in Equations (5-9) and (5-10) to return the state probability distribution at the subsequent time step. Each probability is multiplied by the cost of being in that state, which results in the probability-weighted cost of each state. Each $\{v, a, g\}$ is considered in turn inside a "for" loop, according to the following procedure:

1. The control decisions, now only a function of SOC, are isolated:

$$u_{v,a,g}(SOC) = \pi\big(i_v, i_a, i_g\big)$$

   where i_v, i_a, i_g are the positional indices of the present speed, acceleration and gear states respectively.

2. For the set of SOC_k the probability distribution $\mathbb{P}(SOC_{k+1})$ at the subsequent timestep is calculated – this depends solely on the present control decision, since the vehicle speed is fixed. Note that for any soc_k the probability distribution of SOC at the subsequent timestep may be divided between a maximum of two values. For example, if the calculated soc_{k+1} lay exactly between two grid points the probability would be split equally between the two.

   In the example shown in Figure 5-3 for a given $\{v, a, g\}$ the control decision at soc_k is to use the EM to supply an assist torque. The result of this at the subsequent timestep is that the SOC has been depleted slightly, though because the SOC grid has a relatively coarse discretisation compared to the amount by which SOC changes in one timestep the SOC depletion is not enough to have moved a full grid space; in this example the SOC deviation is in fact one fifth of a grid space. Since the SOC is now between two grid spaces, their probability weightings at time $k + 1$ are calculated as 0.8 and 0.2 in favour of the original SOC value. For simplicity the example assumes that the same control decision is made regardless of $soc_k \in SOC_k$, and as a result the same weightings flow through the table, with the exception of $soc_k = soc_{min}$ where any assist would violate the SOC constraints. In this

$$\mathbb{P}(SOC_{k+1}) = \begin{bmatrix} 1 & 0.2 & 0 & & 0 & 0 & 0 \\ 0 & 0.8 & 0.2 & \cdots & 0 & 0 & 0 \\ 0 & 0 & 0.8 & & 0 & 0 & 0 \\ & \vdots & & \ddots & & \vdots & \\ 0 & 0 & 0 & & 0.8 & 0.2 & 0 \\ 0 & 0 & 0 & \cdots & 0 & 0.8 & 0.2 \\ 0 & 0 & 0 & & 0 & 0 & 0.8 \end{bmatrix} \begin{matrix} \\ \\ \\ \\ \leftarrow \\ \leftarrow \\ \end{matrix} \mathbb{P}(SOC_{k+1})$$

$$\qquad\qquad\quad \underset{soc_{min}}{\uparrow} \qquad\qquad\qquad \underset{soc_k}{\uparrow}$$

Figure 5-3 – Diagrammatic representation of the SOC transition matrix, which at a given $\{v, a, g\}$ is dependent only on the control decision, $\pi(soc_k)$.

case (top-left of the diagram) the control decision is therefore not to assist, and so $\mathbb{P}(SOC_{k+1}|soc_{min})$ contains soc_{min} with probability 1 and is 0 elsewhere.

3. Having obtained the probability distribution $\mathbb{P}(SOC_{k+1})$ we need also to obtain the probability distribution for $\{v, a, g\}_{k+1}$ which is comparatively simple. This is isolated by appropriately indexing into $\mathbb{P}(\{V, A, G\}_{k+1} = \{v, a, g\}_{k+1} \mid \{v, a, g\}_k)$ to return a 3-D probability matrix valid at the present values of $\{v, a, g\}_k$.

4. The two probability matrices must be combined to return

$$\mathbb{P}(X\{V, A, G, SOC\}_{k+1} = x\{v, a, g, soc\}_{k+1} \mid soc_k)$$

which is a 5-D matrix valid at the present values of $\{v, a, g\}$. Perhaps the most obvious solution for this operation would be to extend the concept of scalar expansion and replicate the probability distribution for $\{V, A, G\}_{k+1}$ by the length of SOC in the 4th dimension, and then again in the 5th dimension. Each replica of the original 3-D matrix could then be multiplied by the elements of $\mathbb{P}(SOC_{k+1})$ to give the desired 5-D probability matrix. Although relatively transparent this approach is extremely computationally inefficient and does not take advantage of much of the built-in optimisation of MATLAB. Since $\mathbb{P}(SOC_{k+1})$ is mainly zeroes (see Figure 5-3), as is $\mathbb{P}(\{V, A, G\}_{k+1})$, the operation can be completed far more quickly using sparse matrices.

Furthermore, since we are actually seeking to multiply every element in $\mathbb{P}(\{V, A, G\}_{k+1})$ by every element in $\mathbb{P}(SOC_{k+1})$ we can use the properties of the matrix cross-product. Therefore both matrices are re-arranged into 1-D arrays, where their second (and subsequent) dimensions are nested inside the first dimension. Both arrays are then converted into sparse arrays, and the cross product is carried out as shown in Figure 5-4.

Figure 5-4 - Diagrammatic representation of the calculation of $\mathbb{P}(X_{k+1} = x_{k+1} \mid soc_k)$ using matrix cross-product and sparse matrices.

This procedure is repeated in a loop for each speed, acceleration and gear, and for each repetition the resulting $\mathbb{P}(X_{k+1}|soc_k)$ is stored in a cell array. It is important to keep track of the number and location of dimensions in the result, as this has quickly become rather abstract: We now possess a cell array with a length of $n_v \cdot n_a \cdot n_g$, where n denotes the length of each state vector. Each cell in the array pertains to one combination of $\{v, a, g\}_k$ and contains a matrix representing the probability of progressing into any $\{v, a, g, soc\}_{k+1}$ starting from any soc_k, i.e. $\mathbb{P}(\{V, A, G, SOC\}_{k+1} \,|\, \{v, a, g, soc\}_k)$. Although each matrix should have five dimensions, these have in fact been collapsed into just two to allow storage as a sparse matrix; the first dimension contains $\{v, a, g\}_{k+1}$ nested inside one another, and the second contains soc_{k+1} nested inside soc_k.

5. Once again, we use a "for" loop to cycle through each combination of $\{v, a, g\}$. Inside the loop, the corresponding $\mathbb{P}(X_{k+1}|soc_k)$ matrix is retrieved from the cell array, and the dimensions are rearranged so as soc_{k+1} is moved into the first dimension, and the second dimension then only contains soc_k. This matrix now has dimensions of $n_v \cdot n_a \cdot n_g \cdot n_{soc} \times n_{soc}$. The future expected cost, $J_\pi(x_{k+1})$, is also rearranged into row vector form, so as to have dimensions $1 \times n_v \cdot n_a \cdot n_g \cdot n_{soc}$. These two matrices may now be combined with a cross product, where the probability of arriving at state X_{k+1} at the next timestep, having started from soc_k, is multiplied by the future-expected cost of moving onward from that state. This yields a $1 \times n_{soc}$ array containing the future cost which each present state soc_k may be expected to incur from the next timestep onward under the control policy π.

$$
\underbrace{\begin{bmatrix} \cdots & \cdots & \cdots \end{bmatrix}}_{\substack{J_\pi(soc_k) \\ \text{dim.}\,[1 \times n_{soc}]}} = \underbrace{\begin{bmatrix} \cdots & \cdots & \cdots & \cdots & \cdots \end{bmatrix}}_{\substack{J_\pi(\{v,a,g,soc\}_{k+1}) \\ \text{dim.}\,[1 \times n_v \cdot n_a \cdot n_g \cdot n_{soc}]}} \times \underbrace{\begin{bmatrix} & \cdots & \\ \vdots & & \vdots \\ \vdots & \ddots & \vdots \\ \vdots & & \vdots \\ & \cdots & \end{bmatrix}}_{\substack{\mathbb{P}(X(v,a,g,SOC)_{k+1}|\,SOC_k) \\ \text{dim.}\,[n_v \cdot n_a \cdot n_g \cdot n_{soc} \times n_{soc}]}}
$$

Figure 5-5: Calculation of the future expected cost resulting from following control policy π at each soc_k.

6. The final consideration in the policy evaluation step is the cost incurred instantaneously during the progression from X_k to X_{k+1}. In order to account for this the set of control decisions $u_{v,a,g}(SOC)$ isolated in step (1) is used to retrieve the instantaneous cost $c(x_k, u)$ for each SOC. This is arranged as a $1 \times n_{soc}$ array, and may be summed with the future expected cost $J_\pi(soc_k)$ to create an updated estimate of the expected cost $J_\pi(x_k)$, as defined in equation (5-6).

The process of updating the estimate of future expected cost is recursive: J_π appears on both sides of equation (5-6) and so the solution is approached by iterating several times. Since each iteration is not actually calculated is one pass, but rather consists of the "for" loops described above, the speed of convergence is increased slightly by updating the values used in the old J_π during the "for" loops, i.e. mid-iteration.

For the calculation of J_{π_0} an initial control policy π_0 is required, as well as an initial estimate for the future expected cost, $J_{\pi_0}(x_0)$. These are both set to be entirely zeroes, i.e. neither assist nor regenerate, and the future cost is zero. The calculation of J_{π_0} is therefore time consuming as it must be constructed from nothing – iterations were stopped once every element in the matrix was changing by less than 0.1% per iteration. In subsequent Policy Evaluation steps, after changing the control policy, a minimum of 20 iterations were performed.

Policy Improvement

The purpose of this step is to refine the control policy, approaching the optimal one for the probability-weighted future cost calculated in the Policy Evaluation step. Since the instantaneous cost of any control decision is easily calculated, and the effect this decision has in determining the next state is known, the expected future cost of any decision may be found from J_π. This step is therefore simply a case of picking out the control decision which results in the minimum total (instantaneous plus future) cost for each vehicle state.

The techniques and processes used in the programming of the Policy Improvement step are similar to those described above with regard to the exploitation of sparse matrices and matrix cross-products. The implementation of this step is slightly more

straightforward in that the equation is not recursive, and therefore the optimal control decision at any state x_k is independent of the control decision made in any other state. Therefore at the most fundamental level the task at hand is simply, at every possible state, to evaluate the cost of each possible control decision and return the decision which minimises the cost. Though the more sophisticated MATLAB built-in optimisation functions were trialled it was found to be faster simply to evaluate the cost of every possible control decision by direct enumeration, since the discretisation of the control vector was relatively coarse and so the set of possible control decisions in any state is small.

Although the essence of the problem is simple it is advantageous to increase the number of dimensions being handled simultaneously so as to reduce the number of nested loops, making best use of MATLAB's built-in matrix optimisations. Therefore the *gear* and *SOC* dimensions are processed simultaneously, and the following procedure is repeated inside "for" loops for each *speed* and *acceleration*.

1. J_π is rearranged to have $\{v, a, g\}_{k+1}$ in the first dimension and soc_{k+1} in the second dimension (in fact this is invariant and so is only rearranged once, outside of the loops).

2. The probability distribution of *speed*, *acceleration* and *gear* at the next timestep is isolated from the drive cycle probability matrix: $\mathbb{P}(\{V, A, G\}_{k+1} = \{v, a, g\}_{k+1} \mid g_k)$. This is rearranged so as to contain g_k in the first dimension and $\{v, a, g\}_{k+1}$ nested inside one another in the second dimension. This is converted into a sparse matrix.

3. The probability distribution of *speed*, *acceleration* and *gear* at the next timestep (from step 2) is cross-multiplied by the future expected cost of each state (step 1) as shown in Figure 5-6. The result is a 2-D matrix representing the future expected cost as a function of the present gear, g_k, and the SOC at the next timestep, soc_{k+1}.

$$
\underbrace{\begin{bmatrix} & \cdots & \\ \vdots & \cdots & \vdots \\ & \cdots & \end{bmatrix}}_{\substack{J_\pi^{\{v,a\}_k}(g_k,soc_{k+1}) \\ \text{dim. } [n_g \times n_{soc}]}}
=
\underbrace{\begin{bmatrix} & \cdots & \cdots & \cdots & \\ \vdots & \cdots & \cdots & \cdots & \vdots \\ & \cdots & \cdots & \cdots & \end{bmatrix}}_{\substack{\mathbb{P}(\{V,A,G\}_{k+1}=\{v,a,g\}_{k+1}\mid g_k) \\ \text{dim. } [n_g \times n_v \cdot n_a \cdot n_g]}}
\times
\underbrace{\begin{bmatrix} & \cdots & \\ \vdots & & \vdots \\ \vdots & \ddots & \vdots \\ \vdots & & \vdots \\ & \cdots & \end{bmatrix}}_{\substack{J_\pi(\{v,a,g,soc\}_{k+1}) \\ \text{dim. } [n_v \cdot n_a \cdot n_g \times n_{soc}]}}
$$

Figure 5-6: Weighting the future expected cost at *k*+1 by the probability distribution of $\{V,A,G\}_{k+1}$ to give the future expected cost as a function of g_k and soc_{k+1} based on present vehicle speed and acceleration, $\{v,a\}_k$.
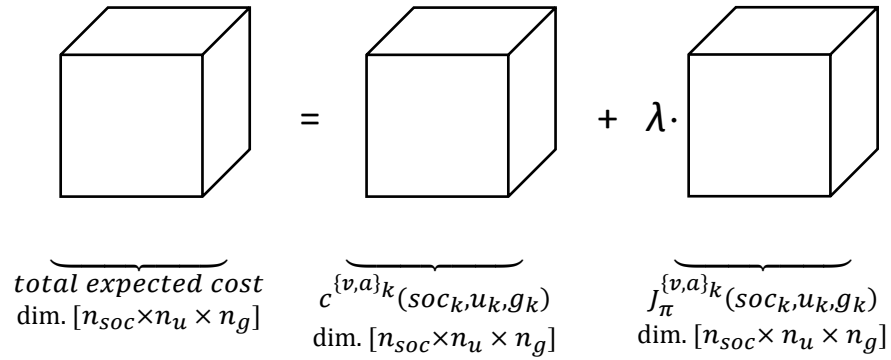
4. From the SOC progression matrix the probability distribution of SOC at the next timestep is isolated for the present speed, such that it now only depends on the present SOC and the control decision:
   $\mathbb{P}(SOC_{k+1} = soc_{k+1} \mid \{soc, u\}_k)$. This may then be cross-multiplied with the transposed result of the previous step to give the future expected cost of any *SOC-control-gear* combination at the present timestep, for $\{v, a\}_k$.

$$
\underbrace{\begin{bmatrix} & \cdots & \\ \vdots & \cdots & \vdots \\ & \cdots & \end{bmatrix}}_{\substack{J_\pi^{\{v,a\}_k}(soc_k,u_k,g_k) \\ \text{dim. } [n_{soc} \cdot n_u \times n_g]}}
=
\underbrace{\begin{bmatrix} & \cdots & \\ \vdots & & \vdots \\ \vdots & \ddots & \vdots \\ \vdots & & \vdots \\ & \cdots & \end{bmatrix}}_{\substack{\mathbb{P}(SOC_{k+1}=soc_{k+1}\mid \{soc,u\}_k) \\ \text{dim. } [n_{soc} \cdot n_u \times n_{soc}]}}
\times
\underbrace{\begin{bmatrix} & \cdots & \\ \vdots & & \vdots \\ & \cdots & \end{bmatrix}}_{\substack{J_\pi^{\{v,a\}_k}(soc_{k+1},g_k) \\ \text{dim. } [n_{soc} \times n_g]}}
$$

Figure 5-7: Weighting the future expected cost of all SOC_{k+1} by the probability distribution of SOC_{k+1} to give the expected cost of any *SOC-control-gear* combination.

5. The instantaneous cost is isolated at the present *speed* and *acceleration*, so as it is now only a 2-D matrix with dimensions of *gear* and *control decision*. Since instantaneous cost is not a function of SOC this is simply replicated in a third dimension n_{soc} times to represent the cost of each *gear-control-SOC* combination at the present $\{v, a\}_k$. The cost is not dependent on SOC, but this is necessary so that the dimensions can be rearranged to match those of the result of the previous step. Instantaneous cost and future expected cost may then be summed directly as shown in Figure 5-8, where the future expected cost is discounted by λ.

$\underbrace{total\ expected\ cost}$
dim. $[n_{soc} \times n_u \times n_g]$

$\underbrace{c^{\{v,a\}_k}(soc_k,u_k,g_k)}$
dim. $[n_{soc} \times n_u \times n_g]$

$\underbrace{J_\pi^{\{v,a\}_k}(soc_k,u_k,g_k)}$
dim. $[n_{soc} \times n_u \times n_g]$

Figure 5-8: Adding the instantaneous cost to the discounted future expected cost gives the total expected cost of each control decision. The decision yielding the minimum total expected cost may then be selected for each $\{soc, g\}_k$.

6. From the result of the previous step the control decision yielding the minimum total expected cost for each $\{soc, g\}_k$ at the present $\{v, a\}_k$ may be read out directly; that is: $\pi_{i+1}(x_k) = \arg\min_u\{total\ expected\ cost\}$. Steps (1)-(6) are repeated for each *speed-acceleration* point.

The policy evaluation and policy improvement steps described must be repeated many times – in this case several thousand times – to converge on the optimal solution. However, the convergence on a solution and the quality of the output depend hugely upon appropriate parameter selection, and during this work practical advice on this was found to be scarce. The methods and plots presented in the following section were found to be informative during the development process.

5.2.2 Choice of Discount Factor

Perhaps the most important parameter to select appropriately is the discount factor, λ, and this issue is not explored or widely reported on in the existing literature, as discussed in Chapter 2. Since this defines how quickly the future costs are discounted, it essentially determines how *future looking* the control strategy is; for example, considering equation (5-5), when $\lambda=0$ the strategy only considers the immediate cost of its decision. In contrast if $\lambda=1$ the strategy would consider the cost at every timestep in the future with equal weighting, and therefore attempt to look infinitely far into the future.

A strategy which discounts future costs quickly will be relatively unaware of the costs required to replenish any battery energy consumed in the present; for this reason strategies built with a low λ tend to be charge depleting. Conversely, strategies with λ→1 are more aware of energy balance costs and so are more charge sustaining. However, although the quality of the solution is better with higher values of λ the number of iterations required increases exponentially. For this reason higher values of λ carry the burden of considerably more computational effort before the solution converges and computation times quickly escalate as a result.

To achieve charge sustaining behaviour it is not mathematically necessary to place a cost on deviation of the SOC from a target or nominal value, as charge sustenance is achievable by selection of an appropriately high value of λ. Nevertheless other researchers have added a cost on SOC deviation, primarily because of other desirable effects such as reducing the likelihood of the SOC being very low when the vehicle is switched off [75], though this may also have been motivated by its enabling the use of a smaller λ. In the ideal case λ would be sufficiently high to ensure the control strategy is adequately charge sustaining, but no higher, therefore not incurring additional computation time unnecessarily.

Determining and examining the charge sustaining properties of a control strategy are not necessarily straightforward, as the response of any control strategy will depend greatly on the drive cycle to which it is subjected. Rather than follow the widely used approach of simulating the controller's response to a particular drive cycle a more advanced method of examining the controller's behaviour was developed in furtherance of the probabilistic nature of the SDP approach. In this approach the probabilistic ratio of positive to negative DC current at each SOC was calculated, in the hypothetical scenario where the SOC is invariant. In other words, this is the ratio of assist to regeneration which is likely to occur at each SOC, calculated by multiplying the probability of each $\{v, a, g\}$ combination with the controller's response in that state, and dividing the sum of positive values by the sum of negative values:

$$Q_I(soc_i) = \frac{\sum \mathbb{P}(v,a,g) \cdot (\pi(v,a,g,soc_i) > 0)}{\sum \mathbb{P}(v,a,g) \cdot (\pi(v,a,g,soc_i) < 0)} \;,\quad i = 1, 2, \ldots n_{soc}. \qquad (5\text{-}11)$$

Where this ratio is greater than 1 the strategy is generally charge depleting, whereas at values less than 1 it is charge gaining over a drive cycle with the same statistical distributions of $\{v, a, g\}$ as the data on which it was based.

Selecting an appropriate value of λ was very much a trial-and-error process and a range of values were experimented with starting from 0.95, as used by several other researchers, and progressively working towards higher values. It became apparent very early on that for this application a value of λ considerably higher than those reported by others was necessary. Figure 5-9 presents the ratio Q_I varying with SOC for a range of λ values, and the importance of λ in determining the equilibrium SOC is evident. These plots shall be referred to as *probabilistic charge sustenance plots*, and the SOC at which the ratio is exactly 1 is the *probabilistic nominal SOC*. It can be seen in Figure 5-9 that increasing λ up to a value of 0.999,999 has a considerable
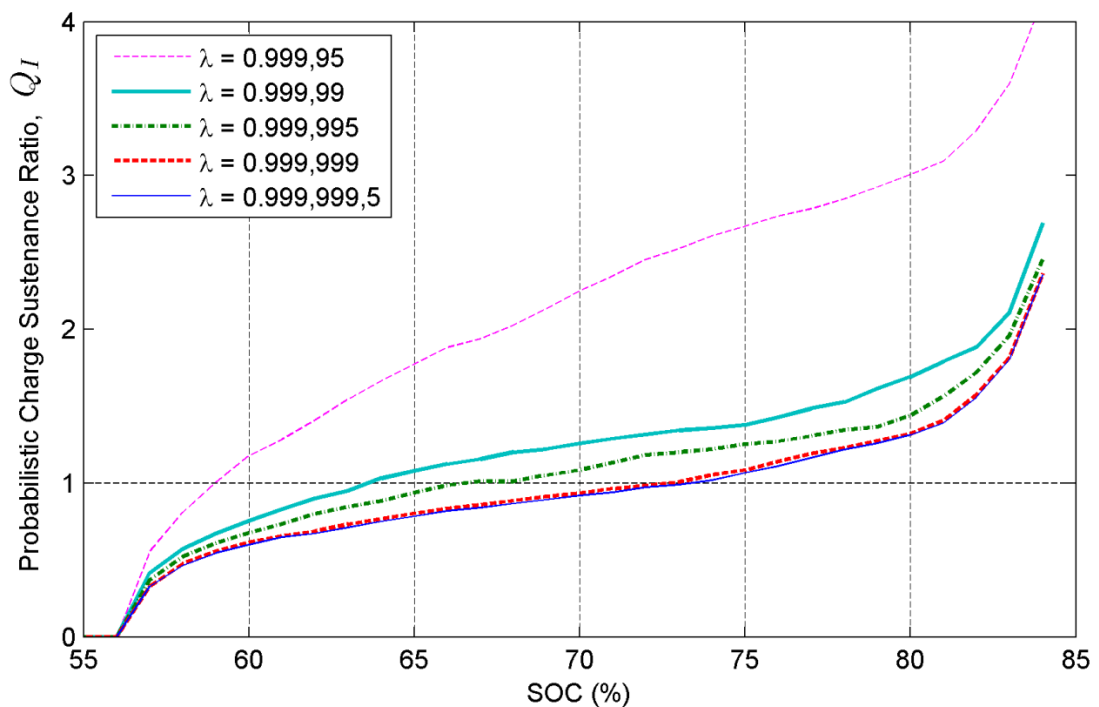

Figure 5-9: Probabilistic charge sustenance plot for different values of λ, showing the ratio of positive to negative current demand as a function of SOC. The control strategy is nominally charge sustaining where $Q_I = 1$. The variation of this function with λ is a result of how "future looking" the strategy is.

effect of the probabilistic nominal SOC, but increasing the value above this does not change the response, and therefore this value was selected.

As identified in Chapter 2 previous researchers have reported using values of λ in the range 0.95-0.995. The considerably higher values required here are most likely because of the use of 10 Hz data where others have used 1 Hz data, meaning the algorithm needs to consider 10 times more data samples and degrade 10 times slower in order to consider the same physical timeframe as the versions implemented by others. Since driving data was available at 10 Hz it was used at this frequency throughout, and not down-sampled, because reduced frequencies would result in considerable loss of information regarding vehicle energy [76].

5.2.3 SOC Adjustment

Probabilistic charge sustenance plots may also prove useful where developers wish to adjust the nominal SOC that a controller operates at. For example, if the $Q_I = 1$ intercept is at 73% SOC, but it is desirable for other reasons that the nominal SOC is 78%, the SOC reported to the controller may simply be offset by 5%. This simple observation offers a great deal of clarity and robustness to the process of developing SDP controllers, which can otherwise be obscured by the difficulty visualising and interrogating the large number of controller dimensions.

5.2.4 Number of Iterations

Complete convergence of the future expected cost J_π would take an extremely long time, however since small changes in the shape of the cost function do not necessarily effect a change in the shape of the control strategy map it may be possible to truncate the iterations at some point; in practice therefore we may permit ourselves to terminate iterations once the shape of the control strategy map is stable. This raises the question of how the truncation threshold, or stability criterion, is defined. Simply comparing successive policies and waiting for them to become identical would take an enormous number of iterations, and does not necessarily give the best representation of the amount of change being undergone. In this work it was observed that there was often some amount of cyclical change where policies cycled through some repeatable series of changes, while the net change over several

iterations was zero. For this reason it was found to be more helpful to define the maximum percentage change of any element in J_π as ε, and curtail iterations when this fell below a threshold. The sum of the absolute difference between the final control strategy π_N and each π_i was then counted and divided by the grid space size to find the number of grid space differences, $n_\Delta = \frac{|\pi_N - \pi_i|}{17}$, used to ensure that no cumulative change was taking place. Note that each value in π may move by more than one grid space. This method was used for a very long sequence of π iterations to investigate at what point cumulative change stopped, and from this information the value of ε was chosen as 0.0008%, which was considered low enough to leave a safe margin between the end of any cumulative change and the conclusion of iterations.

Figure 5-10 shows an example of the exponential convergence of J_π. For λ=0.999,999 the threshold is reached after 5719 iterations of π (note that there are a minimum of 20 iterations of J_π for each iteration of π) and there is no cumulative change in π by this point.

The implementation of SDP achieved during this work took about 30 seconds for each policy evaluation step (which includes 20 iterations of J_π) and a similar amount of time for each policy improvement, running on a desktop computer with an Intel Core i7 CPU at 2.0 GHz. Therefore the evaluation of 5719 iterations of π represents approximately 4 days of contemporary desktop computing. Since the rate of convergence is considerably slower with higher values of λ, and quicker with lower values of λ, it is easy to appreciate the desire to use a λ which is no larger than necessary.

