

Innovative Powertrain Control System for a Premium Vehicle

Thomas LEVERMORE

A thesis submitted in partial fulfilment of the requirements for the degree
of
Doctor of Philosophy

Department of Mechanical and Automotive Engineering
Kingston University
June 2017

Abstract

In order to meet increasingly strict regulations on vehicle emissions, manufacturers are seeking ways to produce vehicles that emit less pollution and consume less fuel. Eco-driving is the optimisation of velocity and gear selection in existing vehicles to reduce fuel consumption and such reductions can be made at relatively low development costs compared to powertrain modification. However, the driving experience of a premium vehicle could be compromised if the vehicle behaviour differs from that which is expected by the driver and the acceptance of such fuel saving measures may be diminished. Therefore, in order to maintain the driving experience the contribution of this work is the development and implementation of an optimal control algorithm based on Dynamic Programming which optimises, in real time, the vehicle velocity and gear selection based on a vehicle and upcoming road model while consideration is given to objective measures of driveability. The algorithm is deployed on a Raspberry Pi miniature computer with connection to the vehicle data network. Fuel savings and time savings are identified with the optimisation algorithm both with and without violating constraints on driveability, first in simulation and finally in a real-time, in-vehicle eco-driving feedback system. Primarily the application of this system is in internal combustion engine passenger vehicles in both urban and extra-urban road situations, however the approach is deliberately flexible to allow development for other powertrain configurations.

Declaration

I declare that while registered for a research degree at Kingston University, I have not been a registered candidate or enrolled student for another award at any other academic or professional institution. No material contained in the thesis has been used in any other submission for an academic award. All material contained in the thesis is my own original work, and that any references to or use of other sources have been clearly acknowledged within the text.

Acknowledgements

I would like to thank my supervisor Prof. Sahinkaya, whose guidance has been invaluable, and whose positivity was very much appreciated. My thanks also to Dr. Zweiri, who brought out the best in my work with thorough and thoughtful analysis. Thanks to my supervision team through the years, Dr. Wang and particularly Prof. Ordys who was the catalyst for this work and provided excellent support to set it on the right path. To Dr. Kock, who went out of his way to be helpful and whose enthusiasm and friendship were much appreciated.

To my industrial supervisor Ben Neaves, it was a pleasure to work with you and the help provided by yourself and Samee Al-Adham was crucial to the outcome of this work. I am incredibly grateful for the opportunity that was provided by yourself and Jaguar Land Rover.

To Mum, your wisdom continues to guide and your memory inspires me. A special thanks to my family, Dad, for believing in me and Alison, for being a good listener. To Claire, for making me smile, even during the tough days. To all my friends, thank you for the support.

Contents

List of Figures	viii
List of Tables	x
Nomenclature	1
Glossary	4
1 Introduction	6
1.1 Hypotheses	7
1.2 Objectives	7
1.3 Contributions	8
1.3.1 Publications	9
1.4 Methodology	9
1.5 Thesis Overview	10
2 Background	12
2.1 Internal Combustion Engine	13
2.2 Environmental Legislation	14
2.3 Fuel Consumption Reduction	15
2.3.1 Eco-driving	17
2.3.2 Coasting	18
2.4 Optimisation in Automotive Control	19
2.4.1 Heavy Goods Vehicles	21
2.4.2 Hybrid Electric Vehicles	21
2.4.3 Stochastic Dynamic Programming	22
2.4.4 Deterministic Dynamic Programming	24

2.4.5	Genetic Algorithm	27
2.5	Vehicle Modelling	28
2.5.1	Longitudinal Model	29
2.6	Performance and driveability	30
2.7	Electronic Horizon	32
2.8	Summary	33
3	Optimal Control Algorithm	34
3.1	Optimisation Problem Formulation	35
3.1.1	Constraints	36
3.1.2	Cost Function	37
3.1.3	Assumptions	39
3.2	Road Model	41
3.3	Vehicle Model	41
3.3.1	Model Type	41
3.3.2	Vehicle Dynamics	42
3.3.3	Engine Model	45
3.3.4	Fuel Consumption	47
3.3.5	Time	49
3.3.6	Discretisation	49
3.3.7	Assumptions	52
3.4	Forward Dynamic Programming	55
3.4.1	Search Space	56
3.4.2	Forward Process	58
3.4.3	Transition Cost	58
3.4.4	Transition Cost with Gear Change	59
3.4.5	Terminal Cost	60
3.4.6	Coasting	60
3.4.7	Driveability	61
3.5	Summary	62
4	Implementation	63
4.1	Programming Language	64
4.2	Structure	65