The final result of the SDP iteration is a stationary policy in the form of a multi-dimensional look-up table. In this case the look-up table has four dimensions, $\{v_k, a_k, g_k, soc_k\}$, and Figure 5-11 presents a cross-section of what this looks like.
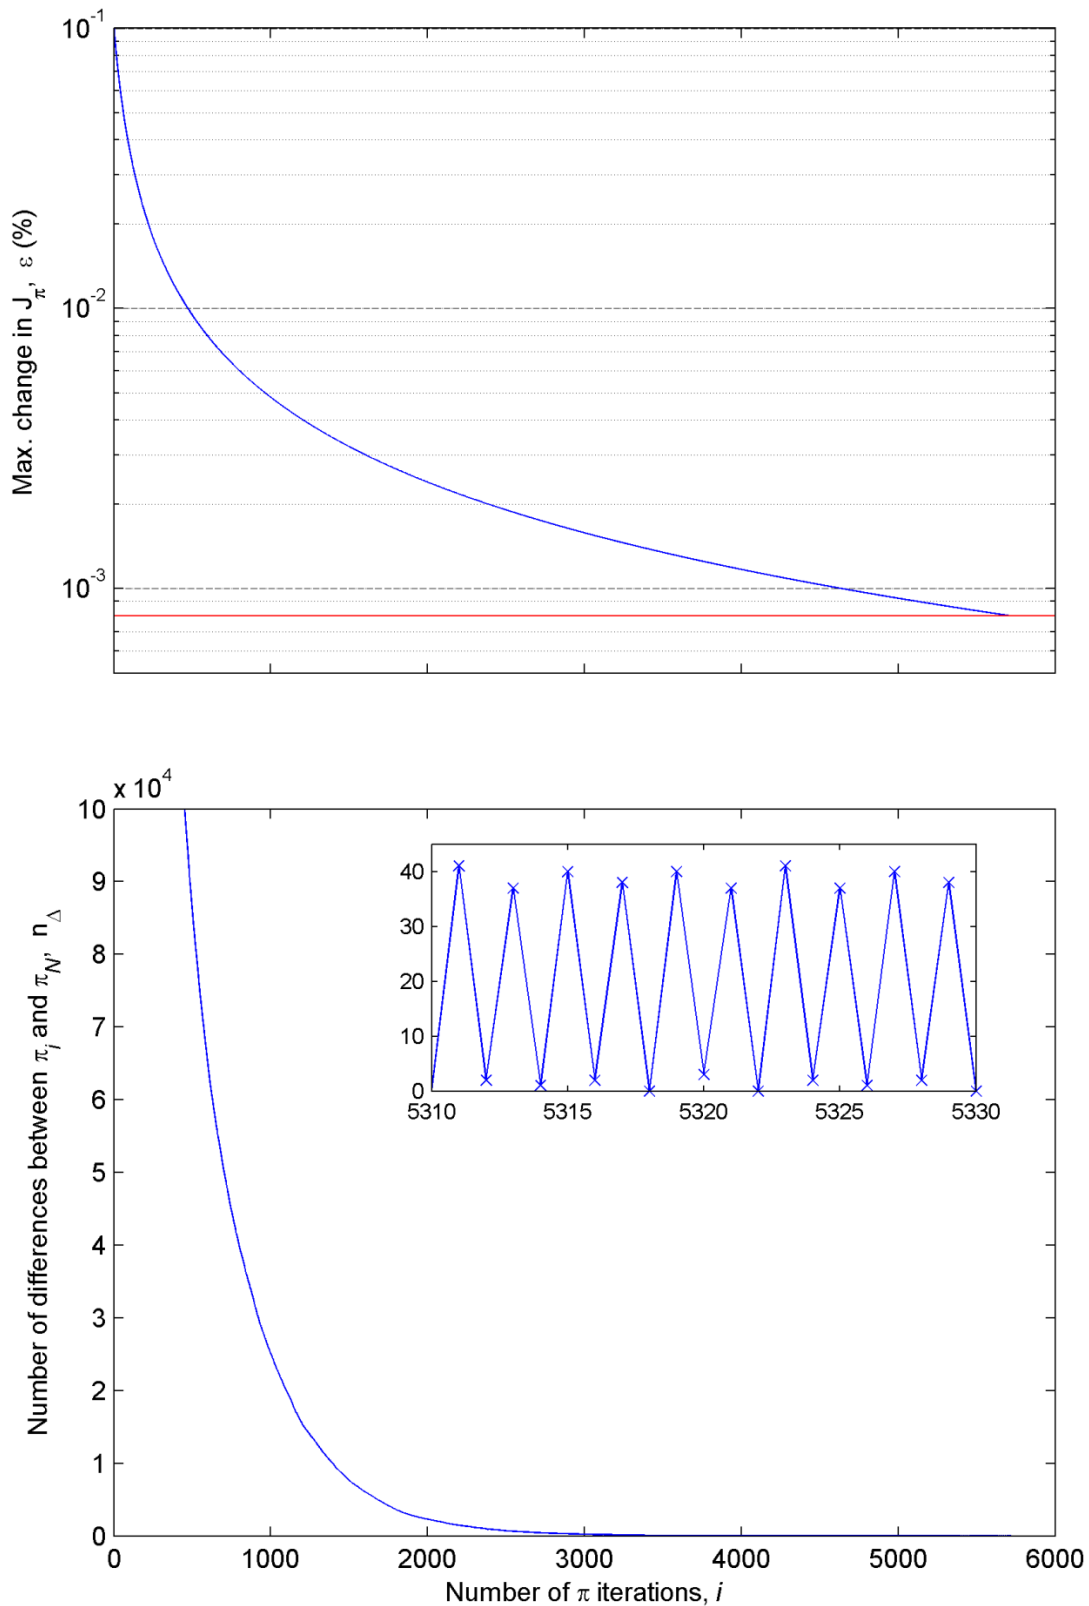
Figure 5-10: Convergence of the future expected cost function and the control policy based on this cost. Iterations are truncated when $\varepsilon \leq 8 \times 10^{-4}$ %. The inset axis shows that at this point the changes in the control policy are cyclical and not cumulative; further iterations are unlikely to yield any significant changes to the salient shape of the policy.
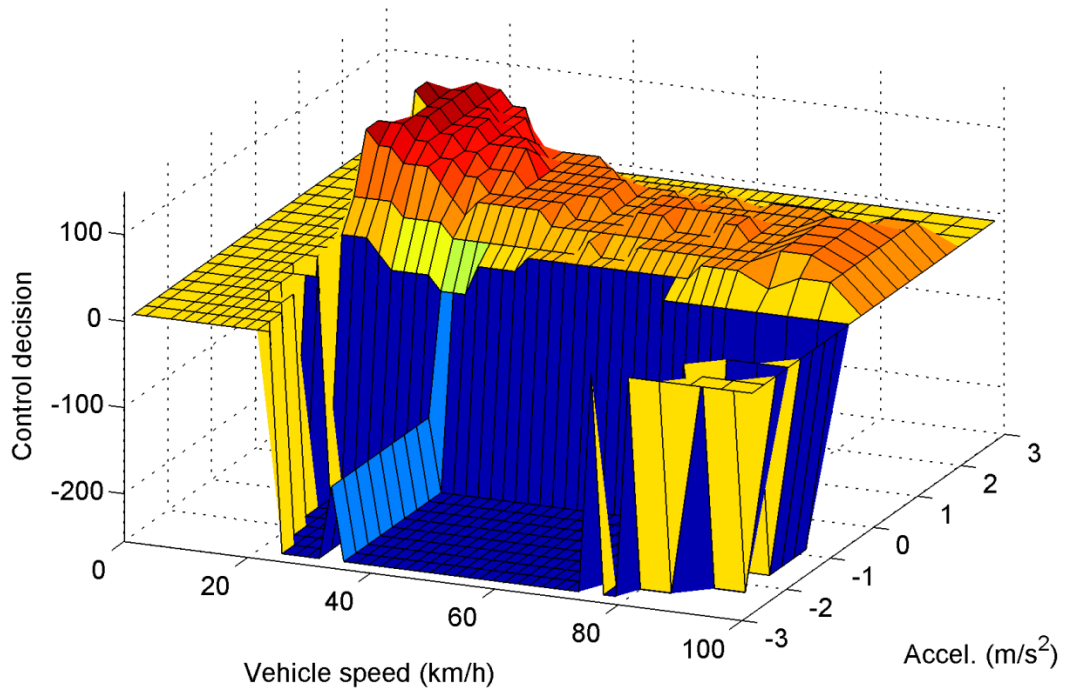
Figure 5-11: SDP control policy surface as a function of vehicle speed and acceleration in 4th gear, at 70% SOC, with α = 0.001. The control decision is a unitless number in the range ±255 which defines the battery assist or regeneration current as a proportion of the maximum available at the present vehicle speed.

5.2.5 Robustness to Varying Driving Patterns

Section 5.2.2 first introduced the idea that the demand of the SDP controller for each SOC can be weighted by the probability of each state occurring. When expressed as a ratio of positive to negative current the result gives an indication of how charge sustaining the strategy is as a function of SOC. Naturally the ratios of positive to negative current demand and the inferred nominal SOC are only good indicators of the controller's average behaviour over an extended period, assuming that the stochastic distributions of the actual driving are close to those used in the design of the controller. It is of course foreseeable that there will be differences between the stochastic distributions of real driving and those historic ones used to design the controller. These differences may be considered under two categories:

(1) Real-time variability: When considering short periods of driving it is inevitable that the distributions observed on the macro scale no longer hold true. As the period under consideration is shortened the variations between

samples will increase, as too will the difference between the macro trends and the trends of each sample.

(2) Paradigm shift: It is possible that the average driving pattern of the vehicle will change permanently, perhaps as a result of the vehicle's purpose being altered, the driver adjusting their style slightly or a different driver altogether.

Both of these situations will result in deviation from the probabilistic nominal SOC, either fleetingly or more permanently. With respect to real-time variability the deviation from the nominal SOC is no cause for concern – on the contrary if there were no deviation this would simply indicate that the system was entirely inactive. The maximum and minimum SOC define the extremes of the operating window and the control strategy will never attempt to operate outside of these limits.

With regard to more fundamental changes in the average use of the vehicle over time it is certainly true that some degree of sub-optimality may develop and that the nominal SOC may drift. For this reason it is interesting to examine the controller's behaviour when a stochastic drive cycle other than that used in its design is assumed. For this purpose the probabilistic charge sustenance characteristics of the SDP controller were calculated using the same procedure as previously set out, but using stochastic driving data from each of vans 4, 6 and 7 in isolation. Each curve therefore describes the controller's behaviour in the situation where the composition of the actual drive cycle differs from the historic set for which the controller was designed. In these cases the differences between actual and historic driving patterns are not enormous since the historic set is actually the combination of the three vans. The situation is therefore illustrative of the condition where a minor change to the vehicle's role has been made, or where the driver has changed. The results are presented in Figure 5-12 and indicate that in such a situation the ratio of energy recovered to energy deployed can change significantly, therefore moving the probabilistic nominal SOC.

In the case of using only Van 4 data Figure 5-12 suggests that less recoverable energy is available in the driving cycle and/or the states resulting in assist events are frequented more often. As a result of this the nominal SOC has fallen to around 63%.

In contrast the opposite is true of Van 6 data, leading to the nominal SOC rising to 81%, while using only Van 7 data seems not to impact the behaviour a great deal. In all cases the responses are considered robust because the controller naturally adjusts to the availability of energy to achieve a charge sustaining behaviour without resting on the limits of the acceptable battery SOC.

Despite the strategy being charge sustaining there may be reasons why it is undesirable for the nominal SOC to have significantly shifted. For example, if the nominal SOC is reduced then the probability of the vehicle being switched off and leaving the battery at a low SOC for an extended period is increased. This may have consequences for the battery's self-discharge rate and subsequent SOH when the vehicle is restarted. If the SDP technique were to be adopted it is likely that this possibility would need to be addressed, though it is unlikely to be problematic. The ideal solution would be to continuously record actual driving data and to refine the optimal control surface based on the most recent use of the vehicle, perhaps when the vehicle is not in use. Should this approach not be feasible (due to hardware performance limitations, for example) simpler adjustments may be made at a higher
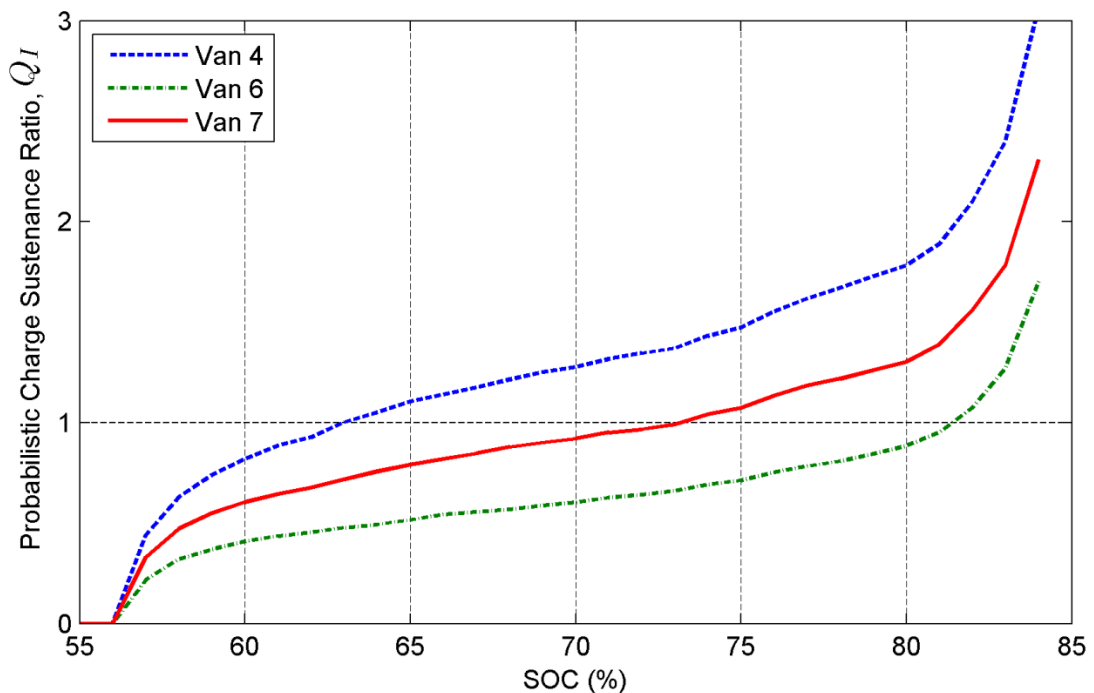

Figure 5-12: Probabilistic charge sustenance plots for the SDP controller when the actual stochastic driving data is based on each van in isolation. The controller is optimised for the combined data of all 3 vans and so its behaviour changes if the actual stochastic driving data differs from the design set.

level in the controller, for example by monitoring the nominal SOC and applying an offset to the SOC reported to the SDP control sub-routine, thereby readjusting the nominal SOC as desired.

5.3 Implementation of ECMS

To properly evaluate the performance and potential of the SDP control strategy best practice would suggest comparison with another accepted and well regarded approach to control; for this reason a comparable ECMS policy was developed. The basis of ECMS optimality is to minimise the instantaneous cost, without any regard to the potential future costs which the present decision might incur. Although in principle this instantaneous optimisation approach is inferior to the global optimisation offered by SDP, the ECMS may in fact perform almost as well, and offers other advantages since it is relatively easy to implement, transparent, and easy to work with. As with other approaches the cost function may be composed of multiple terms, and in ECMS terminology these are weighted using *equivalence factors* which define an equivalent price for the different elements of the cost function in units of the primary cost – fuel.

In the case when the ECMS is deployed in control of a plug-in hybrid vehicle the electricity available in the batteries has an associated physical cost, this being the price paid to charge the batteries from the grid. However in the present case, where much of the energy stored in the battery is recovered from regenerative braking and therefore has no monetary cost, the appropriate price on electricity consumption is not so clear. Leaving the consumption of electrical energy from the batteries without an equivalent cost would simply result in continual use of the electric powertrain to assist, as this always reduces the instantaneous fuel consumption. It is therefore necessary to impose some cost on electricity consumption by means of an equivalence factor that ensures charge sustaining behaviour. For the ECMS control strategy the cost function implemented may then be written as

$$c\big(x_k, \pi(x_k)\big) = f(v_k, a_k, g_k, u_k) + \alpha \cdot C^2 + \gamma \cdot C . \qquad (5\text{-}12)$$

As for the SDP cost function the first two terms represent instantaneous fuel consumption and the cost for high power electrical operation respectively. The third term introduces equivalence between the consumption of electricity stored in the batteries and the consumption of fuel, with equivalence factor γ. It should be noted that the second term, penalising high power use of the electric powertrain, is always positive regardless of the direction of energy flow. In contrast the third term representing the instantaneous cost of electricity consumption can be positive or negative, becoming negative when the batteries are being charged.

Computation of the ECMS control map is far less complicated than that of the SDP, and is adequately described elsewhere, so will not be covered in the same length. Using the same "For Each" models as in the SDP development a cost matrix was derived for each of the three cost terms in the cost function. The cost matrices for high power operation and electricity consumption were expanded as necessary to have dimensions $\{v_k, a_k, g_k, u_k\}$, matching the dimensions of the instantaneous fuel cost. For specified values of the two equivalence factors the three cost matrices were then combined into a single matrix, defining the cost of each control decision in any vehicle state $\{v_k, a_k, g_k\}$. The control map was then obtained simply by retrieving
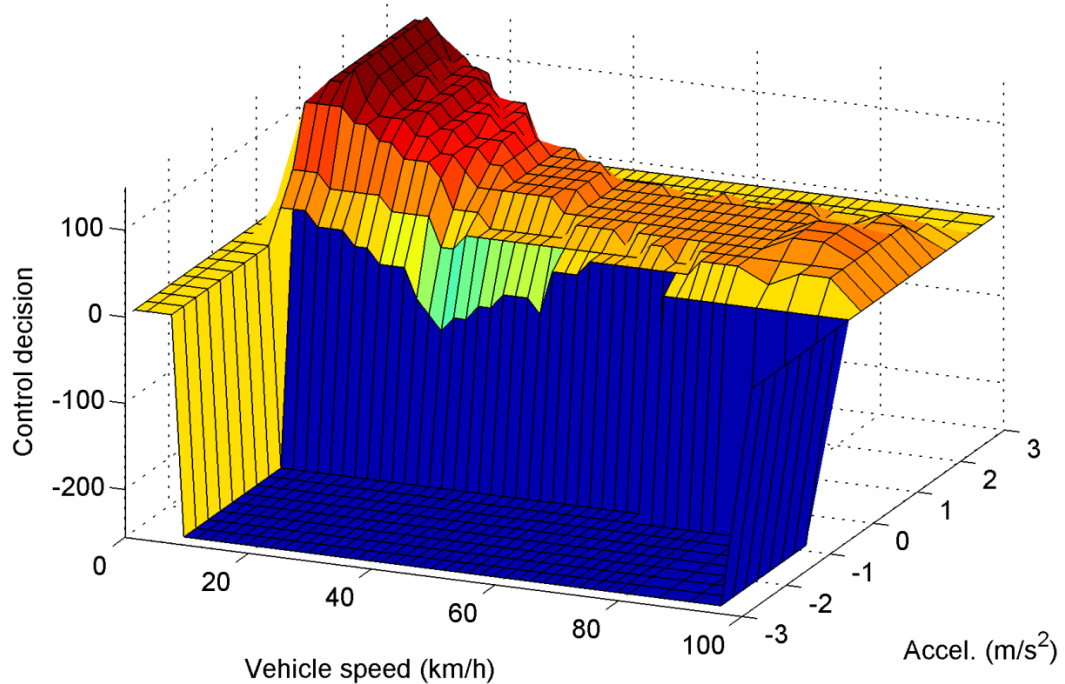

Figure 5-13: ECMS control policy surface as a function of vehicle speed and acceleration in 4th gear, with
α = 0.001. The control decision is a unitless number in the range ±255 which defines the battery assist or
regeneration current as a proportion of the maximum available at the present vehicle speed.

the value of the control decision which minimised the combined cost in each vehicle state. As with the SDP controller all of this was computed off-line such that the resulting control policy could simply be implemented as a look-up table, an example of which is shown in Figure 5-13.

Selection of γ values was made by using the simulation and a golden search optimisation routine to vary γ with the aim of minimising SOC change over a LA92, with the target being to achieve less than ±0.5% SOC difference. This optimisation was carried out for each value of α used, since the control strategies resulting from different α naturally had different charge sustaining behaviour.

5.4 Controller Architecture

Both the SDP and ECMS algorithms output multi-dimensional state loop-up tables which define the optimal response of the hybrid powertrain to the present vehicle state. These look-up tables were implemented as the core decision making block in the centre of a much larger control system architecture which incorporated a substantial amount of low level functionality, both upstream and downstream of the look-up table. A great deal of this low level functionality comprised relatively trivial but important tasks such as input and output filtering and processing, however some more advanced features such as inference of the vehicle context and driver intentions were also necessary for safety and drivability reasons. Functionality outside of the core maps was common between all control policies tested.

One of the most basic examples of safety critical contextual logic is the assurance that when the brake pedal is activated by the driver the demand of the hybrid controller can only be negative (regeneration); it should never be permissible to deliver positive (assist) torque during braking, as this is in direct contravention of the driver's request. By the same token, when the engine is operating very close to its maximum load this may suggest the vehicle is struggling to climb a steep hill or that a rapid acceleration is necessary, when on a slip road preparing to join a motorway for example. In such situations it may be dangerous to demand regeneration, and so this is forbidden.

Whenever the brake pedal is activated it is safe to assume the driver's intention is to slow the vehicle even if this is not adequately reflected in the vehicle's acceleration, due to a gradient for example. For this reason it would seem logical that during braking the control demand should be maximum regeneration, however since the cost function used in this work aims to evaluate different levels of positive and negative aggressivity such crude conditional logic would not be appropriate. Instead, the acceleration input to the control map was overridden with the minimum acceleration defined in the state space, so that the controller's output became its natural response to strong deceleration at that speed, gear and SOC. By following this logic the rate of energy recovery during braking was respectful of the controller's other objectives.

A more subtle example of the contextual logic implemented is the detection of gear upshift events during accelerations, which is of interest for drivability reasons. During any gear shift event the clutch is pressed, mechanically disconnecting the engine from the wheels, hence ensuring that there is no tractive power. As a result of the absence of tractive power the vehicle will immediately begin to decelerate, which may cause the controller to demand regeneration. Although the response of the controller in this situation is entirely logical the inputs to the controller are somewhat misleading as it is not the driver's intention to slow down, nor is regeneration an appropriate control demand. Recovering kinetic energy during an upshift would cause the speed to fall further, therefore requiring more work to re-accelerate the vehicle as soon as the next gear is engaged. In order to mitigate this scenario the control logic monitors the clutch signal, and if activated during an acceleration event triggers an upshift flag; in this case the output of the look-up table is passed through a low pass filter so that the assist torque which was applied during the preceding acceleration is continued through the upshift event, only decaying to about half its magnitude at the end of the shift. For safety reasons this flag expires after two seconds and is overridden by several other use cases, as may be expected, such as brake pedal activation.

5.5 Chapter Conclusions

A set of SDP controllers has been developed in this chapter, which are each designed to operate the HEV powertrain optimally for minimisation of fuel consumption whilst utilising the electric powertrain at different levels of aggressivity, controlled by a weighted term in their cost functions. Each SDP controller is based on realistic stochastic driving distributions extracted from the data collected in Chapter 3, and should therefore be optimised to driving patterns representative of the vehicles' typical real-world usage.

The development of the controllers has been described in detail, with practical considerations and limitations explained and addressed. In particular the management of SOC is not enforced explicitly in the cost function, but is instead managed by appropriate choice of the discount factor, λ. The effect of the discount factor on charge sustenance was therefore examined and methods developed to examine this graphically, thereby aiding selection of an appropriate value.

ECMS controllers were also developed, primarily as a benchmark. The controllers developed in this chapter are implemented both in the whole-vehicle simulation developed in Chapter 4, as well as on a real vehicle, and performance results reported in Chapter 6.

Chapter 6   Hybrid Control Results

In this chapter the HEV controllers are tested, first in simulation, then under controlled conditions on a chassis dynamometer. The trade-off between fuel consumption savings and electrical system stress is investigated through testing of control strategies with a range of values for the equivalence factor α. Finally, one of the SDP controllers is adapted for real-world use and tested on-road. The application of the SDP algorithm to control of a real vehicle is believed to be the most concerted and complete reported to date, as well as the first time that testing of a SDP controller on the open road is reported. For this purpose the controller response surface was transformed by replacing the acceleration state with engine load. This was deemed necessary for safety reasons and because of road gradients in the real world, and this approach is believed novel.

Parts of this chapter have been published in the following separate works:

C. Vagg, C. J. Brace, S. Akehurst and L. Ash, "Model-based Optimal Control of a Hybrid Electric Vehicle Using Stochastic Dynamic Programming." *In: 6th Conference on Simulation and Testing for Automotive Electronics,* 2014-05-15 - 2014-05-16, Berlin, Germany.

C. Vagg, C. J. Brace, S. Akehurst and L. Ash, "Minimizing Battery Stress during Hybrid Electric Vehicle Control Design." *In: 9th IEEE Vehicle Power and Propulsion Conference,* 2013-10-15 - 2013-10-18, Beijing, China.

In order to evaluate the potential of the control strategies to trade-off fuel savings with electrical system stress several variations were built and tested. The cost function on which the SDP control strategies were based was proposed in Section 5.1, and is restated for reference below.

$$c = f + \alpha \cdot C^2 \qquad\qquad (5\text{-}2)$$

The effect of increasing α is to impose a cost on high power use of the electrical system, introducing a preferential bias towards control resulting in more sustained low power operation wherever this detracts least from possible fuel consumption savings. Several control strategies were developed with a range of α; their performance was first evaluated using the vehicle powertrain simulation developed in Chapter 4 before implementing them in hardware and testing on a chassis dynamometer. Results of this testing are presented in the following sections.

6.1 Simulation Results

A range of SDP and ECMS controllers were tested in simulation with different values of α in order to investigate the parameter's effect on fuel consumption and electric powertrain use. It was expected that higher values of α should deter use of the electric powertrain, particularly at high power levels, and that this would have a negative impact on fuel consumption; however it is the shape of the resulting trade-off which is of interest.

Simulations were conducted over a LA92 drive cycle using the bespoke gear shift schedule already described. This drive cycle was selected based on the work presented in Chapter 3 with the intention of testing the hybrid vehicle over a drive cycle representative of its real-world usage. The vehicle simulation used was that presented in Chapter 4, which contains a model of the control strategy and is capable of simulating its response to each vehicle state. Given the drive cycle as an input this simulation is therefore capable of calculating the total fuel consumption of the HEV.

In the case of the SDP controllers charge sustenance over the drive cycle was adjusted by adding an offset to the SOC reported to the controller, as described in

Section 5.2.3, thereby shifting its natural charge-discharge ratio (Q_I). Since the SDP controllers were generally found to be charge depleting the offset applied was in the range -9.5% to -5% SOC. In the case of the ECMS controllers charge sustenance was achieved by optimisation of γ, as described in Section 5.3. All simulations were charge sustaining to within ±0.5% SOC. Parameters used for each control strategy including $α$, γ, and SOC offset are presented in Table 6-1, along with results of the simulations.

As previously discussed the square of the battery C-rate is proposed as a metric of the aggressivity at which the electric powertrain is being operated at, and the stresses to which components are exposed. Figure 6-1 presents the trade-off between fuel consumption and mean square battery C-rate using the data presented in Table 6-1, and the non-linear nature of the relationship suggests that α may be used to reduce aggressive use of the electric powertrain without sacrificing a proportional amount of the potential fuel consumption savings. Using the SDP results as an example, increasing $α$ from 0 to 0.002 g·h² reduces the mean square C-rate from 6.62 h⁻² to 3.24 h⁻², a reduction of 51%; meanwhile the sacrifice of potential fuel savings is 20%.

Table 6-1: Simulation parameters and results for SDP and ECMS controllers over a LA92 drive cycle.

| α (g·h²)	SOC offset (%)	γ (g·h)	ΔSOC (%)	Fuel Cons. (g)	$\overline{|C|}$ (h⁻¹)	$\overline{C^2}$ (h⁻²)	Fuel Saving (%)
Hybrid off							
				1158.63	0	0	
SDP							
0.006	-9.5		0.01	1144.61	0.52	0.71	1.21
0.004	-8.5		-0.09	1139.62	0.75	1.42	1.64
0.003	-8.5		0.13	1137.18	0.87	1.96	1.85
0.002	-7.0		-0.40	1132.00	1.10	3.24	2.30
0.001	-6.5		-0.03	1126.25	1.45	5.94	2.80
0	-5.0		0.13	1125.38	1.60	6.62	2.87
ECMS							
0.006		0.0036	0.29	1141.70	0.65	1.07	1.46
0.004		0.0036	0.02	1138.14	0.82	1.62	1.77
0.003		0.0035	0.02	1133.12	1.04	2.77	2.20
0.002		0.0035	0.00	1127.69	1.33	4.75	2.67
0.0015		0.0035	0.05	1124.87	1.50	6.22	2.91
0.0010		0.0036	0.18	1124.77	1.55	6.43	2.92
0.0005		0.0037	0.46	1123.72	1.69	7.02	3.01
0.00025		0.0038	0.04	1124.02	1.83	7.84	2.99
0		0.0038	0.40	1124.60	2.14	9.77	2.94
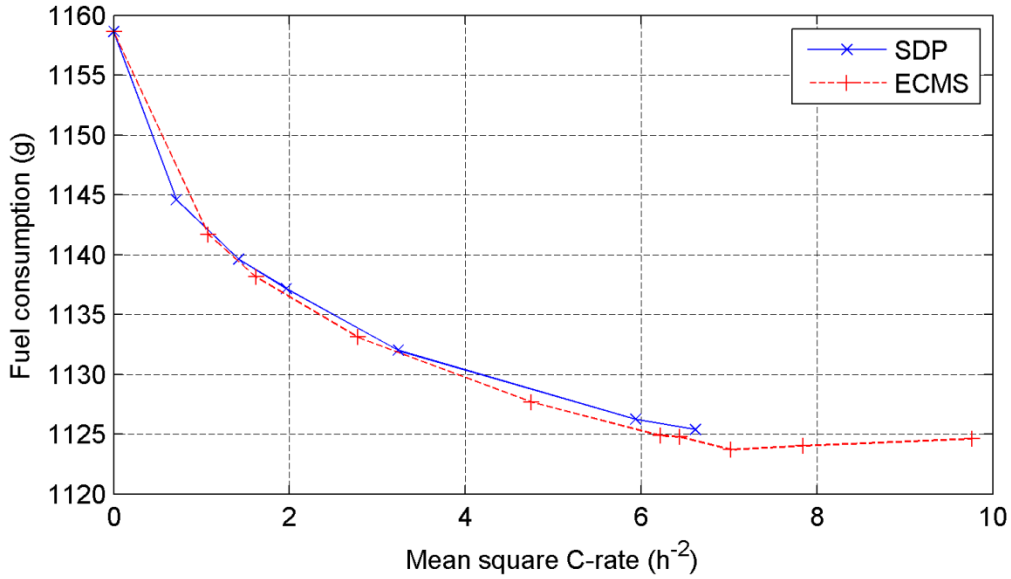
Figure 6-1: Trade-off between mean square C-rate, $\overline{C^2}$, and the fuel consumption achieved in simulation. The non-linear relationship suggests $\overline{C^2}$ can be reduced considerably without sacrificing fuel savings in the same proportion.

Alongside the trade-off with mean square C-rate it is perhaps helpful to consider the equivalent trade-off with the mean *not*-squared C-rate (Figure 6-2). This is enlightening and reassuring because the linear trend with relatively constant gradient for the majority of the range suggests that the fuel benefit per unit battery throughput is roughly constant. In other words, for each kilojoule of energy deployed through
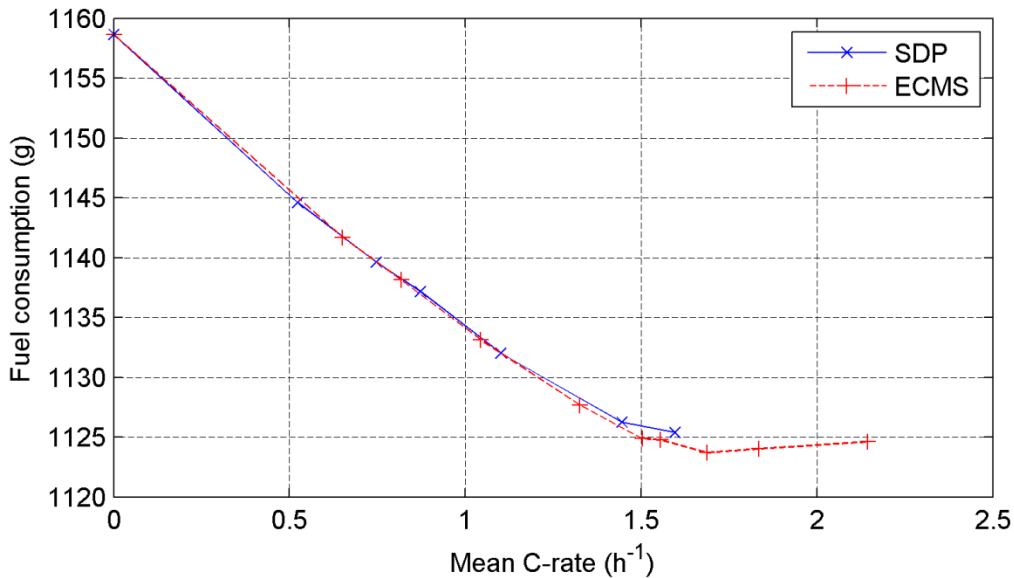

Figure 6-2: Trade-off between mean C-rate, $\overline{C}$, and the fuel consumption achieved in simulation. The generally linear relationship confirms that battery energy throughput is proportional to the fuel savings achieved.

120

the electric powertrain a proportional amount of fuel is saved. This observation confirms that applying cost to the square of the battery C-rate rather than the C-rate directly does not achieve any mysterious improvement in the equivalence between electrical and fuel energy, it simply ensures that when use of the electric powertrain is scaled back it is the highest power events which are sacrificed first, in favour of more consistent low power operation. Although the trend is linear for the majority of the range this is not true at the highest C-rates, which suggests that at this extreme the equivalence between electrical and fuel energy becomes less favourable.

Both the SDP and ECMS control strategies appear able to achieve similar fuel savings, and in doing so operate the electric powertrain at comparable levels of $\overline{C^2}$, therefore exerting similar levels of stress. It is significant however that the curve for ECMS extends considerably further into the high $\overline{C^2}$ region than that for SDP, suggesting that with small values of α the natural response of the ECMS is to apply additional and unnecessary stress to the electric powertrain through high power operation – in this case an additional 39% – while actually achieving sub-optimal fuel savings.

Figure 6-1 suggests the fuel savings achieved by ECMS over a LA92 are slightly better than those achieved using SDP. This is slightly surprising as the SDP algorithm is the more advanced of the two in theory, though its performance is only strictly guaranteed to be optimal when the probabilistic source data exactly matches the test cycle. In reality the savings are extremely similar and it is likely that the unavoidable inaccuracies inherent in modelling make it impossible to make any worthwhile claims on the significance of this observation.

6.2 Chassis Dynamometer Results

Following successful development and demonstration of both types of control strategy the next objective was to implement them as embedded code on production hardware, so that their real-time performance in an operational vehicle could be evaluated on a chassis dynamometer. This process required some preparation of the

control strategies which is described in the next section, followed by a description of the experimental test procedure and the results.

6.2.1 Adaptation for Real-World Use

The principal output of the SDP algorithm is a very large look-up table or control map, in this case four-dimensional, of which exemplary cross sections are shown in Figure 5-11 and Figure 6-4(a). For practical reasons some post-processing of this raw output was found to be necessary, firstly due to memory limitations of the hardware, but also because the control policy may contain occasional discontinuities for reasons explained in the following paragraphs.

For each vehicle state the SDP definition of the optimal control decision relies on the matrix dot-product $\mathbb{P}(X_1 = x_1 | x_0, u) \cdot J_\pi(x_1)$ which represents the probability-weighted future expected cost of decision u, where x_0 is the present vehicle state $\{v_0, a_0, g_0, soc_0\}$. If, however, a particular driving state $\{v_0, a_0, g_0\}$ was never visited during the source data then the state transition probability matrix $\mathbb{P}(X_1 = x_1 | x_0, u)$ is not defined, and in this implementation is simply null (all zeroes). In this case the cost on which the optimal control decision would be selected is exclusively the instantaneous cost. Since the instantaneous cost does not give any consideration to the value of stored electrical energy, though it may penalise high power use, the result is almost always to demand maximum assist during accelerations and cruises, and to demand no regeneration during decelerations, which is clearly illogical. Instead of accepting these decisions it was decided to override the stored control value with a null entry during the control policy refinement, and to re-evaluate the potential effect of the problem off-line.

The very nature of the poorly defined state probability matrix implies that this problem is not a significant one, since if a state was not visited at all during the extensive on-road data collection exercise then it is unlikely to pose any problem during normal operation, and for this reason no remedial action was felt necessary for testing in simulation. Nevertheless for real-world testing these sorts of issues must be taken more seriously, as sudden discontinuities in the control response may be cause for safety concerns. The nature of the problem can be seen in the high

speed, harsh deceleration quadrant of Figure 5-11, but is also illustrated in Figure 6-3.

In general the discontinuities noted in the control policy were in the negative acceleration domain, under braking. This is also the area which is of more concern from a safety viewpoint since under progressively heavy braking the controller's response may switch between zero and maximum regeneration, with the potential to confuse and destabilise the vehicle's ABS. Discontinuities in the positive acceleration domain were not regarded as a significant concern. In order to address the problem in the negative domain the controller's response to increasingly harsh deceleration was considered in every state. The value of this function at the largest deceleration before any unstable behaviour was encountered was then extended and applied universally to all decelerations beyond this. Figure 6-3 shows an example of the raw controller response where the region below -1.67 m/s^2 is considered unstable; the value of the controller response at -1.67 m/s^2 was therefore applied throughout the unstable region as indicated by the corrected response.

Memory allocation was the second significant concern in the embedded implementation of SDP, since the raw look-up table has a size of 48×29×6×31, and therefore a total of 258,912 elements, occupying over 500 kB of memory in single
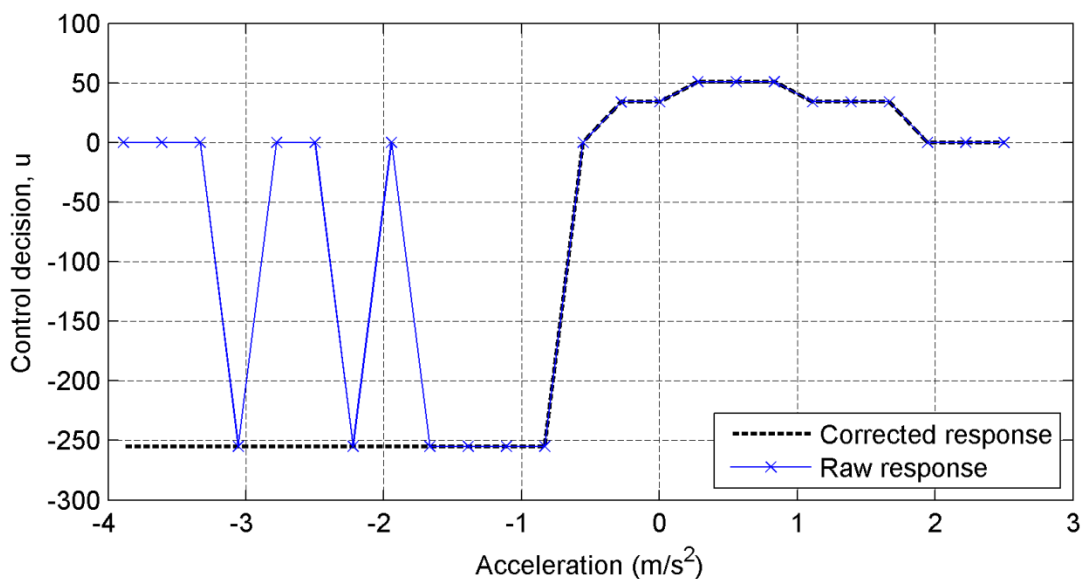

Figure 6-3: SDP controller response as a function of acceleration in 4th gear at v = 90 km/h, SOC = 67%. Undefined state transition probabilities can cause discontinuities in the raw controller response which must be identified and corrected for safety reasons.

precision. Since the target hardware had only 128 kB memory in total, which also needed to store all other functions associated with the hybrid system operation, some significant reduction was necessary. Two steps were taken to drastically reduce the size occupied by the look-up table in memory: a smaller data type was used, and the number of elements was reduced.

Since the vector of possible control decisions was only 31 elements long, comprising the values [-255:17:255], it was possible to translate the stored values into the indices [1:31] thus allowing the entire table to be stored as 8 bit integers. The indices were decoded real-time in the strategy so that the net output was unchanged, but the table occupied a quarter of the physical memory compared to a floating point equivalent.

To reduce of the number of elements in the look-up table a MATLAB script was written to cycle through and remove slices from the v, a, and SOC dimensions. During each execution of the loop this script removed in turn every possible slice from all dimensions of the table and noted the sum of the absolute error when the values in the missing slice were interpolated from surrounding values. The slice which minimized the sum of absolute errors was then permanently deleted from the table. This operation was repeated inside a loop until the map could be stored in less than 90 kB memory. During each execution of the loop the choice over which dimension to delete a slice from was entirely free, so each dimension was not necessarily reduced in size by the same proportion. In fact the a and SOC dimensions were reduced most, while the v dimension remained almost entirely intact. This process was found to be very effective in reducing the memory requirement of the look-up table without significantly diminishing its fidelity to the raw output of the SDP algorithm, as illustrated in Figure 6-4(b).

Other approaches to reducing the size of the look-up table for practical implementation have been suggested, for example Leroy et al [75] proposed using a neural network to model the controller surface. It is likely that this would work equally well and probably enable a far greater reduction in memory than was necessary here; however for this purpose the robustness of retaining calculated values was preferred.
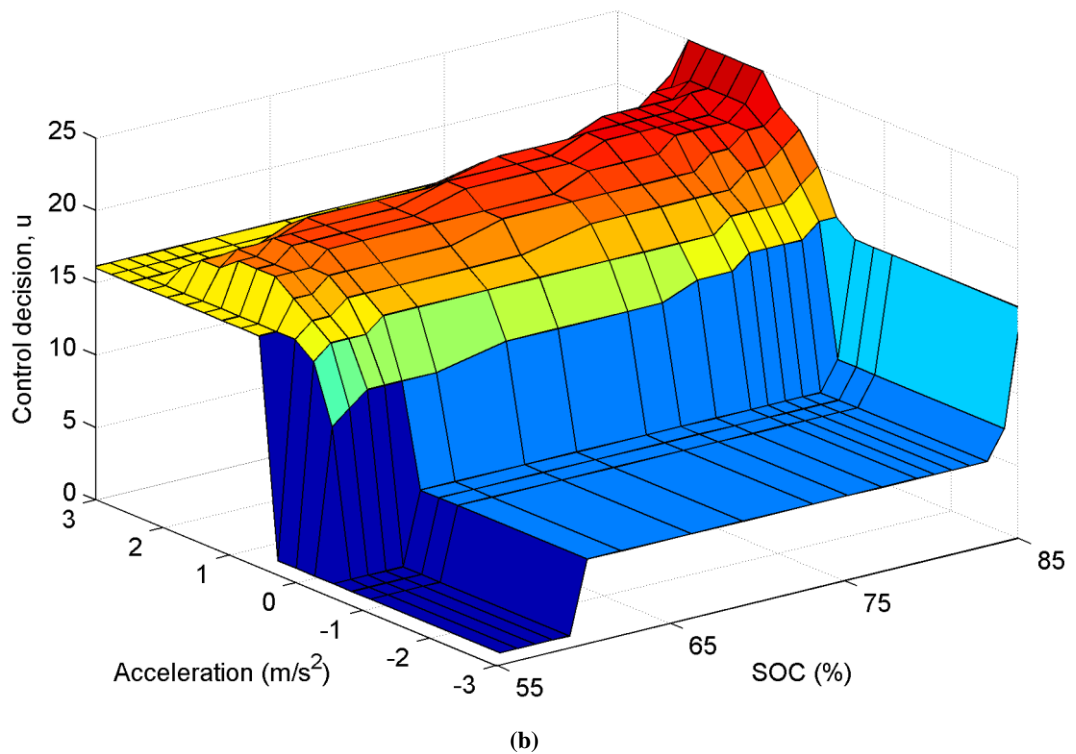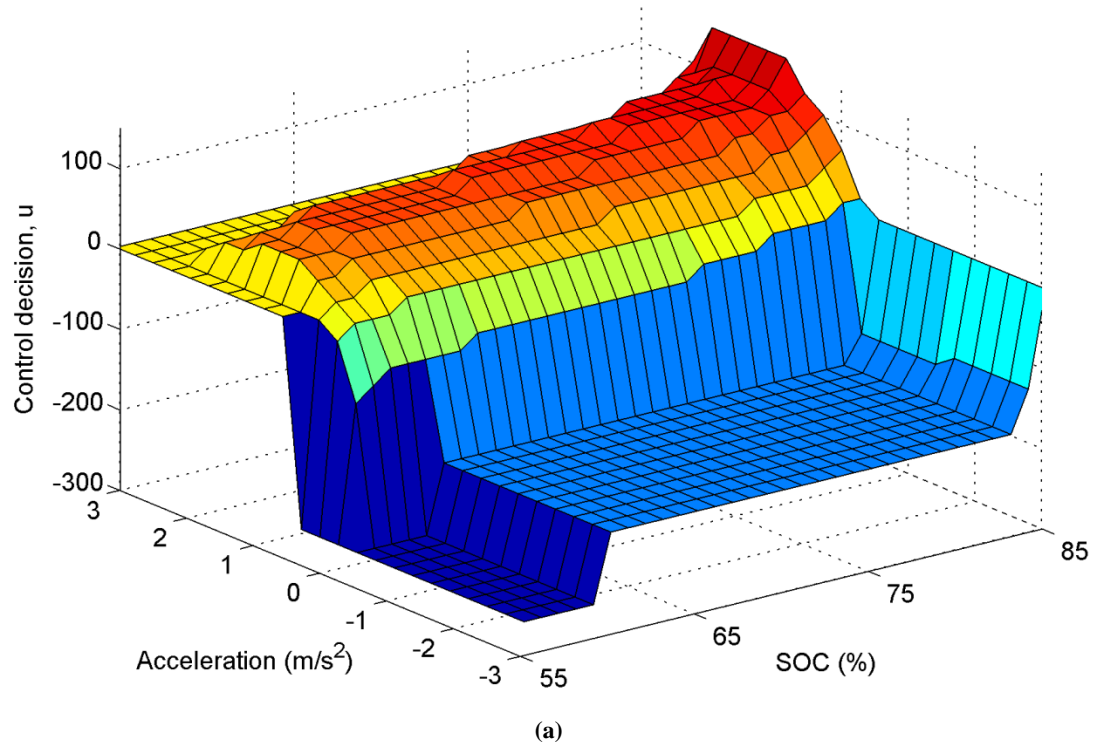
(a)

(b)

Figure 6-4: Control policy surface as a function of *a* and SOC at *v* = 32 km/h and in 4th gear. α = 0.001. In (a) the control decision *u* is a number in the range ±255 which defines the battery assist or regeneration current as a proportion of the maximum available at the present vehicle speed. In (b) the fineness of the state space has been reduced while maintaining the salient shape of the control surface, and the control values have been mapped to the range 1-31 to allow storage as an 8 bit integer in memory.

Several other tactics could also have been used to further reduce memory allocation while maintaining the calculated control values. It is worth noting, for example, that a single four-dimensional look-up table must take the form of a hyperrectangle in which every combination of states is defined. This means that the controller response in gear 1 is defined in the speed range 0–120 km/h, the vast majority of which is clearly unnecessary. Storing separate maps for each gear would allow the speed range for each to be more sensibly defined, thereby reducing waste. Another enhancement that could have been pursued is to make use of the probability data, with the aim of reducing the fidelity of the map more at infrequently visited states, and better maintaining its integrity in the most common regions. These approaches were considered but ultimately not thought necessary in the scope of this work as the process described delivered satisfactory results while maintaining simplicity and robustness.

6.2.2 Test Procedure

Chassis dynamometer testing was carried out at the University of Bath test facility, with the vehicle installation shown in Figure 6-5. The test cell is a temperature controlled environment, and all testing was carried out at 25°C. As previously described all quoted fuel consumption figures were recorded via the industry standard bag analysis method.

Preliminary testing found that over a NEDC the facility is capable of achieving fuel consumption measurement with a repeatability in the order of 1% Coefficient of Variance (CoV) [68]. To achieve the best repeatability it is essential that variables known to significantly affect fuel consumption are properly monitored, such as use of vehicle ancillaries (for example air conditioning and lights), SOC of the vehicle 12 V battery (which affects control of the alternator) and tyre pressures [77]. For this reason the use of ancillaries was constant throughout testing, tyre pressures monitored regularly, and the vehicle 12 V battery was left on trickle charge every night to ensure its SOC was well maintained.

In order to allow multiple tests per day all testing was performed with a hot engine; this introduces engine temperature as an additional and highly influential test condition which must be controlled. A consistent start of test condition for the engine

was achieved by first driving the vehicle at high speed until the engine coolant reached its target temperature of 90°C. As an additional step to the repeatability procedure used in preliminary testing the vehicle was then driven over a EUDC – the extra-urban section of the NEDC – which was followed immediately by the LA92 test, at which point logging of emissions data commenced. This procedure ensured no delay between the conditioning cycle and the test cycle, and as a result the repeatability of the test results is extremely good.

Possibly the greatest source of variability in chassis dynamometer testing is the human driver, and it is essential that the individual in this role is highly experienced in chassis dynamometer testing and able to follow the prescribed speed trace accurately and precisely. This task is far more difficult that might be imagined, even for somebody very experienced in road driving, because the surrounding environment gives no speed cue. In effect the task resembles a computer game much more than regular driving, where the driver's foot-eye coordination – a quality not often developed in daily life – is vital. Highly transient drive cycles such as the LA92 are considerably more difficult to drive accurately than modal drive cycles such as

Figure 6-5: Vehicle installation on the chassis dynamometer test facility. Road speed fan can be seen in the bottom left of the photo, and the driver's aid display in front of the windscreen.
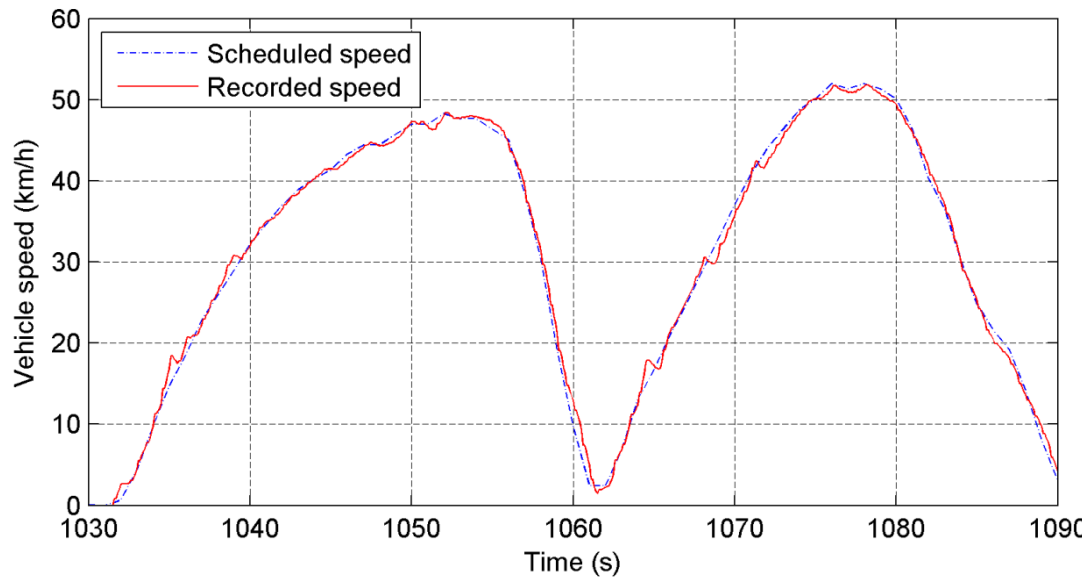
Figure 6-6: An example of the driver's ability to reproduce the LA92 scheduled speed with a good degree of fidelity. The greatest deviations from the trace are the result of gear shifting.

the NEDC. As well as the increased importance of anticipating the vehicle's responsiveness to changes in accelerator pedal position, the LA92 also does not provide the driver with convenient breaks in acceleration during which to change gears, which the NEDC does. The driver used throughout this work is very well accustomed to the task, and the degree of fidelity achieved in following the speed trace is shown in Figure 6-6; the greatest deviations from the scheduled speed are due to the driver compensating for the loss of tractive power during gear shifts, though even these are managed relatively well.

As an alternative to using a human driver the use of a robot driver to further improve repeatability was considered, though ultimately rejected because of the desire to maintain human-like control inputs. For the proper calibration of the hybrid controller it was felt important to develop and test the control strategies with input signals representative of those exhibited by normal drivers. Although robot drivers allow excellent trace following and repeatability the PID control loops they employ to achieve this, and their relatively aggressive handling of gear shifts, are not typical of human behaviour.

The bag analysis emissions measurement system used for this testing was equipped with 3 bags which filled sequentially, switching between bags at pre-programmed

times in the test. Each test is therefore effectively divided into 3 sections, and the total test fuel consumption is the sum of these. A secondary consequence of this is that the three sections may legitimately be regarded in isolation, and so an error in one section does not necessarily invalidate other(s); in fact for similar tests which both contain an error it may even be possible to construct a complete valid test by combining the two if the errors occur in different test sections. This was exploited on two occasions, and therefore 2 of the 23 tests reported are actually re-constructed. In one case this was due to errors in gear selection by the driver, whilst the other was a result of hardware faults causing one or more bags not to record a reading. In both cases the hybrid system behaved similarly in the two tests combined, and so the resulting re-construction did not contain any significant discontinuities in SOC, for example.

Accurate estimation of battery SOC is notoriously difficult, and especially so in the normal SOC working range because the voltage-SOC profile is very flat (Figure 4-10), and so a very small difference in the recorded voltage can dramatically change the inferred SOC. Furthermore the capacitive properties of cells means it is often necessary to leave them completely at rest for some time to allow the open circuit voltage to stabilise; the cells used in the battery pack took around 30 minutes rest to read a representative open circuit voltage. For repeatable performance of the hybrid controller between tests it was essential that the reported SOC at the start of each test be consistent, however for the reasons described it was found to be impractical to try and ensure an identical start of test SOC in absolute terms. For this reason the BMS was allowed to re-estimate the battery SOC between each test, and if necessary the battery was charged or discharged slightly during the high speed pre-conditioning drive in preparation for the subsequent test. The cell voltages were examined to ensure these were indicative of approximately 70% SOC and, once satisfied, the SOC stored in the BMS memory was manually overwritten as 70% to ensure consistency between tests.