4.3	Road Data	65
4.3.1	Road Data Collection	66
4.3.2	Bespoke Road Data Logger	66
4.3.3	Road Data Processing	67
4.4	Algorithm Development	70
4.4.1	Input arguments	70
4.4.2	Road and Vehicle Initialisation	71
4.4.3	Memory Allocation	72
4.4.4	Search Space Initialisation	73
4.4.5	Cost Calculation	74
4.4.6	Terminal Cost Calculation	74
4.4.7	Optimal Path Search	76
4.5	Data Acquisition	76
4.5.1	Controller Area Network	76
4.5.2	Vehicle Position	79
4.5.3	Real Time Traffic Data	79
4.6	Algorithm Implementation	80
4.7	User Interface	81
4.8	Algorithm and Model Configuration	82
4.8.1	Model Configuration	82
4.8.2	Algorithm Sensitivity Analysis	88
4.9	Algorithm Performance	94
4.9.1	Genetic Algorithm	95
4.9.2	Genetic Algorithm Parameters	96
4.10	Summary	99
5	Results	100
5.1	Vehicle Model Validation	100
5.2	Artificial Road Profiles	104
5.3	Real Driving Data	110
5.4	Algorithm Performance	110
5.4.1	Fixed Velocity	112
5.4.2	Speed Limit Following	115
5.4.3	Normalisation	117

5.4.4	Computation Performance	117
5.4.5	Driveability Consideration	119
5.4.6	Real Driver Comparison	122
5.5	Genetic Algorithm Comparison	123
5.6	In-vehicle implementation	125
5.7	Summary	131
6	Conclusion	132
6.1	Hypothesis (a)	132
6.2	Hypothesis (b)	133
6.3	Hypothesis (c)	134
7	Further Work	136
7.1	Commercial eHorizon System Integration	137
7.2	Vehicle Communication Systems Integration	137
7.3	Eco-guidance and eco-cruise control impact	138
7.4	Alternative Powertrain Vehicles	138
7.5	Emissions Model inclusion in Vehicle Control Optimisation	139
	References	141
	Appendices	162
	A Vehicle Data	163
	B Project Plan	165
	C Expanded Results Data	168

List of Figures

2.1	Vehicle model backward-facing calculation sequence	29
2.2	Vehicle model forward-facing calculation sequence	30
3.1	Fuel and time values varying due to fuel/time weighting . . .	40
3.2	Road gradient specification	41
3.3	Vehicle forces	43
3.4	Resistive force variation	45
3.5	Brake Specific Fuel Consumption map	48
3.6	Euler forward and backward discretisation methods comparison.	50
3.7	Dynamic Programming search space with transitions	56
3.8	Transition cost example	59
4.1	System functional decomposition	65
4.2	GPS signal error	67
4.3	Road Data Processing	69
4.4	Dynamic Programming Structure	70
4.5	Cost calculation initial step	75
4.6	GPS and CAN interface board	78
4.7	CAN Bus network simulation	78
4.8	Driver feedback User Interface	83
4.9	Vehicle model sensitivity	84
4.10	Inertial equivalent mass sensitivity	85
4.11	Aerodynamic variation and fuel consumption	87
4.12	Calculation time as a function of horizon length	89
4.13	Total cost difference variation due to calculation frequency . .	91
4.14	Cost of thirty step horizon with varying calculation frequency	92

4.15	Computational load and velocity discretisation interval	92
4.16	Cost variation with velocity discretisation interval	93
4.17	Genetic representation	95
4.18	Genetic algorithm evolution	97
5.1	Engine behaviour during gearshifting	101
5.2	Engine RPM calculation error during gearshift	102
5.3	Engine coolant temperature	103
5.4	Real road route 1 and 2 maps	104
5.5	Real road route 3 and 4 maps	105
5.6	Fuel consumption for four road tests	105
5.7	Artificial road profiles	107
5.8	Artificial roads - Fixed speed results	108
5.9	Artificial roads - Algorithm comparison	109
5.10	Real road gradient distribution	111
5.11	Real road speed distribution	111
5.12	Optimisation results roads A3-A5	113
5.13	Road R3 speed limit and curvature view	123
5.14	Genetic algorithm comparison	125
5.15	Test setup	126
5.16	Test route map	127
5.17	Test drive and optimal velocity profiles	128
5.18	Test drive and optimal velocity profiles considering traffic . .	130
5.19	Test drive and optimal gear profile	130
A.1	Brake Specific Fuel Consumption Map	164
A.2	Engine Maximum Torque Curve	164