Generally the strategies were very good at sustaining the battery SOC over a test, and of all of the reported tests the maximum change in SOC was 73 kJ – about 3%. In the case of the SDP controllers this was achieved without any offsets applied to the SOC reported to the strategy. For ECMS controllers some fine tuning of γ values was

necessary as the values predicted by the simulation tended to result in charge gaining behaviour. Each of the ECMS controllers was actually tested 3 times, but since the first tests were 8% and 11% charge gaining they are not reported. Lowering γ slightly ensured that charge sustenance was improved for the subsequent two tests, in both controller configurations, the results of which are reported. Since the SOC difference over a test, ΔSOC, was relatively little no correction of the fuel consumption values was carried out. In fact even had this have been deemed desirable it would not have been straightforward as no clear trend between fuel consumption and ΔSOC was observed; this suggests that the overall fuel consumption values are more affected by factors such as driver repeatability than by a 3% ΔSOC over a test.

6.2.3 Effect of Equivalence Factor, α

A total of 5 different control strategies were tested, comprising 3 variations of the SDP controller with different values of α, and 2 variations of the ECMS controller. For overall evaluation of the performance of the retrofit hybrid system a baseline condition was also tested with the hybrid system electrically deactivated. In each case a minimum of three repeats were conducted to ensure the controller was behaving repeatability and that the tests' repeatability was within the expected bounds. Average results of each configuration are presented in Table 6-2 and the trade-off between fuel consumption and $\overline{C^2}$ plotted in Figure 6-7 which may be compared with the corresponding model predictions previously presented in Figure 6-1.

The trend observed in the results of the SDP controller is similar to that predicted by the simulation, suggesting that there is a non-linear relationship which allows the battery mean square C-rate to be reduced without compromising fuel savings. For the ECMS results the trend does not align so well with the simulation, though in fact this amounts to one data point (at 5.96 h^{-2}, 1173.2 g with α = 0.0015) not returning the expected result; the other ECMS result (with α = 0) would fit with the predicted trend. It is therefore perhaps not appropriate to suggest too much from this observation, particularly since the number of repeat tests for the ECMS configurations were less than for SDP configurations.

130

For both SDP and ECMS control strategies increasing α is clearly effective in reducing the mean square C-rate, which does not necessarily incur a sacrifice of fuel savings in the same measure. Considering the SDP set of strategies tested it seems that the introduction of a relatively small α value yields highly favourable results, with the data showing that with $\alpha = 0.001$ the mean square battery C-rate is reduced by 13% compared to $\alpha = 0$, without any sacrifice of fuel savings.

Table 6-2: Results of chassis dynamometer testing showing the number of repeats, fuel consumption, Coefficient of Variance of the fuel consumption, change in SOC, and the mean square battery C-rate for a range of SDP and ECMS controllers.

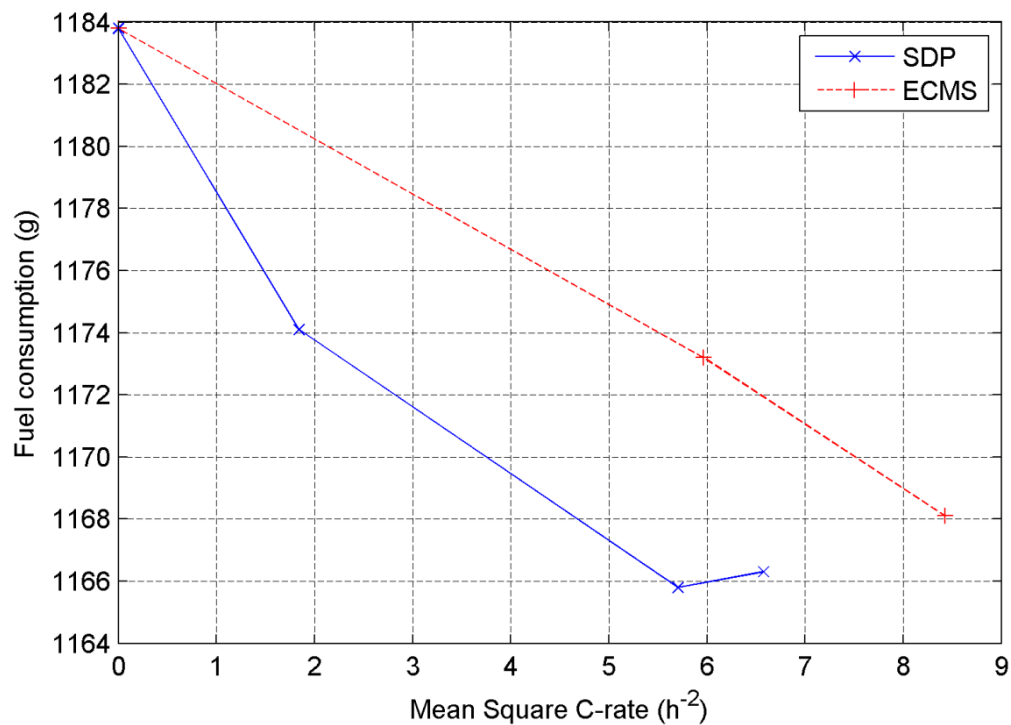| α	# repeats	ΔSOC (% / kJ)	Fuel Cons. (g)	CoV (%)	$\overline{C^2}$ (h^{-2})	Fuel Saving (%)
Hybrid off						
	6		1183.8	0.24		
SDP						
0.003	3	-0.5 / 11.2	1174.1	0.21	1.84	0.82
0.001	5	+2.0 / 43.8	1165.8	0.47	5.70	1.52
0	5	+1.9 / 43.2	1166.3	0.53	6.58	1.48
ECMS						
0.0015	2	+1.4 / 32.3	1173.2	0.20	5.96	0.89
0	2	+0.2 / 4.2	1168.1	0.87	8.43	1.32

Figure 6-7: Dynamometer test results showing the trade-off between mean square C-rate and fuel consumption. The trend for SDP controllers is similar to that predicted in simulation, while the performance of ECMS seems inferior.

6.2.4 Instantaneous Controller Behaviour

Examining the effect of α in more detail the instantaneous controller response for a set of SDP controllers can be compared. For each SDP controller variant the test with fuel consumption closest to the mean for that condition is selected as a representative, and the timeseries of SOC and DC current are presented in Figure 6-8 and Figure 6-9 respectively. The three tests shown have α values of 0.003, 0.001, and 0, fuel consumption of 1173.8, 1163.9, and 1163.9 g, and mean square currents of 1.84, 5.73, and 6.44 h^{-2} respectively.

Figure 6-8 provides an overview of how the control strategies manage the electric powertrain, allowing insight into at what point during the test most energy is deployed in assisting the diesel engine, and where the majority of energy is recovered. Perhaps unsurprisingly this analysis reveals that the three strategies follow similar approaches to the control, absorbing and discharging energy in the same places; there is no great dichotomy or bifurcation in what the overall approach looks like, but the tests with lower α exaggerate the SOC trajectory during the test as a result of their more aggressive use of the electric powertrain.
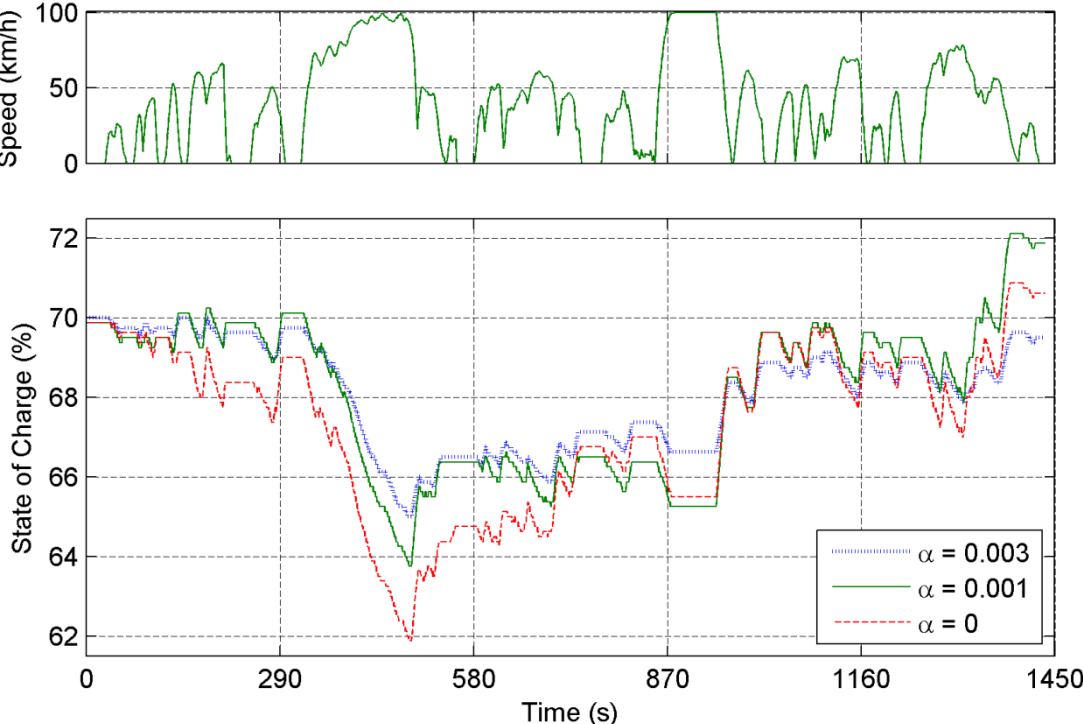

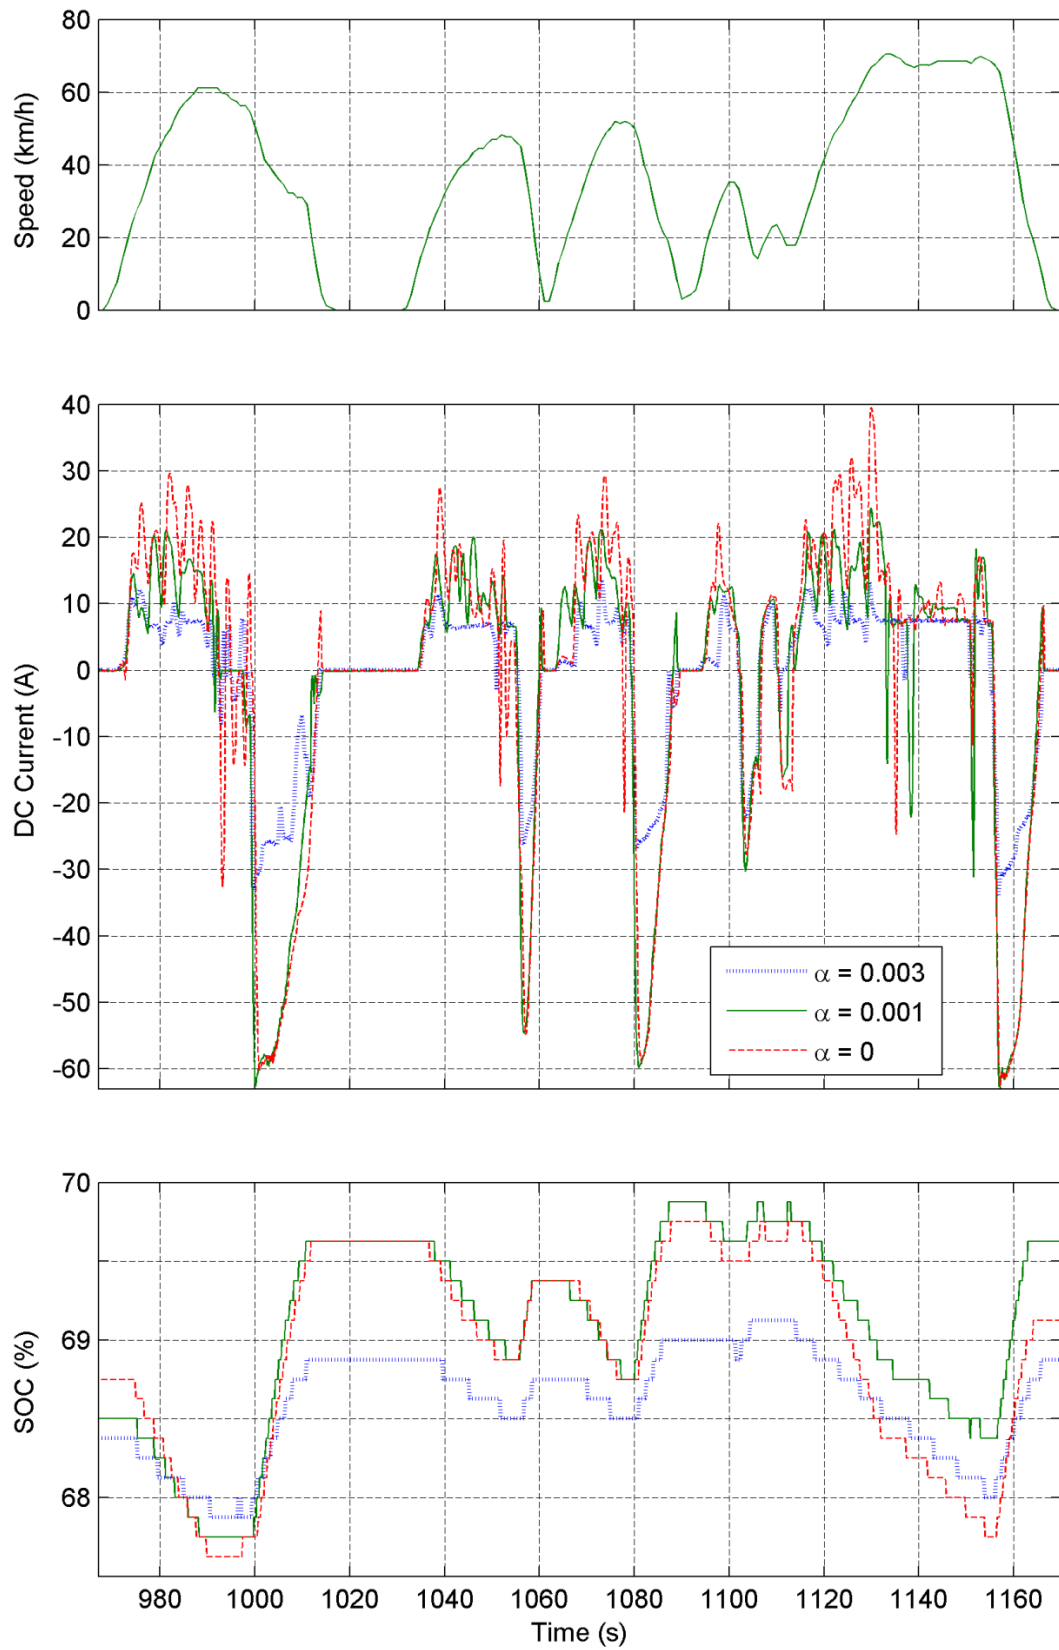**Figure 6-8: SOC trajectories over a LA92 for SDP controllers with different values of α.**

Figure 6-9: Instantaneous response of SDP controllers with different values of α over an excerpt of the LA92.

The conclusions drawn from considering Figure 6-8 are further supported by Figure 6-9, which shows that the control strategies generally cause assist and regeneration in the same places, though as the value of α increases the high current assist events are scaled back towards 10 A, and regeneration events are similarly less aggressive. The tests with $\alpha = 0$ and 0.001 recorded identical fuel consumption, though the test with $\alpha = 0.001$ was more charge gaining and had lower mean square C-rate. The final observation of note from this figure is that once again the extremely transient nature of the LA92 is apparent, with continual change of acceleration and frequent gear changes much akin to real world driving, but very different to the NEDC. As a result of this the assist events in particular are short and highly irregular.

6.3 Road Testing Results

6.3.1 Adaptation for Road Testing

Having successfully demonstrated the feasibility of running the SDP control strategy in real-time in a controlled environment on a chassis dynamometer the final step was to apply this to an on-road test. For all purposes thus far, both in simulation and during dynamometer testing, zero road gradient has been assumed; this assumption is important as it has been possible to directly infer both the driver's intention and the tractive force developed by the engine from the vehicle acceleration. Inference of tractive force from the engine is significant because this is used to determine instantaneous fuel consumption, and therefore in the calculation of the cost function.

Moving towards a production-ready strategy which may be operated on-road the unavoidable existence of road gradients will at times mean that there is a disconnect between acceleration, engine load, and possibly even the driver's intentions, making use of acceleration as a core state variable in the controller less acceptable; in fact in the worst scenarios this could be dangerous. When descending a steep hill for example it may be possible for the vehicle to be accelerating even with the engine at zero load (engine braking). In this situation a control strategy based on acceleration would incorrectly infer that the optimal decision is to deploy electrical energy in an assist event, which is in fact the opposite of the driver's intention.

In order to safely test the SDP control strategy on-road it was felt necessary to completely depart from use of acceleration in the state space, instead replacing this with a measure from which engine tractive force and the driver's intention could be inferred more directly. Of the ECU signals available through the OBD port two were contenders to replace vehicle acceleration: accelerator pedal position, and engine load. Accelerator pedal position may be the most direct indicator of the driver's intention, however it tends to have a highly non-linear relationship to tractive force, often incorporating some hysteresis and which may differ depending on gear. In contrast, for diesel engine vehicles the OBD engine load parameter directly reports the present engine output torque as a proportion of the maximum torque available at that engine speed.

Having already obtained an engine map on the dynamometer which included tractive force and engine load parameters, it was relatively straightforward to establish a direct relationship between acceleration (at zero road gradient) and engine load, allowing a direct substitution in the state space. In this way the control policy was completely transformed so as to be load-based rather than acceleration-based. An
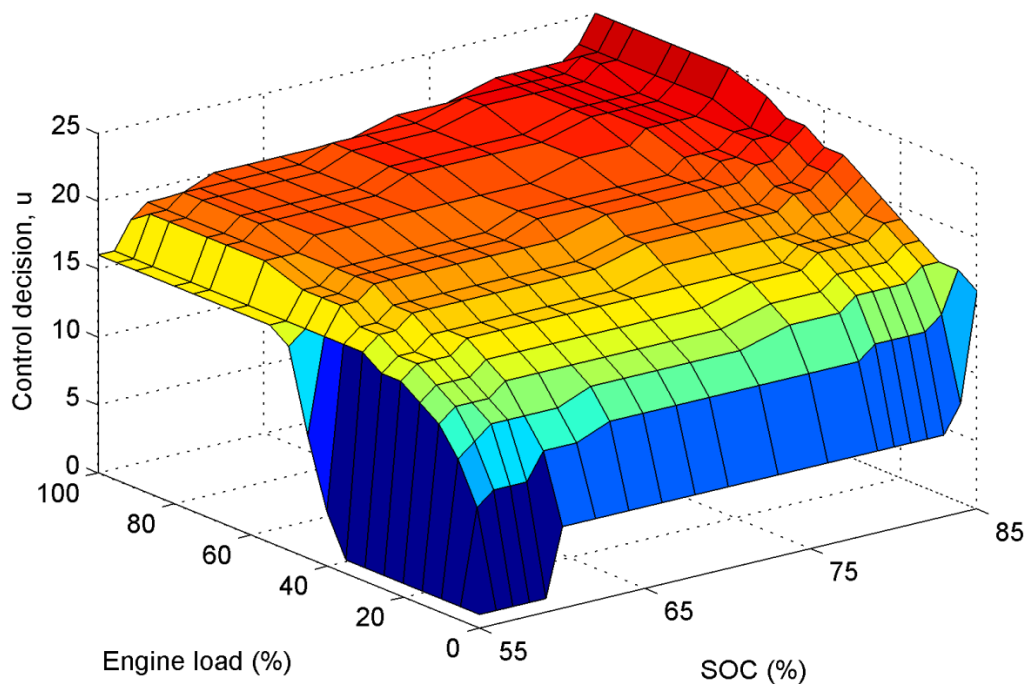

Figure 6-10: SDP control surface having been transformed in the state space, replacing the acceleration dimension with engine load. The cross-section is at v = 32 km/h and in 4th gear, with α = 0.001.

example of the resulting transformed surface is shown in Figure 6-10 which may be compared against the original surface in Figure 6-4.

6.3.2 Analysis of Road Testing

For the road test the control strategy with $\alpha = 0.001$ was used because in the dynamometer testing this achieved maximum fuel savings without over-stressing the electric powertrain. The vehicle was driven on public roads leaving Bath in the direction of Exeter, UK, following the A36, A39, and A368 roads. It is always important to reflect on the representativeness of a test in order to assess the results appropriately and in context. In this case it is particularly important to consider again the level of stochastic similarity between the route driven and the data for which the controller was optimised. To this end a comparison of SAFD plots is again employed, and the SAFD surface for this journey is shown in Figure 6-11. It is clearly evident from comparison between this and Figure 3-8 that, in comparison to the collective data from the 3 vans which were used in the design of the controller, the road test contains a greater proportion of time spent in the region of 50 km/h and less time at around 100 km/h. With this exception the probability distributions are
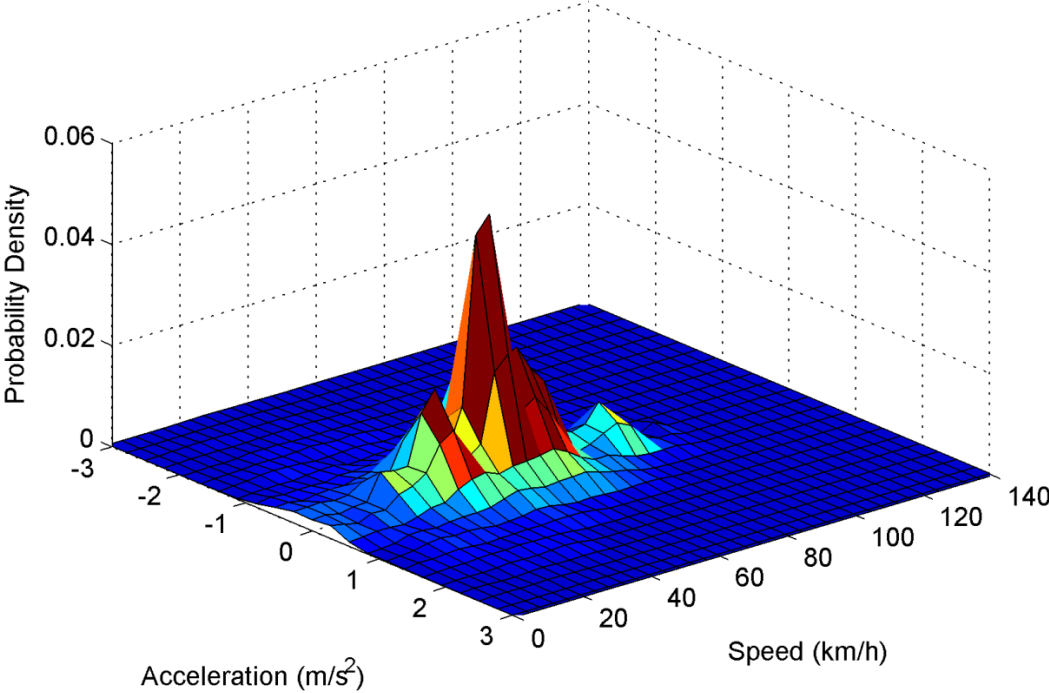

Figure 6-11: SAFD plot of the road test data. The surface resembles those of the LA92 and the data used to design the controller, but with higher probability density around 50 km/h and less around 100 km/h.

136

not dissimilar. The agreement between the road test data and the data on which the controller was based is 67%. This suggests that the road test is broadly representative of the conditions for which the controller was designed, while introducing a healthy degree of discrepancy (which must be expected in the real world) allowing some comment on the robustness of the controller to be made.

Interestingly the level of agreement between the SAFD of the historic data and that of this road test (67%) is similar to the level noted between the historic data set and the LA92 (72%), though the differences are realised in different places. It may therefore be said that although the agreement between the LA92 and road test data is not so strong (60%) both are similarly valid approximations of the original historic data set.

Traces of vehicle speed and battery SOC throughout the road test are shown in Figure 6-12. The transformed control strategy based on engine load appears to operate robustly and to manage battery state of charge easily, maintaining a margin to the maximum and minimum SOC limits. The initial SOC is slightly low, and this is gradually recovered by issuing assist events of reduced magnitude as per the controller's optimality rules. As the SOC increases the assist events become larger, and as a result the SOC approaches and then fluctuates around a nominal SOC. The probabilistic nominal SOC predicted for the acceleration-based version of this
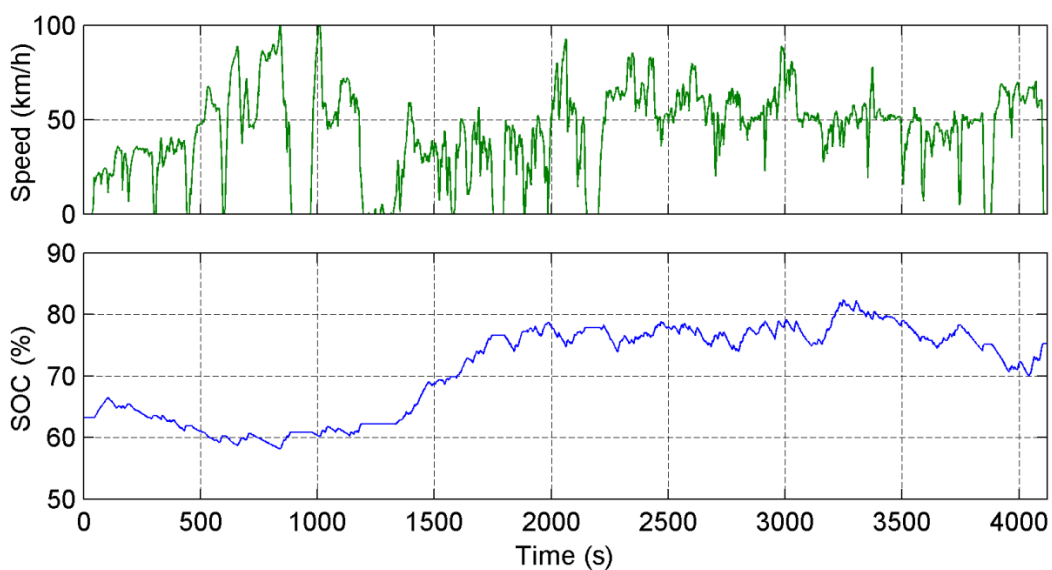

Figure 6-12: SOC trajectory of the SDP controller during a road test between Bath and Exeter, UK.

control strategy was 73% when simulating the LA92 cycle using the vehicle model developed in Chapter 4; behaviour on-road appears broadly consistent with this despite the controller having been translated into a different state space and any deviation from the raw SDP policy that might have be expected as a result of this process. This demonstrates some inherent resilience of the SDP control algorithm to post-processing and external influences, and well as to drive cycles which do not perfectly match the historic data set.

6.4 Chapter Conclusions

In this chapter the HEV controllers developed in Chapter 5 were tested, first in simulation, then under controlled conditions on a chassis dynamometer. Following successful dynamometer testing one of the SDP controllers was adapted for real-world use by replacing acceleration with engine load in the controller state space, and the controller tested on-road.

The trends observed through simulation were also observed in hardware testing. The results show that the trade-off between achievable fuel savings and mean square battery C-rate is indeed an interesting one, being highly non-linear and suggesting that the aggressivity with which the electric powertrain is used may be reduced considerably without significant forfeit of fuel savings. This observation may have profound impacts for system thermal management, as well as battery management and longevity.

Further analyses of the controllers' performances, accuracy of simulations, and potential for interaction between the hybrid system and driver assistance system are presented in Chapter 7.

Chapter 7   Discussion

Testing of both the driver assistance system and the hybrid controllers revealed some interesting and unanticipated effects. For the driver assistance system this included reduced rates of deceleration in addition to the expected reduced rates of acceleration; the hybrid testing exposed some of the weaknesses of relying on the backwards modelling approach. This section includes more advanced investigation into these observations and examines the potential interaction of the two systems when operating simultaneously.

Parts of this chapter have been published in the following separate work:

C. Vagg, C. J. Brace, S. Akehurst and L. Ash, "Model-based Optimal Control of a Hybrid Electric Vehicle Using Stochastic Dynamic Programming." *In: 6th Conference on Simulation and Testing for Automotive Electronics,* 2014-05-15 - 2014-05-16, Berlin, Germany.

7.1 Lightfoot

During development of the *Lightfoot* logic the majority of the development time and effort was invested in design of the IPS-driven logic. It is therefore interesting to note that subjectively speaking the driver is often more conscious of the GSI-driven logic. This is probably in part due to frequency of communication – for most drivers in the early stages after the system is activated the GSI beep will prompt almost every gear change, because most drivers habitually upshift at speeds higher that 2200 rpm. Nevertheless, it is also true that when the advice of the GSI is strictly observed it is much more difficult to drive so aggressively as to generate warnings or violations, because the tractive power of the engine is less at lower speeds and so rates of acceleration are also limited.

There is currently interest from legislating bodies in the ability of GSIs to deliver fuel consumption reductions in the real world [78], and how this can be fairly incorporated into legislative testing. In normal circumstances drivers may choose to ignore the advice of a GSI, or still more likely do not even register it. With this in mind, although it is not strictly the case, it is interesting to view the results as the effect of consistently following the advice of a GSI in the real world, and therefore the maximum possible savings achieved when drivers cannot ignore it.

An inherent assumption in the development of the IPS logic was that the threshold levels for triggering warnings should be independent of the vehicle which the system is fitted in and its maximum power. Effectively the IPS thresholds define an *acceptable limit of acceleration* which is an absolute value with no relation to the vehicle power or mass. In day-to-day life however, it is common to see vehicles accelerating at very different rates, with heavy haulage vehicles often accelerating much more slowly than small sports cars for example. It is therefore a matter of opinion as to whether it makes sense to define an absolute limit, or whether it might be sensible to tailor this to some extent based on vehicle power-to-weight ratio. The argument could be made, for example, that heavy vehicles consume more energy in accelerating and are bigger polluters per vehicle mile, and therefore should be subject to more stringent limits than small vehicles. Nevertheless it was felt that allowing all traffic to accelerate at the same rate makes a great deal of sense, and for the purposes of this work this carries the additional benefit that the functionality of the system is

not dependent on the vehicle specification, and so it may be installed in any vehicle without calibration.

For the commercial success of the *Lightfoot* system a key factor which has not been discussed is driver acceptance. During the trials there was no formal collection of feedback from the drivers involved through interviews or surveys for example, however during informal conversations drivers generally reported the system as being helpful and fair. Similarly the fleet managers of the companies involved did not report any dislike or problems with acceptance of the system by the drivers.

During the trials no consequences were attached to the number of violations a driver accumulated, though when implemented across a whole company fleet the expectation is that some policy would be adopted. This raises a further question and opportunity for research with regard to what form such a policy should take to maximise the effectiveness of the system, for example whether it should be based around a reward or punishment structure. Furthermore it is noted that the system's current implementation is largely based around providing criticism for poor driving, rather than constructive feedback. A re-evaluation of this along with a carefully considered management policy would mitigate the risk of any driver acceptance problems.

7.1.1 Effect on Recoverable Energy

It is well known that the amount of energy which may be recovered by a hybrid or electric vehicle may be heavily affected by the way in which it is driven [79]; for example sudden braking will mean that brake energy is released at high power levels which may exceed the power capacity of the energy regeneration system. Any energy released at powers above the limit of the KERS must be dissipated by the mechanical brakes as heat. As a result of the power limitations of the recovery system it is likely in most hybrid powertrains that lower rates of deceleration would result in a greater proportion of braking energy falling within the recoverable envelope, increasing the amount of energy which may be captured and consequently enabling the system to deliver greater fuel savings.

Although the IPS-driven logic only considered positive values of IPS (sometimes called Relative Positive Acceleration) it should be noted that the value is equally well defined in the negative domain, giving some indication of braking aggressiveness. Figure 7-1 compares the distribution of braking power during each phase of the trial, showing that during the Live phase the incidence of high power braking was reduced with respect to the Baseline phase, while the incidence of lower power decelerations was increased. This is an interesting finding because no part of the logic implemented was designed so as to encourage more gentle decelerations in any way. It seems that the act of accelerating more gently, and perhaps the slightly more relaxed mentality of the driver which results from this, inadvertently also caused a shift in braking behaviour. This finding has at least two practical consequences:

1. Reduced braking rates would suggest that drivers are better anticipating the road in front of them, and adopting an all-round smoother driving style. This is likely to have safety implications and an accompanying fall in accident rates might be anticipated, though long-term data collection would be required to substantiate this.

2. The proportion of braking energy which may be captured by the KERS may be increased, further enhancing the efficacy of the hybrid powertrain. In this way the interaction between the systems may mean that the fuel savings
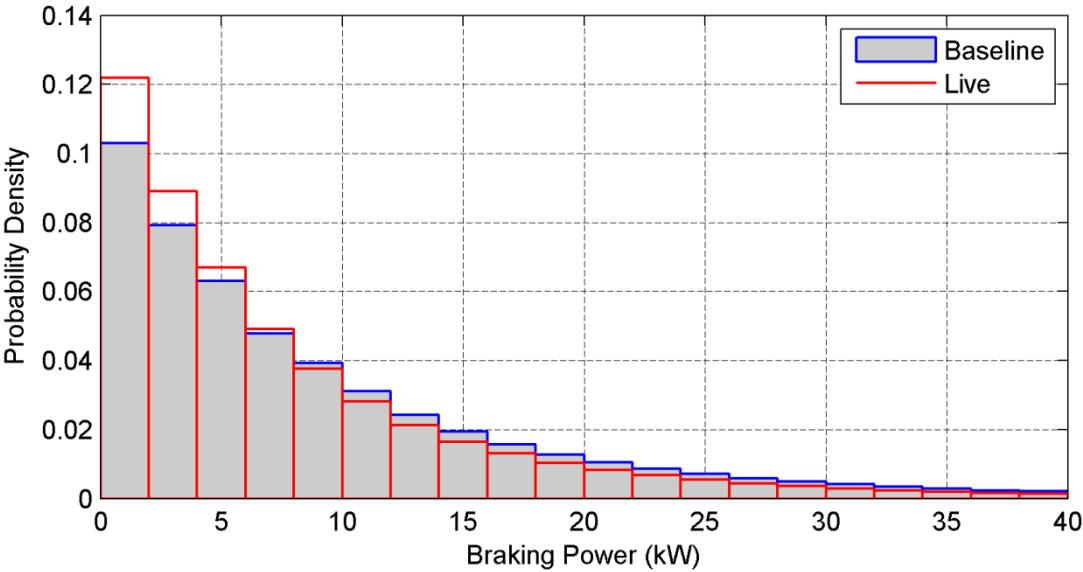

Figure 7-1: Distribution of braking power during decelerations. In the Live phase drivers were more likely to decelerate gently (reducing the braking power) even though this was not deliberately encouraged.

achievable with both operating together exceed the sum of their individual contributions.

With regard to the second observation, it may be that reduced rates of braking mean a greater proportion of deceleration energy becomes recoverable; however, although the probability distribution of deceleration power shifted in favour of lower powers the *absolute* quantity of deceleration energy released per kilometre actually fell from 380.6 kJ/km in the Baseline phase to 334.9 kJ/km in the Live phase of the trial as a result of the smoother driving style adopted. The question concerning absolute recoverable energy is therefore slightly different to that concerning the proportion of braking energy which may be recovered.

In order to investigate further the potential effect of *Lightfoot* on energy recoverable by the hybrid system the recoverable energy per kilometre was calculated as a function of the energy recovery power limit, as shown in Figure 7-2. The calculation of these data assumes that with an energy recovery system rated at 15 kW, for example, any deceleration energy released at powers less than or equal to 15 kW is
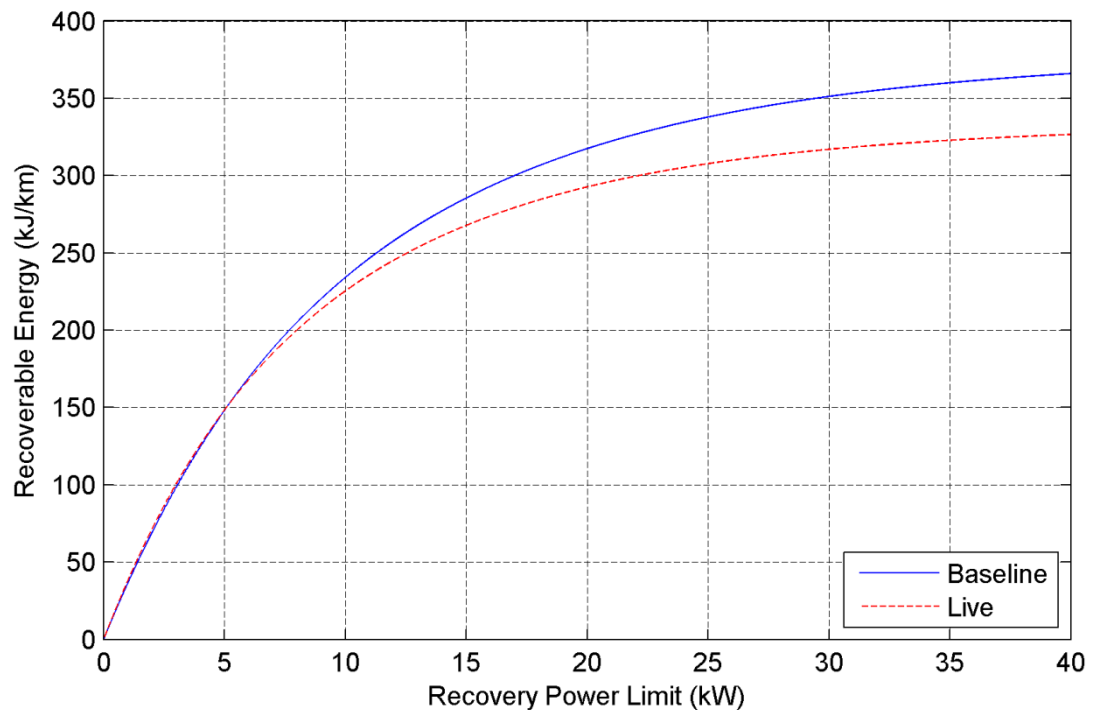

Figure 7-2: Braking energy recoverable per kilometre during each phase of the trial. During the Live phase less energy was available for recovery because of the smoother driving styles adopted; however at recovery rates below about 10 kW the overall reduction in energy availability was largely offset by the shift towards lower deceleration rates, meaning there was little net change.

recovered entirely, while for all powers above this threshold energy is recovered at the 15 kW limit, with mechanical brakes also employed to make up the remaining braking power required.

Examination of Figure 7-2 shows that overall there is less recoverable energy per kilometre with *Lightfoot* active, however it also reveals that at low rates of energy recovery (below approximately 10 kW) this is offset by the already described shift towards lower rates of deceleration, meaning there is little net change. It should be noted however that the situation is in fact far more complicated than this graph may suggest, because with no fuel being injected the engine alone may provide as much as 10 kW of braking power (engine braking), meaning that for the Live trial up to 67% of the energy which is in theory available for recovery could have been absorbed by the engine. In practice this is likely to be an overestimate because the engine is sometimes disconnected from the wheels during braking by means of the clutch. The situation is in fact even less clear-cut than this because the KERS actually recovers energy at the same time as engine braking and not only when the braking power exceeds engine friction. A side-effect of the hybrid system's operation would probably therefore be to introduce some bias towards faster rates of deceleration. This coupling of effects makes it very difficult to draw quantitative conclusions, however the simple observations drawn from Figure 7-2 offer worthwhile insight.

In summary, although it is extremely difficult to quantify the effects of interaction between *Lightfoot* and the hybrid system it appears probable that reduced rates of deceleration would result in a greater proportion of deceleration energy being captured by regenerative braking, though the absolute quantity of energy recovered per kilometre may be less because of the smoother driving style. It should be reiterated that these findings were observed even though no attempt to encourage gentle braking was made, and that if this were actively encouraged then more energy may be recoverable.

7.2 Hybrid Control

For several years SDP has been proposed in literature as a state-of-the-art solution to the optimal control problem, with researchers demonstrating in simulation how the approach could be applied to a vehicle. During the course of applying this to a real vehicle a great deal was learnt about how this should be approached; this is generalised into a procedure for robust implementation of the algorithm and presented in Section 7.2.1. Following this the performances of the SDP and ECMS controllers relative to one another are discussed.

7.2.1 Generalised SDP Implementation Process

Development of a robust SDP control strategy which may be implemented on a real vehicle and expected to perform reliably as well as optimally is a substantial undertaking; it requires collation of drive cycle information with data extracted from models of powertrain components, as well as considerable trial and error in the course of parameter selection. In order to illustrate this procedure a flowchart of the process is presented in Figure 7-3.

Both an accurate state transition model for the vehicle and a good representation of the vehicle's typical driving cycle are required in order to make SDP worthwhile. This information is combined in the SDP implementation, and values of the discount factor λ and the policy convergence threshold ε must then be determined by trial-and-error, ensuring that the probabilistic charge sustenance profile is acceptable, as indicated by the first feedback loop in Figure 7-3. It should also be noted that varying the equivalence factors in the cost function will alter the charge sustaining properties and so, though these effects are likely to be considerably less than from varying λ, it may be necessary to repeat the exercise for several values of these parameters.

As a result of the controller's performance in simulation, and possibly some real-world validation, it may prove necessary to apply a SOC offset to achieve the desired nominal SOC during operation. Finally, although not ideal, it may be necessary to transform the state space of the controller to one more suited to real-world implementation, as will be discussed in the following section; this is represented in the second feedback loop in Figure 7-3.
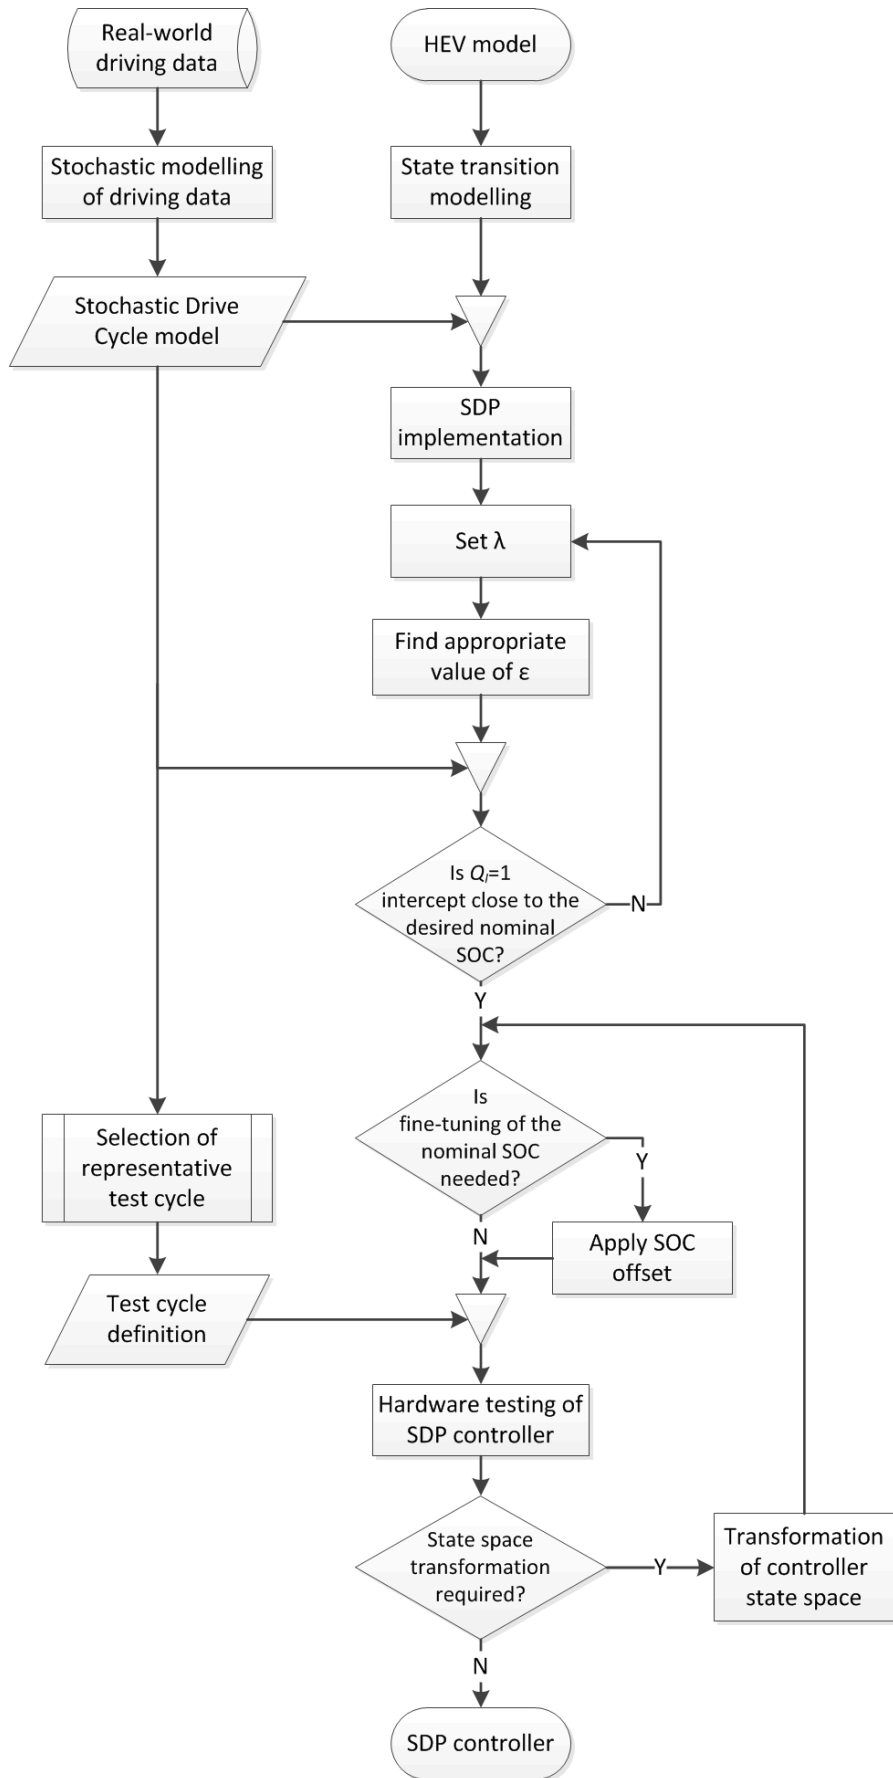
Figure 7-3: Generalised process for the robust design of SDP controllers.

The calculation of probabilistic charge sustenance profiles is dependent on the driving cycle probabilities, and in Section 5.2.5 some consideration was given to the effect of using driving cycle probabilities which differ from the design set. It would be interesting as an item of further work to carry out a more rigorous analysis of this to determine the effects on the control and probabilistic charge sustenance of a sudden change in duty cycle, perhaps as a result of a reallocation of a vehicle's usage within the company, or a change in driver. A natural progression of this would be to consider ways in which the controller could identify shifts in driver behaviour on-line and self-adapt to better suit these.

7.2.2 On-Road SDP Implementation

In this work a stochastic model of the typical drive cycle as a Markov process was used, with vehicle acceleration as one of the states. In the road-ready version of the control strategy the acceleration state was replaced with engine load. This transformation between states is not ideal and will undoubtedly have incurred some loss if information, and therefore of optimality. For example, there will have been situations where an acceleration-based controller would have assisted or regenerated, where the revised load-based controller did not. While this is precisely the reason for having made this change, it also means that the strictly optimal policy which was the output of the SDP algorithm has subsequently been *changed*, and must therefore no longer be optimal. Of course the reality of the situation is less absolute since the original policy was only optimal among those sharing the level of information available, which did not include road gradient for example; it was therefore already sub-optimal among controllers for which road gradient could have been available.

An alternative approach to the SDP problem formulation may have been to describe the stochastic drive cycle in terms of engine load from the outset, thereby absorbing road gradient and acceleration into a single state and avoiding the necessity for a fifth probabilistic state. This was not pursued here because the vehicles used to collect drive cycle data did not all have the same engine model and so the load signals would not have been comparable, nor was the fuel consumption map based on load available for these engines. Furthermore since the engine load signal can be

extremely volatile it is possible that this may cause unexpected problems in its stochastic modelling.

Nevertheless the behaviour of the controller when operating on-road seemed robust, with the net charge sustaining properties appearing undiminished. This implementation of SDP to a physical vehicle, the subsequent on-road application, and the lessons learned therein represent an incremental advance on the work of others. The better incorporation of road gradient in the stochastic model to produce a policy which is more globally optimal is an opportunity for future work.

7.2.3 Comparison of SDP and ECMS

Having successfully implemented a SDP control strategy on-vehicle it is clear that the approach has great potential for practical application. Nevertheless the behaviour of the ECMS controller is also of interest and since this work is one of very few in which SDP and ECMS controllers are implemented side-by-side the opportunity should be taken to make a direct comparison.

Regarding the two control strategies' relative performance it is clear from the holistic analysis that SDP and ECMS are able to achieve similar levels of fuel saving, but that the SDP controller is able to do so while exerting less stress on the powertrain. Although differences exist between the simulated and measured results this overall trend is apparent in both, and is consistent with the finding of Liu and Peng [31] who observed through simulation that the ECMS can tend to oscillate between heavy assist and regeneration. This behaviour is a result of the characteristics of reciprocating internal combustion engines, which tend to be most efficient when operating under heavy load where the volumetric efficiency is highest (both for spark ignition and compression ignition engines). For this reason *both* strategies ought to have a similar tendency to regenerate electricity heavily in an effort to move the engine's operating point into the high load region, and to later use that energy to reduce the required power output of the engine. Both strategies are aware of the engine efficiency map; the difference between the strategies' approaches is in fact solely due to the value they assign to stored energy in the context of the driving mission, as is highlighted by Equations (7-1) and (7-2), in which the strategies

ordinary notation have been combined so that their cost functions can be compared analogously.

$$J_{SDP}(x_k) = f(x_k, u) + \lambda \sum_{x_{k+1} \in X} \mathbb{P}(X_{k+1} = x_{k+1} | x_k, u) \cdot J_\pi(x_{k+1}) \qquad (7\text{-}1)$$

$$J_{ECMS}(x_k) = f(x_k, u) + \gamma \cdot C \qquad (7\text{-}2)$$

For both strategies the optimal control response at every state is found according to

$$\pi^*(x_k) = \arg \min_u \{ J(x_k) \} . \qquad (7\text{-}3)$$

In the case presented it is assumed that there is no cost on powertrain stress ($\alpha = 0$), and the equations have been laid out so as to facilitate direct parallels being drawn. The first term is common between both cost functions and represents the instantaneous fuel consumption, which will always be reduced by assistance from the EM and increased by regeneration; however the second terms, which offset the instantaneous effects by implying some worth of the electrical energy, differ considerably. In the case of ECMS the price of electrical energy is assumed constant and defined by the equivalence factor γ, whilst in the case of SDP a far more sophisticated approach is taken which incorporates the likely availability of electrical energy in the future, and whether the stored energy might be better spent at other opportunities. The global result of this variable electricity price is that the SDP strategy is better able to target the use of stored energy to minimise cost incurred over a drive cycle, including battery stress where this is included in the cost function.

Quantifying the difference in performance between the SDP and ECMS is an important contribution of this work, and it certainly seems from the test results that SDP has substantial advantages over ECMS. For example, by directly comparing the $\alpha = 0$ cases with reference to Table 6-2 and Figure 6-7 we may conclude that the ECMS tests had a mean square C-rate 28% higher than the SDP, while achieving 10% *less* fuel consumption savings.

Although the difference in performance between the two controller groups is quite apparent when considering the overall results it is worth noting that differences between their instantaneous behaviour are not so easy to identify. Figure 7-4 presents
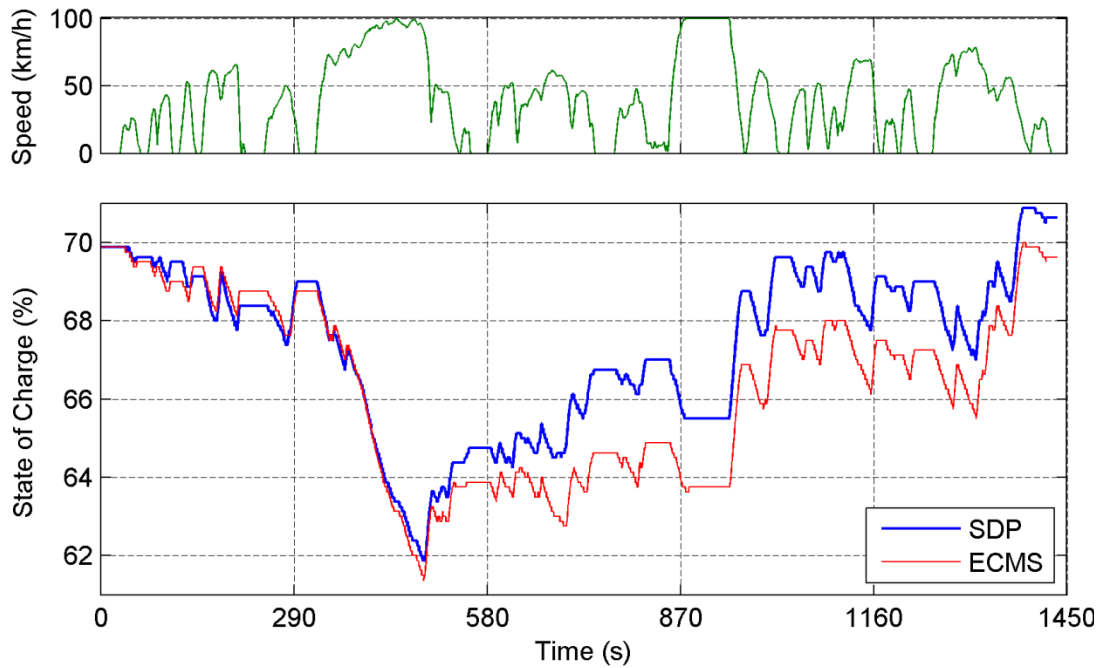
Figure 7-4: Comparison between SOC trajectories of one SDP and one ECMS test. The tests are generally similar, regenerating and deploying energy in the same places, they have similar mean square C-rates of 6.44 h^{-2} and 6.20 h^{-2} for the SDP and ECMS controllers respectively, but have very different fuel consumptions of 1163.9 g and 1174.9 g respectively.

SOC trajectories for two tests: one SDP with $\alpha = 0$, and the other ECMS with $\alpha = 0.0015$. Both tests have a similar mean square C-rates (6.44 h^{-2} and 6.20 h^{-2} respectively) while the fuel consumptions achieved differ substantially (1163.9 g and 1174.9 g respectively). The SOC trajectories of the two tests are extremely similar, and therefore the controllers are generating and consuming energy in roughly the same way.