List of Tables

2.1	CO ₂ Emissions Standards Comparison	15
2.2	Global optimisation method comparison	23
2.3	Dynamic Programming velocity optimisation applications . .	26
4.1	Hardware Comparison	64
4.2	Algorithm input arguments	71
4.3	Vehicle constants	72
4.4	Search space variables	73
4.5	Cost variable structure	76
4.6	Vehicle data from CAN bus	79
5.1	Vehicle model fuel consumption validation	106
5.2	Fixed velocity and unconstrained DP results roads A1-A5 . .	114
5.3	Fixed velocity and unconstrained DP costs on R1-R4	115
5.4	Speed limit following and DP costs on R1-R4	116
5.5	Speed limit following and normalised DP costs on R1-R4 . . .	117
5.6	Computation time for real road	119
5.7	Fixed velocity and constrained DP results roads A1-A5 . . .	120
5.8	Fixed velocity and constrained DP results real roads	121
5.9	Real driver compared to DP results	122
5.10	Optimisation algorithm settings	124
5.11	Test drive section results	129
A.1	Vehicle Model Data	163

C.1 Fixed velocity and unconstrained DP fuel, time and cost on
R1-R4 169

C.2 Fixed velocity and constrained DP fuel, time and cost on R1-R4170

Nomenclature

α	Road angle of inclination (rad)
\dot{m}	Mass fuel flow (g s^{-1})
ω_e	Engine rotational speed (rpm)
ω_w	Wheel rotational speed (rad s^{-1})
ρ_a	Density of ambient air (kg m^{-3})
τ_{eb}	Engine Brake torque (N m)
τ_w	Torque at wheels (N m)
A_f	Frontal area of vehicle (m^2)
c_d	Coefficient of drag (-)
c_r	Coefficient of rolling friction (-)
F_d	Combined drag forces (N)
F_a	Force due to aerodynamics (N)
F_d	Force due to additional disturbances (N)
F_g	Force due to gravity (N)
F_r	Force due to rolling resistance (N)
g	Gear (-)
g_k	Gear selection at step k (-)

g_{max}	Highest gear (-)
g_{min}	Lowest gear (-)
h	Elevation change (m)
J_c	Cost associated with comfort (-)
J_f	Cost associated with fuel use (-)
J_t	Cost associated with journey time (-)
m_v	Vehicle mass including inertial equivalent mass (kg)
N_g	Number of gear intervals (-)
N_s	Number of distance intervals (-)
N_v	Number of velocity intervals (-)
p	Absolute pressure (Pa)
R_{dr}	Driveshaft ratio (-)
R_{spec}	Specific gas constant ($\text{J kg}^{-1} \text{K}^{-1}$)
R_{tr}	Transmission ratio (-)
r_w	Wheel radius (m)
s	Distance on path (m)
s_0	Initial distance (m)
s_k	Distance at step k (m)
s_{max}	Maximum distance (m)
sfc	Specific fuel consumption ($\text{g kW}^{-1} \text{h}^{-1}$)
T	Absolute temperature (K)
t	Time (s)
u_g	Gear Control Input (-)

u_v	Velocity Control Input (m s^{-1})
v	Velocity of vehicle (m s^{-1})
v_k	Velocity at step k (m s^{-1})
v_{max}	Maximum velocity (m s^{-1})
v_{min}	Minimum velocity (m s^{-1})
x_k	State Vector at step k (-)

Glossary

ADASIS	Advanced Driver Assistance System Interface Specification
BMEP	Brake Mean Effective Pressure
BSFC	Brake Specific Fuel Consumption
CAFE	Corporate Average Fuel Consumption
CAN	Controller Area Network
DEAP	Distributed Evolutionary Algorithms in Python
DEM	Digital Elevation Model
DP	Dynamic Programming
ECMS	Equivalent Consumption Minimization Strategies
EPA	Environmental Protection Agency
GA	Genetic Algorithm
GHG	Greenhouse Gas
GPS	Global Positioning System
GSI	Gear Shift Indicator
HEV	Hybrid Electric Vehicle
HGV	Heavy Goods Vehicle

ICE	Internal Combustion Engine
IVS	In-Vehicle System
MISRA	Motor Industry Software Reliability Association
NEDC	New European Driving Cycle
NVH	Noise Vibration and Harshness
OEM	Original Equipment Manufacturers
SUV	Sports Utility Vehicle
V2V	Vehicle to Vehicle