The current profiles during middle portion of the test, where the greatest SOC difference is accumulated, are shown in Figure 7-5. Even during this section it is clear that the controllers are operating in broadly the same fashion, both tending to assist and regenerate in the same places, with the SDP controller showing only a slight propensity towards more energy recovery and less assist throughout.

The distinct lack of any great differences between controller responses suggests that there may be ways in which the ECMS could be slightly enhanced, without great effort, so as to achieve more SDP-like behaviour. Essentially the goal of such an exercise would be to replicate the variable electricity price used in SDP while avoiding much of its inherent complexity, therefore yielding a more straightforward,
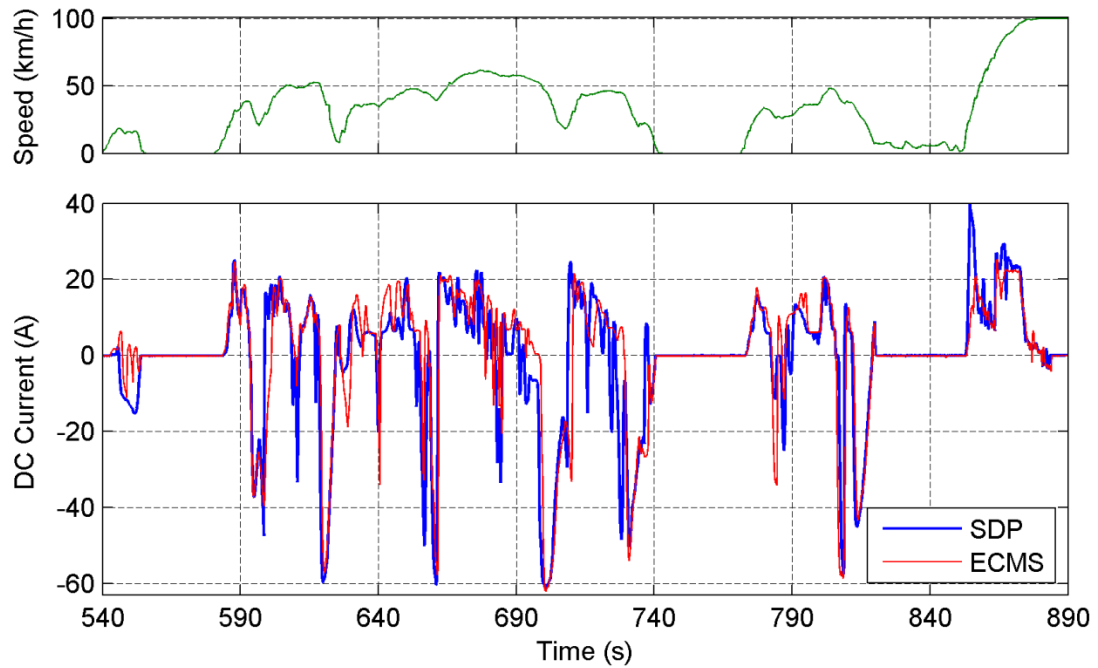
Figure 7-5: Comparison between battery current profiles of one SDP and one ECMS test. The tests are generally similar, with the SDP controller regenerating only slightly more around 550 s and 695 s.

practical and less cumbersome solution. One possible approach, for example, might be to assume a crude link between vehicle speed and the likelihood of urban, rural or motorway driving, and to adjust the price of electricity based on this; this approach could be viewed as some form of ECMS augmented with a Fuzzy Logic driving pattern recognition module such as that developed by Liaw and Dubarry [36].

Since it is difficult to ascertain any great differences between the current profiles of the two tests in consideration the reason for the difference in measured fuel consumption must be questioned. It is likely that when examining differences in this order of magnitude other confounding factors, such as driver repeatability, are important. Overall the trends observed using averaged results are as expected from simulation, and so the results are regarded as robust. Nevertheless it would certainly be a worthwhile exercise to repeat the investigation using an electrical powertrain with a larger power capacity.

7.2.4 Physical Effects of Mean Square C-rate

A major assumption in this work is that a reduction in the mean square C-rate would result in reduced heat generation in the motor and battery pack, which in turn has system level benefits such as reduced cooling requirements, reduced BMS workload, reduced battery degradation and/or reduced battery capacity requirement. Although the reasoning behind these hypothesised effects is regarded as well-grounded it is worthwhile examining the data for any sign of their manifestation. Figure 7-6 presents the mean temperature rise of the battery pack and motor for all tests as a function of mean square C-rate. It should be noted that no particular procedures were put in place to maximise the quality of these data, because to do so while also monitoring engine temperature and battery SOC would have complicated the testing too much; nevertheless the data show clear trends. In the case of the battery it is not unreasonable to conclude that the temperature rise is directly proportional to mean square C-rate. For the motor temperature a similar conclusion may be reached, though it should be noted that the trend line will not pass through the origin because electrical resistance is not the only source of heat generation; even with zero electrical activity some heat is generated by friction in the motor bearings and by magnetic cogging.

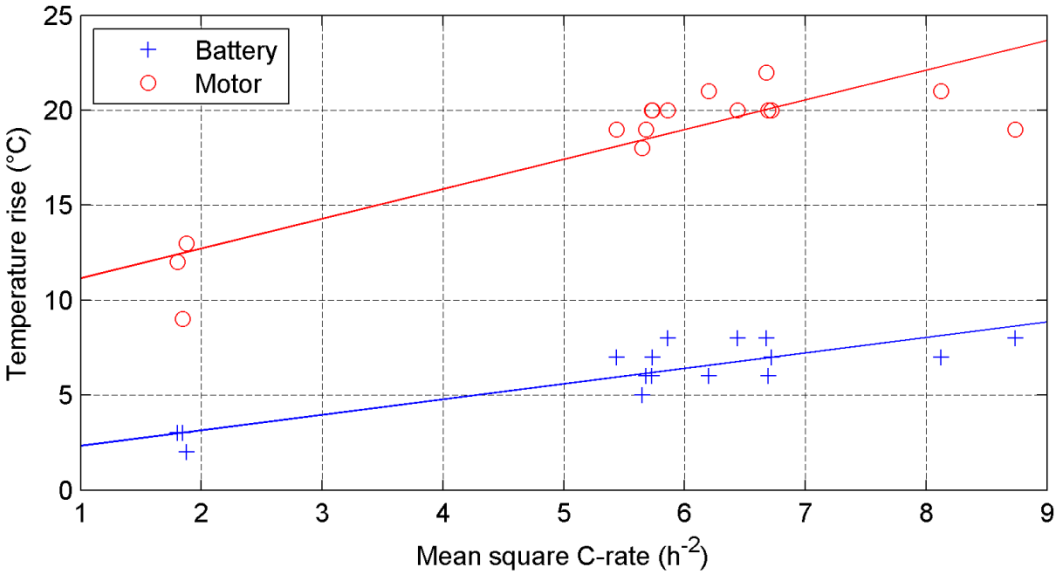As well as the bulk temperature rise of the battery it is also hypothesised that,

Figure 7-6: Temperature rise in the motor and battery pack during each LA92 test as a function of mean square C-rate, with lines of best fit.

perhaps over a longer drive cycle, a trend linking high mean square C-rates to larger cell temperature differentials could be established. This is an important metric because cell temperature differentials have the effect of creating differences between internal resistances, which in turn cause difficulties for the battery management system. Unfortunately it was not possible to observe these effects during the test program because temperatures of each cell in the battery pack were only measured with a resolution of one degree centigrade, and because for its protection the battery was operated relatively conservatively within known safe working limits.

7.2.5 Correlation between Simulation and Dyno Results

Although the absolute quantity of fuel consumed during a test and the general trends observed when varying the value of α were both predicted by the simulation to good levels of accuracy, comparison of the percentage fuel savings predicted and those achieved during testing reveals a disappointing correlation. For example, in simulation a maximum fuel saving of 2.87% was predicted with the retrofit hybrid system active and under SDP control; however during dynamometer testing the average fuel saving for the same controller was only 1.48%. This difference clearly warrants some investigation.

In order for the hybrid system to achieve any fuel saving it is assumed that the driver's target speed is achievable without the assistansce of the EM, and therefore that any assistance from the hybrid system would result in the driver activating the accelerator pedal less. As previously described the retrofit system actually increases the installed power of the vehicle, and so ensuring that this power is not simply exploited is part of the rationale for also having the driver behaviour tool installed alongside the hybrid system. Much of the discrepency observed between the simulated and measured fuel savings is likey to be because of this fundamental assumption being erroneous in the case of the LA92, since its high power demands meant that in several places the vehicle struggled to achieve the scheduled speed trace even with full activation of the accelerator pedal.

Figure 7-7 shows an excerpt of the acceleration from 0–100 km/h at 850 seconds into the LA92 test; the scheduled speed is shown with a heavy black line, alongside the speed achieved for each of the 6 baseline (hybrid off) tests. As can be seen, for

approximately 15 seconds the driver was applying full accelerator pedal demand, lifting off momentarily only to change gears, and so the accelerator pedal trace resembles an on-off demand. In addition to this the driver frequently started the acceleration slightly ahead of the scheduled speed in order to carry some speed in reserve, yet despite this consistently struggled to meet the prescribed test speed from about 860 seconds on. For this reason when the hybrid system was activated the additional power available was not in fact used exclusively to save fuel, but was at least in part used to supplement the engine power and therefore help the driver better achieve the speed trace.

Although difficult to quantify it is suspected that power deficit during hard accelerations, which invariably correspond to regions of assist, was a large contributor to the discrepency between simulated and measured performances.
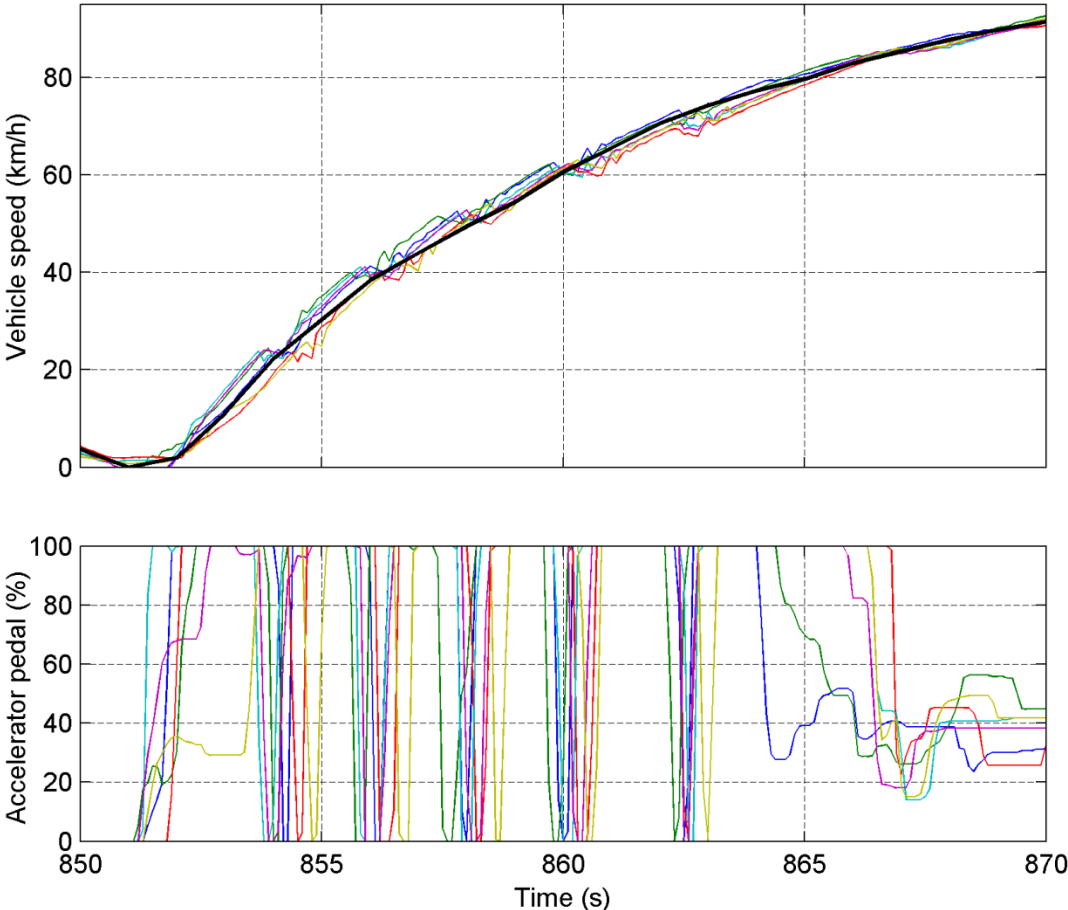

Figure 7-7: Speed and accelerator pedal position during a harsh acceleration in the LA92. The bold black line is the scheduled test speed; the data series presented represent all 6 baseline (hybrid off) tests. The high power requirement of the test cycle meant the vehicle often struggled to keep to the speed trace even with maximum activation of the accelerator pedal.

154

Nevertheless is it possible that other factors also played major roles, for example it may be that the transmission efficiency of the EM pully system, assumed in simulation to be 98%, was not this high. Similarly the motor controller used on-vehicle was a smaller model than that used during motor mapping, and so its MOSFETs may have been operating nearer their load capacity and therefore at slightly lower efficiency.

7.2.6 Limitations of Backwards Modelling

As described in Chapter 4 the simulation is arranged with a backwards architecture, such that it is assumed that the speed trace is perfectly achieved, and the fuel consumed in doing so is calculated. This approach has limitations, of which the power deficit observed during dynamometer testing (Figure 7-7) is exemplary. Since the simulation assumed that the speed profile was perfectly followed it was incapable of detecting that this would have been unachievable in practice.

Two problems are inherent: Firstly, the transient engine performance differs from steady state performance, due to turbocharger lag for example, and so even though the accelerator pedal was fully depressed the engine did not achieve a torque output equal to the maximum achieved at steady state. Secondly, the absence of tractive power during gear shifts meant that more power was required transiently to make up for the temporary deficit. These issues highlight the limitations not only of the backwards modelling approach, but also of quasi-static engine modelling.

Some alternative simulation architectures, such as forward simulation and backward-forward hybrid simulation, do exist however these do not necessarily have the same potential for model-based control design exploited here. Furthermore these simulation architectures may not in fact have resolved the limitations in this case, since the power requirements *were* within the capability of the engine according to the quasi-static engine model – a transient engine model would therefore also be required.

In addition to problems during accelerations it was found that the simulation also had shortcomings in modelling the interaction between the hybrid system and the vehicle during braking phases. Specifically, it was found that the controller's decision to
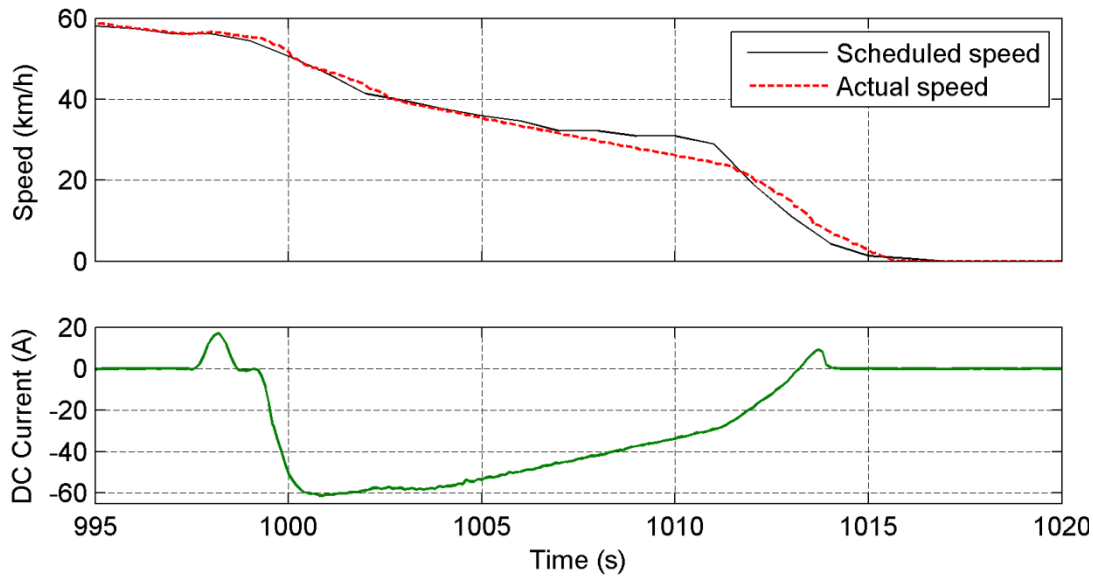
Figure 7-8: During regeneration the actual vehicle speed occasionally fell below the scheduled speed, which the driver was unable to correct without violating the gear shift schedule. A backwards simulation is unable to predict this.

regenerate would in some instances cause the vehicle speed to fall below the scheduled speed due to added braking torque, however it was not always possible for the driver to correct for this. As an example, Figure 7-8 presents a deceleration during which the drive cycle gear shift schedule specified neutral gear be engaged at 1003 seconds; from this time on it was therefore impossible to supply positive tractive force using the engine without violating the gear shift schedule. Initially the driver employed the mechanical brakes to decelerate the vehicle, and the control strategy responded appropriately by demanding maximum regeneration current. As the vehicle speed fell below the scheduled speed the driver released the mechanical brakes, however seeing a constant deceleration the hybrid control strategy continued to regenerate electricity. The regeneration was therefore self-perpetuating and the braking force generated by the EM was enough to cause the vehicle speed to fall well below the scheduled speed.

In the case of real-world driving such an episode would not present any real problem, and so to some extend this is a non-issue generated by the peculiarity of following a prescribed speed trace. Nevertheless this raises questions as to how this would be dealt with during a legislative test. These sorts of issues are not foreseeable using a backwards model because the speed is explicitly defined by the driving cycle.

156

Similarly the control strategies were not augmented with enough intelligence to understand the effect of regeneration demand on vehicle acceleration, for example.

7.3 Chapter Conclusions

Detailed analyses of the data collected during both the driver assistance and hybrid control elements of this work yield noteworthy observations. For the *Lightfoot* data this includes the observation that rates of deceleration have been inadvertently reduced in addition to the intended effect on rates of positive acceleration, and the potentially beneficial influence this could have on recoverable energy. For the HEV controller development a generalised procedure for good-practice implementation of the SDP algorithm has been proposed. Whilst the trends resulting from HEV testing shown here are believed robust it is also clear that there is plenty of potential for further development and refinement. Repeating the study with a more heavily hybridised vehicle would certainly improve the confidence in results, and doing so with test procedures and instrumentation that would allow system state of health metrics to be recorded would be yet more interesting. The shortcomings of the backwards modelling approach, quasi-static modelling, and testing over a fixed velocity trajectory are also important observations.

Improvements can be envisaged in terms of how each system is conceived (including the possibility for interaction), the details of each system's design, and also the procedures used to evaluate performance. Further detailed conclusions and possibilities for further work are outlined in Chapter 8.

Chapter 8   Conclusion

This chapter summarises the main findings of this work, namely the 7.6% reduction in fuel consumption achieved through driver behaviour modification and the 13% reduction in electrical powertrain stress achieved through hybrid control optimisation. Opportunities for further work are identified in each field separately, as well as the possibility of greater synergy between systems.

8.1 Hybrid Control

This work has examined the real-world potential for predictive control of HEVs using Stochastic Dynamic Programming. A set of control strategies were developed for a hybrid LCV based on recorded data from similar vehicles operating in their daily commercial routine. Of the common legislative test cycles the LA92 was found to be the closest statistical match. The control strategies were tested in simulation before being implemented on-vehicle and evaluated both on a chassis dynamometer and on-road.

The completeness of this work in designing a SDP controller using a representative real-world historic data set, applying this control strategy to a physical vehicle and testing it over a suitable test cycle is unique. During this process several practical challenges were resolved and the methods used here are offered to others wishing to implement the algorithm; these included selection of an appropriate test cycle, determining a suitable value for the discount factor (λ) and number of necessary policy iterations, determining and adjusting the nominal SOC which a controller yields by use of probabilistic charge sustenance plots and a simple method for compressing the size of the resulting look-up table so as to be implementable on hardware with restricted memory allocation.

A novel cost function was used during the optimisation which incorporated a measure of battery stress and electrical heating. By varying the weighting of this component of the cost function the trade-off between fuel consumption and electrical component stress could be examined. It was found that in simulation a 20% sacrifice in potential fuel savings could yield a 51% reduction in electrical stress, and during dynamometer testing a 13% reduction in stress was achieved without any sacrifice of fuel savings. The stress metric showed potential to reduce thermal loading on the battery and motor which will yield dividends in real-world operation, or alternatively a corresponding reduction in the size of the battery pack carried, yielding weight and cost savings. Further benefits for battery management are also likely, but were not detectable within the scope or resource of this work.

Performance of the SDP controller set was directly compared against a set of similar ECMS controllers and found to achieve a lower mean square battery C-rate, though

the fuel consumption savings for both types of controller were similar. It was found to be difficult to pinpoint exactly how the SDP achieved superior results to the ECMS, no doubt partly because the power capacity of the hybrid system was relatively small in comparison to the total installed vehicle power. Both controller types appear to follow a similar regime, suggesting that it may be possible to achieve behaviour very similar to the SDP using ECMS, which is considerably more straightforward and easier to calibrate.

For on-road testing a SDP controller was adapted by replacing the acceleration state, which the controller's decision was based on, with engine load. The control surface was correspondingly transformed into this state-space using the correlation between engine load and acceleration at zero road slope. Although this process is not ideal and will have incurred some loss of optimality the resulting controller appeared to operate robustly and managed SOC well.

8.2 Driver Behaviour

In the course of this work a driver assistance system was developed with the aim of ensuring the additional power installed in a hybridised vehicle is not abused, while also facilitating a reduction in fuel consumption through eco-driving techniques. The device encourages drivers to restrict their rates of acceleration and enforces the advice of a GSI, therefore reducing the average engine speed.

When applied to a test fleet of 15 light commercial vehicles the fuel consumption of the fleet, weighted by the distance travelled by each vehicle, was reduced by 7.6%. It was noted that the savings of individual vehicles/drivers varied considerably, with the maximum saving being 12.0%. Changes in driver behaviour and fuel consumption were achieved without any impact on average vehicle speeds.

The device presented represents an improvement on those developed by other researchers because its relative simplicity allows easy integration into vehicles, using only signals available on the vehicle CAN-bus accessed via the OBD port, thereby avoiding the need for dedicated sensors. Furthermore the device is safe for real-time

use as it does not require the active attention of the driver and so is likely to introduce minimal additional cognitive loading.

8.3 Further Work

The behaviours of SDP and ECMS controllers were found to be broadly similar and the primary difference between their cost functions was shown to be in the value they assign to electrical energy. Whereas SDP assumes a value based on a stochastic model of future operating conditions the ECMS uses a fixed value model. It may be possible to produce SDP-like behaviour by adopting an Adaptive ECMS (A-ECMS) approach where the equivalence factor is allowed to vary slightly, a basic version of which was proposed by Musardo et al. [80]. This approach may yield a better compromise between optimality and ease of use than does SDP.

During trials of the driver assistance system a considerable range of fuel savings were recorded, from a minimum of 0.43% to a maximum of 12.0%. There are several plausible reasons as to why some drivers were able to achieve substantial savings whilst others were not: it could be that some drivers were very conservative to begin with, or it may be that the drive cycles of some vehicles allowed greater savings than others. Further analysis is required to establish the mechanisms of fuel saving, as well as the reasons for the large range of observed savings.

A logical development of the driver assistance system would be to introduce a degree of adaptive behaviour such as that proposed by Wada et al. [81], so that the sensitivity of the system is no longer fixed during operation. This would allow the system to continually encourage drivers at an appropriate level, driving continuous improvement without becoming irritating.

Finally, it was noted that as well as causing drivers to accelerate more gently the driver advisory system also inadvertently caused drivers to decelerate more slowly. Although this could in principle have increased the amount of braking energy per kilometre which the hybrid system was able to capture, analysis showed that this was in fact almost exactly offset by a reduction in total braking energy throughput. Nevertheless it may yet be that a concerted effort to integrate rate of deceleration into

the system logic could further increase the proportion of recoverable energy, making the total fuel savings achieved by the two systems more than the sum of their independent contributions. These observations give rise to the next set of research questions which may be explored in furtherance of the work presented here.

Appendix 1 – LA92 gear shift schedule

Table A1: Gear shift points for the LA92 determined using the gear shift survey method.

Time (s)	Gear	Time (s)	Gear	Time (s)	Gear
0	0	499	4	1001	0
25	1	503	5	1027	1
36	2	511	4	1035	2
54	0	526	0	1039	3
67	1	535	1	1050	4
75	2	553	0	1058	0
84	1	576	1	1061	1
86	2	585	2	1065	2
89	3	588	3	1068	3
97	4	596	2	1071	4
104	0	599	3	1076	5
112	1	602	4	1083	0
122	2	607	5	1091	1
124	3	621	0	1095	2
126	4	625	1	1117	3
137	0	627	2	1120	4
147	1	632	3	1124	5
154	2	647	4	1128	6
156	3	666	5	1160	0
158	4	673	6	1169	1
161	5	683	5	1193	0
165	6	703	0	1195	1
173	5	707	2	1203	2
179	4	710	3	1205	3
185	5	714	4	1208	4
190	6	732	0	1219	0
207	0	734	2	1244	1
211	1	739	0	1250	2
215	0	767	1	1254	3
242	1	774	2	1257	4
249	2	785	3	1260	5
253	3	818	0	1268	6
256	2	820	1	1283	5
262	3	854	2	1286	6
271	4	856	3	1332	4
277	5	858	4	1359	5
286	0	860	5	1374	0
317	1	862	6	1391	1
325	2	955	0	1402	2
327	3	964	1	1420	0
330	4	973	2		
333	5	976	3		
337	6	979	4		
495	2	983	5		
497	3	989	6		

Appendix 2 – Journal Paper 1 (Proc. IMechE)

The following article written by the author includes some of the work presented in this thesis. The article appeared in the peer-reviewed journal *Proceedings of the Institute of Mechanical Engineers Part D: Journal of Automobile Engineering*.

Article information

Authors:	C. Vagg, C. J. Brace, S. Akehurst, R. Wijetunge, and L. Ash
Title:	Development of a new method to assess fuel saving using gear shift indicators
Journal:	Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering, vol. 226, pp. 1630-1639
Publication year:	2012

Development of a new method to assess fuel saving using gear shift indicators

Christopher Vagg[1], **Christian J Brace**[1], **Roshan Wijetunge**[1],
Sam Akehurst[1] and **Lloyd Ash**[2]

Abstract

European regulations set the emissions requirements for new vehicles at 130 g CO_2/km, with an additional 10 g CO_2/km to be achieved by additional complementary measures, including gear shift indicators. However, there is presently little knowledge of how much fuel or CO_2 could actually be saved by the introduction of gear shift indicators, and there is no consensus on how these savings should be quantified. This study presents a procedure which allows these savings to be quantified over a New European Driving Cycle, and explores the trade-off between fuel savings and drivability. A vehicle model was established and calibrated using data obtained from pedal ramp tests conducted at steady speed using a chassis dynamometer, significantly reducing the time required to generate a calibration data set when compared with a steady-state mapping approach. This model was used for the optimisation of gear shift points on the New European Driving Cycle for reduced fuel consumption subject to drivability constraints. During model validation the greatest fuel saving achieved experimentally for a warm engine was 3.6% over the New European Driving Cycle, within the constraints imposed using subjective driver appraisal of vehicle drivability. The same shift strategy for a cold start driving cycle showed a fuel saving of 4.3% over the baseline, with corresponding savings in CO_2 of 4.5% or 6.4 g CO_2/km. For both hot and cold tests the savings were made entirely in the urban phase of the New European Driving Cycle; there were no significant differences in fuel consumption in the extra-urban phase. These results suggest that the introduction of gear shift indicators could have a substantial impact, contributing significantly towards the 10 g CO_2/km to be achieved by additional complementary measures when assessed in this way. It is not clear whether these savings would translate into real world driving conditions, but for legislative purposes an assessment procedure based on the New European Driving Cycle remains a logical choice for simplicity and continuity.

Introduction

Regulatory framework

UN/ECE Regulation No 83[1] describes specific testing procedures for determining fuel consumption and tailpipe emissions of passenger and light commercial vehicles during Type Approval. The Type 1 test is used to verify tailpipe emissions and fuel consumption during the New European Driving Cycle (NEDC) on a chassis dynamometer after a cold start. For vehicles with manual transmission the Driving cycle defines both the speed trace of the vehicle, as well as the points at which the driver must change gear. The NEDC may be regarded as comprising two parts: urban and extra-urban, as defined in Annex 4a of UN/ECE Regulation No 83. Part 1 consists of four repetitions of the Elementary Urban Cycle, and Part 2 consists of a single Extra-Urban Cycle, as shown in Figure 1.

Following concerns that motor vehicle manufacturers would not meet their commitments to reduce new car CO_2 emissions to 140 g CO_2/km by 2008–9, the

---

[1]University of Bath, Bath, UK
[2]Ashwoods Automotive Ltd, Exeter, UK

Corresponding author:
C Vagg, University of Bath, Bath BA2 7AY UK.
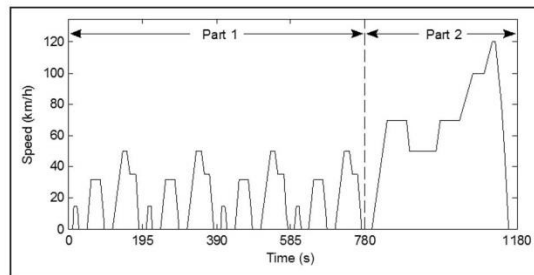Email: C.R.M.Vagg@bath.ac.uk

Figure 1. The New European Driving Cycle (NEDC) broken down into Part 1 (Urban) and Part 2 (Extra-Urban).

Table 1. CO_2 savings achieved by following the GSI.[4]

	% Reduction in CO_2		
	Part 1	Part 2	NEDC overall
BMW Mini	6.45	2.15	4.16
VW Golf	0.12	−0.23	−0.07
Ford Transit	13.53	1.65	6.68

European Union issued Regulation (EC) No 443/2009.[2] This Regulation sets the average emissions requirement for the new car fleet in the European Community in 2012 at 130 g CO_2/km, with an additional 10 g CO_2/km to be achieved by additional complementary measures such as gear shift indicators (GSIs). A GSI is defined by Regulation (EC) No 661/2009[3] as 'a visible indicator recommending that the driver shift gear', and under the Regulation[3] are a requirement for all cars Type Approved from 1 November 2012.

Interestingly, under Annex 14 of UN/ECE Regulation No 83, hybrid electric vehicles (HEVs) fitted with GSIs have been permitted to ignore the standard gear shift points and instead use the speed trace designed for vehicles with automatic transmission with gear shifts effected by the driver as advised by the GSI. For those HEVs equipped with a GSI it is reasonable to speculate that some of the fuel savings made during a Type 1 test result from the exemption from the standard gear shift schedule.

There is not yet an agreed method to determine the real fuel or CO_2 savings that will be delivered by GSIs, and, therefore, whether the intended 10 g CO_2/km savings to be delivered partly by GSIs are realistic. The Department for Transport commissioned a report, undertaken by AEA and Millbrook Proving Ground, to assess the potential impact of GSIs and to propose a method to quantify this.[4] This report highlighted the lack of high quality published data on the subject, and also noted that the savings achieved would vary considerably for vehicles of different mass, fuel, usage patterns and so on. One of the proposed methods to quantify the savings achieved by a GSI was to perform Type 1 tests, first following the standard legislative gear shift points and then those suggested by the GSI, and to compare the two sets of results. This was undertaken for three vehicles, for which results are summarised in Table 1. For each test configuration one test was completed (a total of two tests per vehicle), and the nominal repeatability of the chassis dynamometer with an experienced user and driver (2% coefficient of variance (CV)) was deemed sufficient.

Ngo et al.[5] showed fuel savings of 11.2% in Part 1 of the NEDC when optimising the shift strategy of a

Mitsubishi Colt equipped with automated manual transmission (AMT). Casavola et al.[6] showed savings of 9.1% over a NEDC when optimising the shift strategy of a small car in a simulation environment.

Objectives

This study responds to the present state of legislation and the needs identified in the literature to achieve the following objectives:

- *Objective 1*: Propose a method to determine the optimum gear shift schedule for reduced fuel consumption subject to drivability constraints, to aid future development of GSI logic.
- *Objective 2*: Follow the procedure set out by AEA/Millbrook in order to add to the body of high quality published data on the efficacy of GSIs, thereby responding to the need identified.

Methods

In order to meet the defined objectives, this study presents the development of a simulation tool based on empirical test results and supported by subjective appraisal of vehicle drivability. The resulting simulation is used for gear shift strategy optimisation on the NEDC, and determines the fuel savings which are possible through this alone. The method used to optimise the gear shift schedule is explained, which allows the trade-off between reduced fuel consumption and drivability to be explored. Empirical test results for the validation data sets are presented, following the procedure outlined by AEA/Millbrook.

The test vehicle used in this investigation was a Citroen Berlingo, a light commercial vehicle, for which full specifications can be seen in Table 2. Experiments were carried out on the chassis dynamometer test facility at the University of Bath. Two methods of emission measurement were used: tailpipe gas sampling (CO_2 tracer method), and bag analysis. In the former, a continuous sample of tailpipe gas is passed through emissions analysers in real-time and recorded at 10 Hz. In the latter, a continuous sample of the same gas is collected in a bag throughout each part of the test. The contents of each bag are passed through analysers post-test in order to infer the total mass of emissions throughout the test. Bag analysis is typically more accurate, but does not allow any insight into the instantaneous emissions
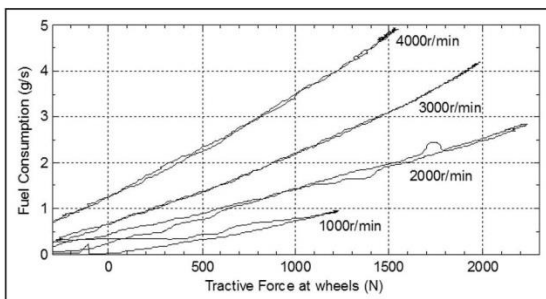
169

Table 2. Specifications of the test vehicle.

Vehicle type	Citroen Berlingo
Vehicle model	LX 625 L1
Fuel	Diesel
Engine	1.6HDi 90hp
Kerb weight (kg)	1354
Transmission – Gear 1	11/38 (0.290)
Transmission – Gear 2	15/28 (0.536)
Transmission – Gear 3	32/37 (0.865)
Transmission – Gear 4	45/37 (1.216)
Transmission – Gear 5	50/33 (1.515)
Transmission – Final Drive	17/73 (0.233)

production.[7] In both cases a carbon balance approach was used to find the fuel consumption. The test facility is temperature controlled, and was maintained at 25°C throughout all tests. Brace, Burke and Moffa[8] demonstrated that during chassis dynamometer fuel consumption measurement three of the most influential variables are vehicle battery charge, oil level and tyre pressures. Regular vehicle checks were performed in order to limit the effects of these variables.

Vehicle mapping

A backward facing computational model of the vehicle driveline was built using Simulink. In order to determine the fuel consumption characteristics of the vehicle a series of mapping tests was conducted. These tests were performed on the chassis dynamometer rather than by testing the engine separately. The advantage of this approach is that a direct link may be drawn between tractive force at the wheels and fuel consumption, accounting for gearbox and tyre losses in the measurement, rather than adding these in later. Ideally, tests should then be performed in each gear in order to account for differing efficiencies and losses, and, therefore, populating steady-state maps would be a considerable undertaking. For this reason the driveline was mapped by using the dynamometer in closed-loop speed control mode at a range of speeds, and applying a throttle ramp to achieve the full range of tractive

effort. Each throttle ramp was applied over 2 min, held briefly, and then released over a period of 2 min. An example of the data collected by this method in fourth gear is presented in Figure 2; some speeds have been omitted for clarity.

The advantage of this method is that it provides a fuller picture of engine performance than may be obtained through steady-state testing. However, whilst the traces are very repeatable at high speeds and loads, at low speeds and loads there is considerable hysteresis, which produces ambiguity. Another observation from this exercise is that the engine does not normally produce positive tractive force until the throttle signal is > 25% and the tractive force production is typically saturated at approximately 70% throttle. Consequently although 4 min of transient data is collected for each speed, only about 45% of this is useful information. Furthermore, since the NEDC is a relatively low power drive cycle, a large portion of this usable map is never visited. If this exercise were to be repeated, the authors would suggest considering these facts, and perhaps modifying the throttle ramp trace to ensure accuracy is maximised in the areas of the engine map which are of most interest.

The fuel consumption data recorded by the tailpipe gas sampling (CO_2 tracer) method were parameterised as a set of quadratic surfaces, with one surface per gear. Manipulation of these surfaces allowed plots such as those presented in Figures 3 and 4 to be constructed which are useful when considering possible gear shift strategies. Note that the engine torque has been calculated from the tractive force at the wheels, assuming 100 per cent transmission efficiency. For this reason it is referred to as *net engine torque*, and represents the useful torque delivered by the engine to the wheels.

Model validation

In order to validate the model a set of heuristic gear-shift schedules was designed with reference to Figures 3

Figure 2. Fuel consumption data obtained in fourth gear using the throttle ramping technique.
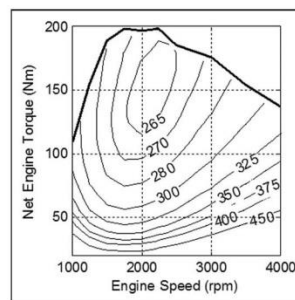
Figure 3. Brake Specific Fuel Consumption (BSFC) plot and torque limit curve constructed from third gear data. Contours show BSFC in units of g/kWh. Net engine torque represents useful engine torque delivered to the wheels, accounting for driveline inefficiencies.

and 4, to give a range of shifting behaviour. Each upshift was defined at a road speed, which implies an engine speed, as defined by the vehicle transmission ratios. Since the objective of minimising fuel consumption inevitably leads to upshifting as soon as possible, a number of constraints were placed on the possible gear shift schedules: each upshift must occur at a higher road speed than the previous, and a minimum post-upshift engine speed was defined in order to ensure drivability.

The speed trace defined in UN/ECE Regulation No 83 for manual vehicles includes steady speed periods of 2 s, which correspond to gear shift points. Since gear shift points were to be moved throughout this investigation, it was decided instead to use the speed trace defined for vehicles with an automatic transmission, which does not include these sections; this is the procedure used for testing HEVs with a special shifting strategy. The gear shift points for the standard (baseline) schedule used here are fixed at the *times* defined in UN/ECE Regulation No 83. Since the shape of the speed trace for automatic vehicles is subtly different to that for manual vehicles, the upshift speeds quoted with respect to this speed trace differ very slightly to the standard upshift speeds using the manual trace; these differences are not considered significant. Using the standard gear shift schedule (applied to the speed trace for automatic vehicles) the engine speed after an upshift is typically in the region 1300–1700 r/min for the vehicle described in Table 2. Three modified schedules were designed such that the minimum engine speed after upshifting was greater than 1300, 1200 and 1100 r/min. These schedules shall be referred to as A, B and C for convenience, and the standard schedule shall be referred to as STD. Downshifts were effected when the vehicle speed dropped 0.5 km/h below each upshift speed to allow hysteresis and avoid excess shift busyness, and the de-clutch points were unaltered from the standard schedule. Table 3 shows the road speed (km/h) at which gears 2 to 5 were engaged for each schedule.

It is standard test procedure for the NEDC to be conducted from a cold start, and as such it is preferable

Table 3. Upshift speeds (km/h) for each gear schedule. Note that the STD shift speeds are found by taking the time stamp of the shift in the manual trace, and applying this to the automatic speed trace.

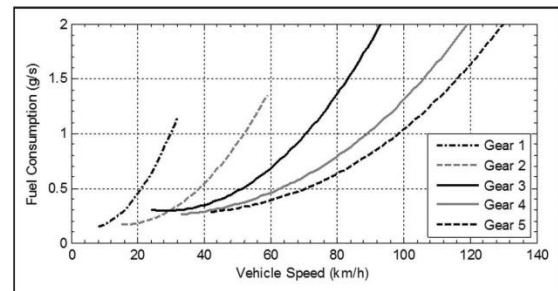|	Gear shift schedule			
	STD	A (1300)	B (1200)	C (1100)
Gear 2	15.1	21.5	21.5	19.3
Gear 3	32.7	32.0	30.0	30.0
Gear 4	50.1	46.9	41.9	36.9
Gear 5	68.5	65.3	65.3	50.0

Figure 4. Constant velocity fuel consumption of the vehicle in each gear, predicted by the quadratic engine model.

to conduct one test per day. However, the simple engine model used here is only able to simulate a warm engine, and it was, therefore, decided to test the above schedules from a hot start, with reference to engine sump oil temperature. Therefore, the vehicle was driven at a high speed until the engine oil temperature reached 90°C immediately prior to each test.

The effect of modifying the shift speeds on gear selection is demonstrated in Figure 5, which shows traces of vehicle speed divided by engine speed in the first Elementary Urban Cycle of the NEDC for each gear shift schedule. During the second microtrip, schedule STD only uses gears 1–2, while schedules A, B and
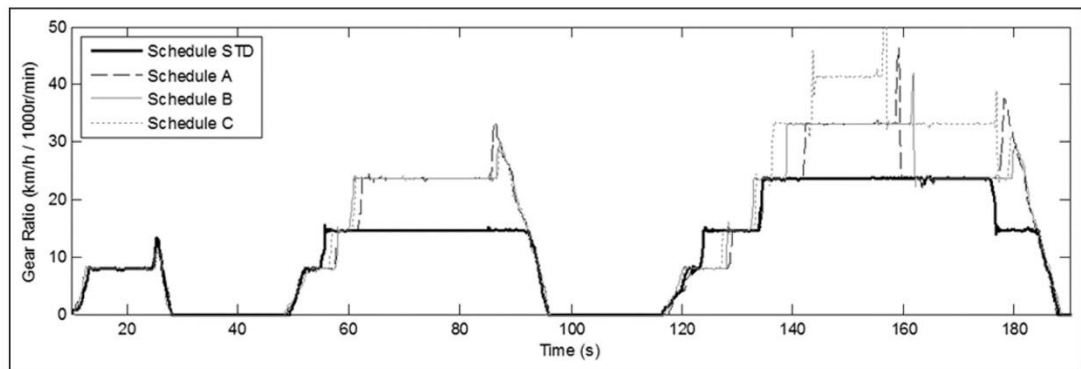

Figure 5. Traces showing the gears engaged during the first Elementary Urban Cycle.

Table 4. Fuel consumption changes for hot start tests.

	STD	A	B	C
Number of repeats (n)	6	4	5	5
Fuel used (g)	445.5	439.4	435.4	429.6
Standard deviation (SD) (g)	4.90	4.31	5.23	2.71
Coefficient of variance (CV) (%)	1.10	0.98	1.20	0.63
Fuel saving over STD (%)	–	1.4	2.3	3.6
Statistical significance (p)	–	0.0767	0.0089	< 0.001

Table 5. Fuel consumption changes broken down into Part 1 and Part 2 of the NEDC.

	Part 1 (urban)				Part 2 (extra-urban)			
	STD	A	B	C	STD	A	B	C
Fuel used (g)	195.7	186.5	183.6	178.1	249.8	252.9	251.8	251.5
Std deviation (SD) (g)	2.13	2.40	3.15	1.05	3.51	2.67	4.67	1.69
CV (%)	1.09	1.29	1.71	0.59	1.41	1.06	1.85	0.67
Fuel saving (%)	–	4.7	6.2	9.0	–	−1.2	−0.8	−0.7
Statistical significance (p)	–	< 0.001	< 0.001	< 0.001	–	0.19	0.45	0.35

C all make use of gear 3. (A microtrip is defined as beginning when the vehicle is accelerated from rest to having positive velocity, and ends when the vehicle returns to rest. Thus the Elementary Urban Cycle consists of three microtrips and the NEDC consists of 13 microtrips.) Similarly, in the third microtrip, schedule STD only uses gears 1–3, while schedules A and B make use of gear 4 and schedule C also uses gear 5. The spikes which are frequently seen towards the end of a microtrip are a result of the driver de-clutching and a subsequent drop in engine speed. These should not be misinterpreted as a brief upshift.

Subjective driver feedback was collected for each gear shift schedule to assess the drivability of the vehicle. This feedback indicated that schedule C was close to the acceptable limit of drivability; it was possible to achieve the speed trace without the engine nearing the stall condition, but the vehicle was becoming noticeably more difficult to drive due to reduced torque reserve. (Torque reserve is the difference between current engine torque and maximum engine torque at the current engine speed. In test cycle driving it is important to have adequate torque reserve so that, if the current vehicle speed falls below the required vehicle speed, it can be quickly corrected. In real-world driving torque reserve determines how responsive the vehicle feels to a step input in throttle.) In order to maintain a realistic driving style with a suitable torque reserve it was decided not to further reduce the upshift speeds.

Results

Cumulative fuel use and emissions figures quoted from this point onwards are measured using the bag analysis method for increased accuracy. Instantaneous fuel consumption measured according to the CO_2 tracer method was used only to check whether the behaviour of the model was sensible.

Hot start tests

Test results for the complete NEDC performed from a hot start are presented in Table 4 for each gear shift schedule. As can be seen, moving the gear shift points has a significant impact on fuel consumption, and lowering the minimum engine speed constraint reduces the total fuel use. During the engine mapping work undertaken for model development it was found that the vehicle engine was more efficient when operating at low speed and high load (Figure 3). This finding is typical of diesel engines and so it is not surprising that reducing engine speed causes a reduction in fuel use. Table 4 also shows *p*-values resulting from two-tailed unpaired *t*-tests comparing schedules A, B and C with schedule STD, showing that the fuel savings for schedule A are not significant at the 95% confidence level, whilst those for B and C are significant with confidence levels of 99% and 99.9%, respectively.

Breaking down the results into urban (Part 1) and extra-urban (Part 2) sections of the NEDC allows further insight into where the majority of the fuel savings are made. These data are presented in Table 5. As can be seen, all of the savings are in fact made during the urban part where the maximum fuel saving is 9.0% (p < 0.001) for schedule C. In contrast, statistical analysis suggests that differences in the extra-urban part are not significant at the 90% confidence level and, therefore, must be disregarded. The result that fuel is saved in the urban phase but not in the extra-urban phase is due to the extra-urban phase containing fewer gear changes and predominantly making use of higher gears; this will be explored fully in the discussion. It can also
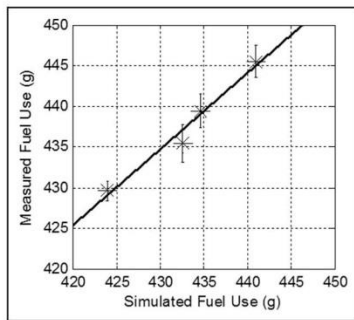
Figure 6. Simulated and measured fuel consumption results. The trend line is described by $y = 0.938x + 31.5$. Error bars for measured results show $\pm$ standard error of the mean (SEM).

Table 6. Comparison between simulated and measured fuel consumption (FC) results over a NEDC.

	Measured FC (g)	Simulated FC (g)	Error (%)
STD	445.5	441.0	−1.01
A	439.4	434.6	−1.08
B	435.4	432.7	−0.63
C	429.6	424.0	−1.31

be seen from Figure 5 that schedules A, B and C actually cause gear 2 to be engaged slightly later than schedule STD, and, therefore, this upshift speed could have been reduced further.

Results predicted using the simulation are given in Table 6 and compared graphically with empirical test results in Figure 6. It can be seen that the simulation predicts the behaviour of the vehicle with a good degree of accuracy, especially when the relatively crude nature of the model is considered. Furthermore, the engine mapping techniques did not allow for any speed transient effects of the engine, such as those of the turbocharger, to be characterised.

Figure 7 compares the measured and simulated instantaneous fuel consumptions, and it can be seen

that there is a good correlation. Note that the simulation output trace consists of many relatively linear sections and sudden changes as it is based on the ideal speed trace input to a backwards model. When the actual recorded speed trace is fed into the model, many of the features recorded in the measured fuel consumption trace are reflected in the simulated trace. Fine tuning of the model to the instantaneous fuel consumption signal generated from experimental testing must be approached with caution, due to the difficulties encountered when estimating tailpipe volume flowrate using the CO_2 tracer method.[7,9] The spikes in fuel consumption observed at tip-outs and also at 290 s and 385 s are fictitious, and are the result of precisely these difficulties. Similarly, the step response of this measurement technique can be poor: clear examples of this are exhibited when idle fuelling is resumed following deceleration-fuel-shut-off at 225 s, 295 s and 390 s. For this reason comparison of the traces in Figure 7 should be regarded only as a guide to check that the instantaneous model behaviour is sensible. A better impression of how the model performs instantaneously (particularly during transients) could have been gained if access to the engine control unit (ECU) fuelling signal were available. Unfortunately, this signal was not available, and the quality of the results obtained in spite of this may be regarded as a strength of the approach taken.

Figure 8 compares cumulative fuel consumption of the simulation with empirical cumulative fuel consumption, measured by the bag analysis technique, for a test with schedule STD. The first two bag-measured points are taken half way through and on completion of Part 1; and the final point is taken on completion of Part 2. It can be seen that the trace produced by the simulation passes very close to the bag values, correctly predicting the proportions of fuel used in the urban and extraurban phases.

Cold start tests

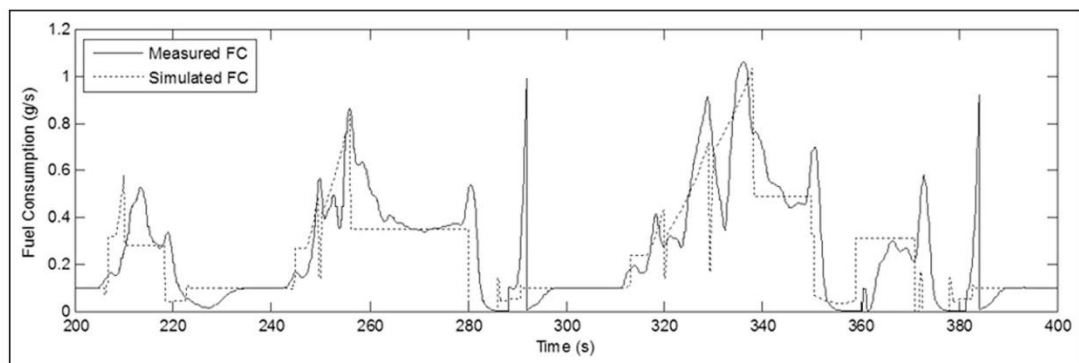The test procedure followed above to determine the savings achieved by following a GSI could be

Figure 7. Measured and simulated fuel consumption (FC) during the second Elementary Urban Cycle of a test with schedule STD.

Table 7. Fuel consumption changes for cold start tests.

	Part 1 (urban)		Part 2 (extra-urban)		Total NEDC	
	STD	C	STD	C	STD	C
Number of repeats (n)	9	3	9	3	9	3
Fuel used (g)	229.8	208.0	264.3	264.8	494.1	472.8
Std deviation (SD) (g)	2.40	1.87	3.92	0.62	4.89	1.65
CV (%)	1.04	0.90	1.48	0.24	0.99	0.35
Fuel saving (%)	–	9.5	–	−0.2	–	4.3
Statistical significance (p)	–	< 0.001	–	0.8281	–	< 0.001
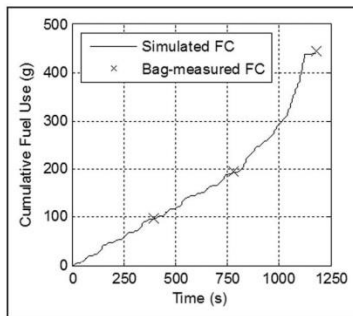
Figure 8. Cumulative simulated fuel consumption and empirical 'bag' values for fuel consumption for the same test presented in Figure 7.

followed using cold or hot tests. Although the model may not be used to predict the absolute fuel consumption for a cold test, it is likely that the shift schedules derived for hot tests maintain some validity for cold tests. Since the current legislative test is from a cold start it was decided to select the gear shift schedule from A–C which was found to give the best fuel economy improvement for hot tests, and to apply this to a cold start test (in line with the procedures for a Type 1 test). Though this schedule may not be regarded as strictly optimal in any sense when applied to a cold engine, the results are still of interest and so are presented separately.

Table 7 shows the results of Schedules STD and C implemented as cold start tests. Fuel savings of 4.3% ($p < 0.001$) were achieved over the whole NEDC, and a reduction of 9.5% ($p < 0.001$) was observed during Part 1. Once again the differences in Part 2 (extra-urban) are not significant at the 90% confidence level and, therefore, must be disregarded.

Discussion

Overall it can be seen that the repeatability of the tests is very good, with the CV nominally around 1% over a complete drive cycle. During the tests the vehicle was controlled by a human driver, which inevitably introduces some variation. An alternative would have been to use a robot driver, which will be considered for future test programmes.

The accuracy of the model is extremely encouraging, given its simplicity. As previously mentioned, the mapping tests revealed considerable hysteresis in engine characteristics, particularly at low loads, such as those experienced during cruise phases. This hysteresis adds a degree of uncertainty into the model at operating conditions which are frequently visited. More advanced modelling techniques could be used to fully characterise the driveline, but it is likely that the net result on fuel savings would be minimal.

As previously mentioned, the empirical tests showed significant fuel savings in Part 1 (urban phase), but no
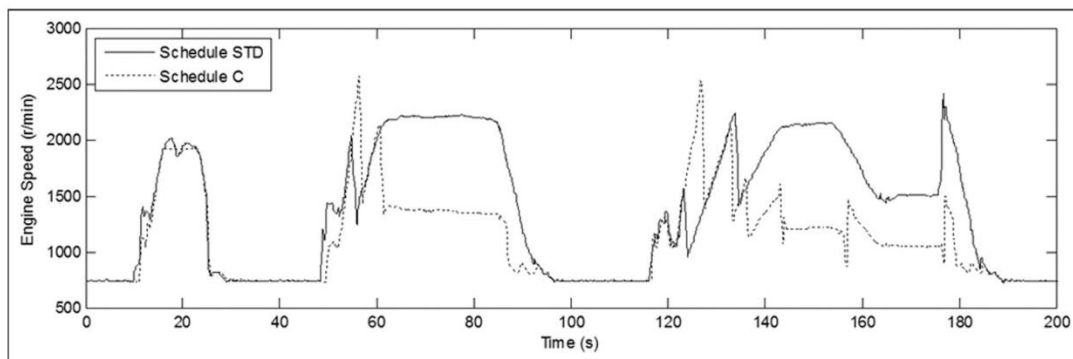
Figure 9. Engine speed traces for schedules STD and C during the first Elementary Urban Cycle.
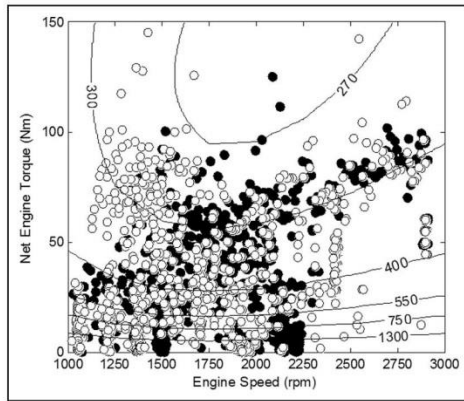
Figure 10. Engine operating points for schedules STD (●) and C (○) during an NEDC. Contours show BSFC in units of g/kWh, as in Figure 3.

Table 8. Upshift speeds (km/h) for the optimal shift strategy.

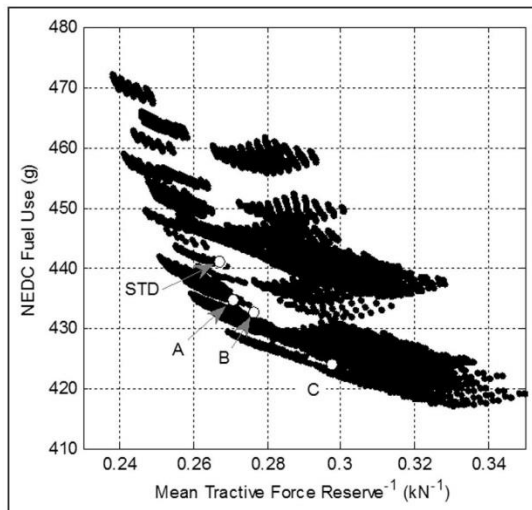|	Gear shift schedule		
	STD	C	OPT
Gear 2	15.1	19.3	16.5
Gear 3	32.7	30.0	30.4
Gear 4	50.1	36.9	36.9
Gear 5	68.5	50.0	48.1

Figure 11. The Pareto-optimal front arising from the attempt to minimise fuel consumption and maximise drivability.

change in Part 2 (extra-urban phase); this is probably due to several reasons. In Part 1 there are many more transients and a much lower proportion of time is spent cruising at constant speed (Figure 1). A consequence of this transience is that there are more gear changes than in Part 2 (16 compared with 6 for schedule STD, not
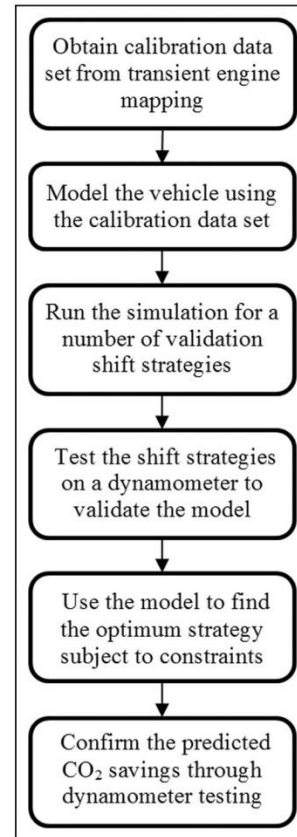

Figure 12. The recommended procedure to determine the optimum shift strategy.

including changing from neutral into gear 1) and, for this reason, there is more opportunity for adjustment. Furthermore, the changes made to the gear shift schedules in Part 2 are only slight: according to schedule STD, the steady speed section at 50 km/h followed by the acceleration from 50 to 70 km/h in Part 2 of the NEDC should be performed in gear 4, but in schedule C the vehicle is kept in gear 5. This is the only significant difference between any of the four schedules in this Part of the test – schedules STD, A and B are extremely similar. Figure 4 shows that the differences in fuel consumption between gears 1 to 4 are much more pronounced than the difference between gears 4 and 5 for a given road speed. Indeed, at 50 km/h the difference between gears 4 and 5 is negligible, and it is possible that inaccuracies in the mapping technique mean that this difference does not actually exist. For the acceleration from 50 to 70 km/h the increased road load may be expected to favour use of gear 5; however, the mapping set neglects the effects of speed transients, such as those of the turbocharger, which may be significant here. Any differences between gear shift schedules in Part 2 will also be diluted by the fact that more fuel is used than in Part 1.

175

Schedules A–C allow fuel to be saved by reducing the engine speed – Figure 9 shows that for schedule C the engine speed is reduced considerably for a large portion of the Elementary Urban Cycle. Since the power requirement of the NEDC remains unchanged, a reduction in engine speed necessitates an increase in net engine torque. This shifts the operating point of the engine towards the top left of the brake specific fuel consumption (BSFC) map, which tends to yield a higher efficiency, as shown in Figure 10. Note that the original data were sampled at 10 Hz, but the number of data points in Figure 10 has been reduced by a factor of five for clarity.

Having established confidence in the model, the upshift points were discretised, and an exhaustive search was conducted to find the optimal shift strategy. The upshift points were moved in steps of 50 r/min in both directions, to a maximum of 200 r/min above the baseline and a minimum of 500 r/min below the baseline. Observing the same constraints as before, a total of 13,060 viable gear shift schedules were generated, which could be run in a total of three hours on a desktop PC with the model running simulations at 10 Hz. The lowest fuel consumption achieved which also maintained a minimum engine speed of 1100 r/min was 420.3 g by schedule OPT, as shown in Table 8. Using the equation of the trend line in Figure 6 we may infer a real world fuel consumption of 425.6 g representing a 4.5% fuel saving over an NEDC – not far off that demonstrated by schedule C. Figure 11 shows the Pareto-optimal front that arises when seeking to optimise the gear shift strategy for fuel consumption whilst maintaining good drivability (torque reserve). The positions of schedules STD, A, B and C on this plot are indicated, and it can be seen that schedule C lies close to an optimal solution.

Procedure generalisation

In order to generalise the procedure set out above, Figure 12 formalises the process and lays out the logical progression of activities in order to identify the optimal shift strategy and estimate the CO_2 savings it would yield. This process, as described throughout this paper, has several significant strengths.

- A relatively simple backward facing model can be used to adequately represent the vehicle performance over a NEDC.
- Transient engine mapping can be used to obtain a calibration data set quickly (and, therefore, cost effectively). This could be achieved with as little as 1–2 days use of a chassis dynamometer if only two or three gears are mapped, and data is extrapolated to approximate the other gears.
- This procedure does not require access to vehicle manufacturer data, such as that available through on-board diagnostics (OBD), and so could be used in the design of retro-fit GSIs.

- The NEDC is a drive cycle that industry is well acquainted with, despite its weaknesses. Therefore, when investigating the potential impact of a technology for legislative purposes, figures presented with reference to the NEDC are well understood.

Conclusion

This paper has presented a test procedure which may be used to estimate the maximum fuel savings possible by moving the gear shift points in the NEDC, subject to drivability constraints. The procedure proposed, and the results presented from its implementation, are of use to legislative authorities in the development of appropriate measures of GSI efficacy, as well as to those involved in GSI design and vehicle modelling.

Application of this procedure to the vehicle examined yielded fuel savings of 3.6% over a hot-start NEDC (9.0% in the urban phase). Simulation work undertaken suggested that 4.5% fuel savings could be achieved, though a trade-off between fuel savings and drivability is inevitable. For a cold-start NEDC (Type 1 test) fuel savings of 4.3% were achieved during dynamometer testing, which equates to 6.4 g CO_2/km or 4.5% CO_2 saving. This was achieved by upshifting earlier than the prescribed NEDC gear shift points, so lowering the engine speed. These results support the findings of others, and suggest that the introduction of GSIs could make a significant impact.

These savings, estimated over a NEDC, could make a substantial contribution towards the 10 g CO_2/km to be achieved by additional complementary measures, though there remains a significant gap to be made up by other technologies. However, it is not clear whether these savings would translate into real world driving conditions, firstly because the NEDC is widely regarded as unrepresentative of real driving, and secondly because a GSI only provides advice, which real drivers may choose to ignore.[4] Other drive cycles, such as the ARTEMIS cycles, may give more representative results, but for legislative purposes the NEDC remains a logical choice for simplicity and continuity, especially considering that a similar procedure is presently permitted for HEVs. The savings that can be achieved by eco-driving techniques, including following the advice of a GSI, in the real world will be the subject of future work.

Automotive Ltd. No writing assistance was received in the preparation of this manuscript.

Conflicts of interest statement

The authors declare that there is no conflict of interest.

References

1. UN Economic and Social Council. UN/ECE Regulation No 83 Rev.4.2011.
2. European Commission. Regulation (EC) No 443/2009. OJ L 140, 5.6.2009, p.1.
3. European Commission. Regulation (EC) No 661/2009. OJ L 200, 31.7.2009. p.1.
4. Norris J, Walker H, Stones P, et al. Assessing the efficacy of gear shift indicators. Report, AEA/Millbrook, Oxford, 2011.
5. Ngo DV, Hofman T, Steinbuch M, et al. Improvement of fuel economy in Power-Shift Automated Manual Transmission through shift strategy optimization – an experimental study. In: *Vehicle Power and Propulsion Conference*, Lille, France, 1–3 September 2010, pp.1–5. USA: IEEE Computer Society.
6. Casavola A, Prodi G and Rocca G. Efficient gear shifting strategies for green driving policies. In: *American Control Conference 2010*, Baltimore, MD, 30 June–2 July 2010, pp.4331–4336. USA: IEEE Computer Society.
7. Hawley JG, Bannister CD, Brace CJ, et al. Vehicle modal emissions measurement – techniques and issues. *Proc IMechE Part D: J Automobile Engineering* 2004; 218(8): 859–873.
8. Brace CJ, Burke R and Moffa J. Increasing accuracy and repeatability of fuel consumption measurement in chassis dynamometer testing. *Proc IMechE Part D: J Automobile Engineering* 2009; 223(9): 1163–1177.
9. Bannister CD, Wallace F, Hawley JG, et al. Predicting instantaneous exhaust flowrates in a constant volume sampling system. *Proc IMechE Part D: J Automobile Engineering* 2007; 221(12): 1585–1598.

Appendix

Abbreviations

BSFC	brake specific fuel consumption
CV	coefficient of variance
FC	fuel consumption
GSI	gear shift indicator
HEV	hybrid electric vehicle
NEDC	New European Driving Cycle
SD	standard deviation
SEM	standard error of the mean

Appendix 3 – Journal Paper 2 (IEEE Trans.)

The following article written by the author includes some of the work presented in this thesis. The article appeared in the peer-reviewed journal *IEEE Transactions on Intelligent Transportation Systems*.

Article information

Authors: C. Vagg, C. J. Brace, D. Hari, S. Akehurst, J. Poxon, and L. Ash

Title: Development and Field Trial of a Driver Assistance System to Encourage Eco-Driving in Light Commercial Vehicle Fleets

Journal: IEEE Transactions on Intelligent Transportation Systems, vol. 14, iss 2, pp. 796-805

Publication year: 2013

References

[1]   A. M. K. P. Taylor, "Science review of internal combustion engines," *Energy Policy,* vol. 36, pp. 4657-4667, 2008.

[2]   Department of Energy & Climate Change, "2012 UK Greenhouse Gas Emissions, Final Figures," Department of Energy & Climate Change, London, 4 February 2014.

[3]   *Regulation (EC) No 692/2008.*

[4]   *UN/ECE Regulation No 83 Rev.4*, 2011.

[5]   A. Thornton*, et al.*, "Climate Change and Transport Choices: Segmentation Study – Interim Report," 7 December 2010.

[6]   Department for Transport, "Public attitudes towards climate change and the impact on transport: 2011," 26 January 2012.

[7]   The Carbon Trust, "The Carbon Trust Standard Rules v1.3," ed: The Carbon Trust, 2010.

[8]   Department for Transport, "Licensed cars by years since first registration, Great Britain, annually from 2000," VEH0207, Ed., ed. London: Vehicle Licensing Statistics, 2014.

[9]   Department for Transport, "Licensed light goods vehicles by years since first registration, annually from 2000," VEH0407, Ed., ed. London: Vehicle Licensing Statistics, 2014.

[10]   Department for Transport, "Road traffic (vehicle miles) by vehicle type in Great Britain, annual from 1949," TRA0101, Ed., ed. London: Department for Transport, 2014.

[11]   S. Sasaki, "Toyota's newly developed hybrid powertrain," in *Power Semiconductor Devices and ICs, Proceedings of the 10th International Symposium on*, Kyoto, Japan, 1998, pp. 17-22.

[12]   F. R. Salmasi, "Control Strategies for Hybrid Electric Vehicles: Evolution, Classification, Comparison, and Future Trends," *Vehicular Technology, IEEE Transactions on,* vol. 56, pp. 2393-2404, 2007.

[13]   S. G. Wirasingha and A. Emadi, "Classification and Review of Control Strategies for Plug-In Hybrid Electric Vehicles," *Vehicular Technology, IEEE Transactions on,* vol. 60, pp. 111-122, 2011.

[14]   L. A. Zadeh, "Fuzzy Logic and Approximate Reasoning (In Memory of Grigore Moisil)," *Synthese,* vol. 30, pp. 407-428, 1975.

[15]   G. Gerla, "Fuzzy Logic Programming and Fuzzy Control," *Studia Logica: An International Journal for Symbolic Logic,* vol. 79, pp. 231-254, 2005.

[16]   H.-D. Lee and S.-K. Sul, "Fuzzy-logic-based torque control strategy for parallel-type hybrid electric vehicle," *Industrial Electronics, IEEE Transactions on,* vol. 45, pp. 625-632, 1998.

[17]   N. J. Schouten*, et al.*, "Fuzzy logic control for parallel hybrid vehicles," *Control Systems Technology, IEEE Transactions on,* vol. 10, pp. 460-468, 2002.

[18]   R. E. Larson and J. L. Casti, *Principles of Dynamic Programming: Part 1 - Basic Analytic and Computational Methods* vol. 7. New York: Marcel Dekker, 1978.

[19]   C.-C. Lin*, et al.*, "Control System Development for an Advanced-Technology Medium-Duty Hybrid Electric Truck," presented at the SAE 2003 International Truck and Bus Meeting and Exhibition, Fort Worth, Texas, 2003.

[20]   C.-C. Lin*, et al.*, "Power management strategy for a parallel hybrid electric truck," *Control Systems Technology, IEEE Transactions on,* vol. 11, pp. 839-849, November 2003.

[21]   S. Ichikawa*, et al.*, "Novel energy management system for hybrid electric vehicles utilizing car navigation over a commuting route," in *2004 IEEE Intelligent Vehicles Symposium,*, Parma, Italy, 2004, pp. 161-166.

[22]   L. Johannesson*, et al.*, "Assessing the Potential of Predictive Control for Hybrid Vehicle Powertrains Using Stochastic Dynamic Programming," *Intelligent Transportation Systems, IEEE Transactions on,* vol. 8, pp. 71-83, March 2007.

[23]   T. S. Kim*, et al.*, "Model predictive control of velocity and torque split in a parallel hybrid vehicle," in *2009 IEEE International Conference on Systems, Man and Cybernetics*, San Antonio, TX, United States, 2009, pp. 2014-2019.

[24]   C.-C. Lin*, et al.*, "A stochastic control strategy for hybrid electric vehicles," in *2004 American Control Conference*, Boston, MA, United States, 2004, pp. 4710-4715.

[25]   I. Kolmanovsky*, et al.*, "Optimization of powertrain operating policy for feasibility assessment and calibration: stochastic dynamic programming approach," in *2002 American Control Conference*, Anchorage, AK, United States, 2002, pp. 1425-1430 vol.2.

[26]   E. D. Tate Jr.*, et al.*, "Shortest path stochastic control for hybrid electric vehicles," *International Journal of Robust and Nonlinear Control,* vol. 18, pp. 1409-1429, 2007.

[27]   S. J. Moura*, et al.*, "A Stochastic Optimal Control Approach for Power Management in Plug-In Hybrid Electric Vehicles," *Control Systems Technology, IEEE Transactions on,* vol. 19, pp. 545-555, 2011.

[28]   D. F. Opila*, et al.*, "Real-Time Implementation and Hardware Testing of a Hybrid Vehicle Energy Management Controller Based on Stochastic Dynamic Programming," *Journal of Dynamic Systems, Measurement, and Control,* vol. 135, pp. 21002-1 - 21002-11, 2012.

[29]   G. Paganelli*, et al.*, "Control development for a hybrid-electric sport-utility vehicle: Strategy, implementation and field test results," in *2001 American Control Conference*, Arlington, VA, United States, 2001, pp. 5064-5069.

[30]   G. Paganelli*, et al.*, "Optimizing Control Strategy for Hybrid Fuel Cell Vehicle," presented at the SAE 2002 World Congress, Detroit, Michigan, 2002.

[31]   J. Liu and H. Peng, "Modeling and Control of a Power-Split Hybrid Vehicle," *Control Systems Technology, IEEE Transactions on,* vol. 16, pp. 1242-1251, November 2008 2008.

[32]   J. Scoltock, "Focus hybrids: Temperature curve," *Automotive Engineer,* vol. 37(3), pp. 30-31, April 2012.

[33]   M. Dubarry*, et al.*, "A roadmap to understand battery performance in electric and hybrid vehicle operation," *Journal of Power Sources,* vol. 174, pp. 366-372, 2007.

[34]   G. L. Plett, "Extended Kalman filtering for battery management systems of LiPB-based HEV battery packs: Part 1. Background," *Journal of Power Sources,* vol. 134, pp. 252-261, 2004.

[35]   Z. Li*, et al.*, "Modeling the capacity degradation of $LiFePO_4$/graphite batteries based on stress coupling analysis," *Journal of Power Sources,* vol. 196, pp. 9757-9766, 2011.

[36]   B. Y. Liaw and M. Dubarry, "From driving cycle analysis to understanding battery performance in real-life electric hybrid vehicle operation," *Journal of Power Sources,* vol. 174, pp. 76-88, 2007.

[37]   C. Masjosthusmann*, et al.*, "A vehicle energy management system for a Battery Electric Vehicle," in *2012 IEEE Vehicle Power and Propulsion Conference*, Seoul, 2012, pp. 339-344.

[38]   C. Masjosthusmann*, et al.*, "Smoothing the battery current in battery electric vehicles using a quantitative feedback theory designed MIMO-controller," in *2013 IET Hybrid and Electric Vehicles Conference*, London, 2013, pp. 1-6.

[39]   Z. Amjadi and S. S. Williamson, "Power-Electronics-Based Solutions for Plug-in Hybrid Electric Vehicle Energy Storage and Management Systems," *Industrial Electronics, IEEE Transactions on,* vol. 57, pp. 608-616, 2010.

[40]   S. M. Lukic*, et al.*, "Energy Storage Systems for Automotive Applications," *Industrial Electronics, IEEE Transactions on,* vol. 55, pp. 2258-2267, 2008.

[41]   B. Y. Liaw*, et al.*, "Modeling capacity fade in lithium-ion cells," *Journal of Power Sources,* vol. 140, pp. 157-161, 2005.

[42]   M. Mousavi*, et al.*, "Optimal design of an air-cooling system for a Li-Ion battery pack in Electric Vehicles with a genetic algorithm," in *2011 IEEE Congress on Evolutionary Computation*, New Orleans, LA, 2011, pp. 1848-1855.

[43]   A. Affanni*, et al.*, "Battery choice and management for new-generation electric vehicles," *Industrial Electronics, IEEE Transactions on,* vol. 52, pp. 1343-1349, 2005.

[44]   M. Dubarry and B. Y. Liaw, "Development of a universal modeling tool for rechargeable lithium batteries," *Journal of Power Sources,* vol. 174, pp. 856-860, 2007.

[45]   L. Lu*, et al.*, "A review on the key issues for lithium-ion battery management in electric vehicles," *Journal of Power Sources,* vol. 226, pp. 272-288, 2013.

[46]   G. Paganelli*, et al.*, "Equivalent consumption minimization strategy for parallel hybrid powertrains," in *IEEE 55th Vehicular Technology Conference, 2002*, Birmingham, AL, United States, 2002, pp. 2076-2081 vol.4.

[47]   S. J. Moura, "Techniques for Battery Health Conscious Power Management via Electrochemical Modeling and Optimal Control," Ph.D, Mechanical Engineering, University of Michigan, University of Michigan, MI, United States, 2011.

[48] S. Ebbesen*, et al.*, "Battery State-of-Health Perceptive Energy Management for Hybrid Electric Vehicles," *Vehicular Technology, IEEE Transactions on,* vol. 61, pp. 2893-2900, 2012.

[49] L. Serrao*, et al.*, "Optimal energy management of hybrid electric vehicles including battery aging," in *American Control Conference (ACC), 2011*, 2011, pp. 2125-2130.

[50] E. Vinot*, et al.*, "Optimal management of electric vehicles with a hybrid storage system," in *2010 IEEE Vehicle Power and Propulsion Conference*, Lille, France, 2010, pp. 1-6.

[51] J. Färnlund and C. Engström, "The representativeness of driving cycles in real-world traffic," Rototest AB, Rönninge (Sweden) Publication 2001:35 E, 2001.

[52] R. Daniel*, et al.*, "Analysis of US and EU Drive Styles to Improve Understanding of Market Usage and the Effects on OBD Monitor IUMPR," presented at the SAE 2009 World Congress, Detroit, MI, United States, 2009.

[53] B. Adornato*, et al.*, "Characterizing naturalistic driving patterns for Plug-in Hybrid Electric Vehicle analysis," in *2009 IEEE Vehicle Power and Propulsion Conference*, Detroit, MI, United States, 2009, pp. 655-660.

[54] European Conference of Ministers of Transport / International Energy Agency, "Making Cars More Fuel Efficient: Technology and Policies for Real Improvements on the Road," in *European Conference of Ministers of Transport*, 2005, pp. 33-35.

[55] M. S. Young*, et al.*, "Safe driving in a green world: A review of driver performance benchmarks and technologies to support 'smart' driving," *Applied Ergonomics,* vol. 42, pp. 533-539, 2011.

[56] B. Beusen*, et al.*, "Using on-board logging devices to study the longer-term impact of an eco-driving course," *Transportation Research Part D: Transport and Environment,* vol. 14, pp. 514-520, 2009.

[57] A. E. af Wåhlberg, "Long-term effects of training in economical driving: Fuel consumption, accidents, driver acceleration behavior and technical feedback," *International Journal of Industrial Ergonomics,* vol. 37, pp. 333-343, 2007.

[58] M. van der Voort*, et al.*, "A prototype fuel-efficiency support tool," *Transportation Research Part C: Emerging Technologies,* vol. 9, pp. 279-296, 2001.

[59]   C. Wu*, et al.*, "A fuel economy optimization system with applications in vehicles with human drivers and autonomous vehicles," *Transportation Research Part D: Transport and Environment,* vol. 16, pp. 515-524, 2011.

[60]   L. Nouveliere*, et al.*, "Energy saving and safe driving assistance system for light vehicles: Experimentation and analysis," in *2012 9th IEEE International Conference on Networking, Sensing and Control,*, Beijing, China, 2012, pp. 346-351.

[61]   C. J. G. van Driel*, et al.*, "Directives for New-Generation Fuel Economy Devices," Centre for Transport Studies of the University of Twente, Enschede, The Netherlands Rep. 2002R-004/VVR003, 2002.

[62]   H. Larsson and E. Ericsson, "The effects of an acceleration advisory tool in vehicles for reduced fuel consumption and emissions," *Transportation Research Part D: Transport and Environment,* vol. 14, pp. 141-146, 2009.

[63]   E. Ericsson, "Variability in urban driving patterns," *Transportation Research Part D: Transport and Environment,* vol. 5, pp. 337-354, 2000.

[64]   J. Tulusan*, et al.*, "Direct or indirect sensor enabled eco-driving feedback: Which preference do corporate car drivers have?," in *2012 3rd International Conference on the Internet of Things*, Wuxi, China, 2012, pp. 39-46.

[65]   E. Ericsson, "Independent driving pattern factors and their influence on fuel-use and exhaust emission factors," *Transportation Research Part D: Transport and Environment,* vol. 6, pp. 325-345, 2001.

[66]   I. Fomunung*, et al.*, "A statistical model for estimating oxides of nitrogen emissions from light duty motor vehicles," *Transportation Research Part D: Transport and Environment,* vol. 4, pp. 333-352, 1999.

[67]   E. K. Nam*, et al.*, "A Comparison of Real-World and Modelled Emissions Under Conditions of Variable Driver Aggressiveness," presented at the 82nd Annual Meeting of the Transportation Research Board, Washington, DC, United States, 2003.

[68]   C. Vagg*, et al.*, "Development of a new method to assess fuel saving using gear shift indicators," *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering,* vol. 226, pp. 1630-1639, 2012.

[69]   H. C. Watson, "Vehicle Driving Patterns and Measurement Methods for Energy and Emissions Assessment," Bureau of Transport Economics, Canberra,1978.

196

[70]   *Code of Federal Regulations, Title 40: Protection of Environment,* §86.128-00, 2013.

[71]   R. D. Wilson, "Advisory Circular No. 72A - Shift Points for Manual Transmission Test Vehicles," US Environmental Protection Agency: Office of Air and Radiation, Washington, DC,1984.

[72]   J. M. Miller, *Propulsion systems for hybrid vehicles.* London: Institution of Electrical Engineers, 2004.

[73]   R. Burke*, et al.*, "Critical evaluation of on-engine fuel consumption measurement," *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering,* vol. 225, pp. 829-844, 2011.

[74]   International Organization for Standardization, "Heavy-duty engines — Measurement of gaseous emissions from raw exhaust gas and of particulate emissions using partial flow dilution systems under transient test conditions," vol. ISO 16183:2002, ed, 2002.

[75]   T. Leroy*, et al.*, "Towards real-time optimal energy management of HEV powertrains using stochastic dynamic programming," in *2012 IEEE Vehicle Power and Propulsion Conference*, Seoul, 2012, pp. 383-388.

[76]   A. Corti*, et al.*, "Vehicle's energy estimation using low frequency speed signal," in *2012 15th International IEEE Conference on Intelligent Transportation Systems*, Anchorage, AK, United States, 2012, pp. 626-631.

[77]   C. J. Brace*, et al.*, "Increasing accuracy and repeatability of fuel consumption measurement in chassis dynamometer testing," *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering,* vol. 223, pp. 1163-1177, 2009.

[78]   J. Norris*, et al.*, "Assessing the Efficacy of Gear Shift Indicators," AEA / Millbrook, Oxford, Rep. ED49820, 2011.

[79]   C. Bingham*, et al.*, "Impact of driving characteristics on electric vehicle energy consumption and range," *IET Intelligent Transport Systems,* vol. 6, pp. 29-35, 2012.

[80]   C. Musardo*, et al.*, "A-ECMS: An Adaptive Algorithm for Hybrid Electric Vehicle Energy Management," in *2005 44th IEEE Conference on Decision and Control and 2005 European Control Conference*, Seville, Spain, 2005, pp. 1816-1823.

[81]   T. Wada*, et al.*, "Proposal of an eco-driving assist system adaptive to driver's skill," in *2011 14th International IEEE Conference on Intelligent Transportation Systems*, 2011, pp. 1880-1885.