Chapter 1

Introduction

The increasing level of CO₂ in the Earth's atmosphere is known to be one of the leading causes of global warming [1]. Combustion of diesel automotive fuel produces CO₂ along with H₂O, CO, H₂, O₂, NO_x, N₂, and unburned hydrocarbons [2]. In order to reduce the impact of passenger vehicle CO₂ emissions on global warming, regulations have been put in place by governments across the globe [3] and are increasingly being tightened. Similarly, local air pollution due to passenger vehicle emissions is an area increasingly being regulated, particularly the emissions of CO, NO_x and particulate matter (PM) [4]. The health effects of such local air pollution are beginning to be understood by governments [5]. In addition to regulation of emissions, there is an economic incentive to reduce fuel consumption in passenger vehicles. With finite resources for hydrocarbon fuels, the cost to vehicle owners has risen and fuel economy is now one of the biggest selling points for new vehicles, as seen in [6] between 2002 and 2007 when US gasoline prices increased from \$1.75 to \$2.86 per gallon and US market share for large sports utility vehicle (SUV) decreased from 18.3% to 12% as this class of vehicle represents some of the least fuel efficient vehicles. Depending on the class of vehicle, the importance placed by purchasers on fuel economy differs but it is clear that it plays a role in the majority of vehicle purchasing decisions and in customer satisfaction during ownership.

In the executive, luxury and SUV segments of the vehicle market the minimum absolute fuel consumption that is achievable is restricted by the

requirements for cabin space, engine power and overall quality of components which all add to the weight of the vehicle. The above mentioned market segments will be grouped under the label “premium” for the purposes of this thesis.

The aim of this project is to study the problem of real world fuel economy and propose, implement and test a method for improving fuel economy while maintaining the driving experience of a premium vehicle. The key questions for the project are then :

- what affects driving experience in relation to longitudinal vehicle behaviour?
- what fuel saving measures are possible without compromising the driving experience?
- how can these measures be controlled to maximise the benefits within the constraints of maintaining driving experience?

1.1 Hypotheses

To investigate this problem, a number of hypotheses are considered in this work as follows :

- (a) Fuel savings can be made by utilising optimal control methods to control vehicle speed and gear selection in real-time, based on instantaneous vehicle and road data.
- (b) Driving experience, in relation to a vehicle’s longitudinal performance, can be quantified and applied to an optimisation algorithm.
- (c) Fuel savings can be made as above, without compromising the driving experience and this can be verified across a range of real driving data.

1.2 Objectives

To investigate the hypotheses stated above the following objectives are to be achieved :

- Review literature to identify fuel savings possible by optimal control of vehicle velocity.
- Develop an optimisation algorithm that uses an accurate and efficient passenger vehicle model to optimise velocity and gear selection for a given journey, in real-time.
- Develop an objective metric for longitudinal driveability of a vehicle based on the correlation of subjective assessments of driveability and quantifiable vehicle data.
- Incorporate driveability metrics as constraints on the optimisation algorithm.
- Develop a catalogue of real world driving scenarios using real data that encompass a range of roads and driving characteristics to test the algorithm.
- Implement real-time optimisation algorithm on hardware and deploy in a production vehicle.

1.3 Contributions

After a thorough literature review and ongoing study of relevant new publications the contributions of this project to the field are as follows :

- Development of a dynamic programming algorithm to optimise, in real-time, vehicle velocity and gear selection to minimise real world fuel consumption and journey time by utilising road and traffic information, applicable for all legal speed limit roads.
- Production of a sensitivity analysis of vehicle model and algorithm variables quantifying the influence on optimisation results.
- Integration of objective driveability constraints into the optimisation strategy to balance fuel savings with ensuring driver satisfaction.
- Deployment of real-time vehicle velocity and gear optimisation algorithm on widely available hardware.

1.3.1 Publications

In the course of this research the following articles were produced for conference proceedings and journal publication :

- Levermore, T., Ordys, A., and Deng, J. “A review of driver modelling”, 2014 UKACC International Conference on Control (CONTROL). IEEE, 2014.
 - T. Levermore - Literature research, manuscript production
 - A. Ordys - Conceptual advice, manuscript reviewing
 - J. Deng - Manuscript reviewing
- Levermore, T., Ordys, A., and Deng, J. “A review of fuel efficient control strategies in the automotive industry”, International Conference on Modern Auto Technology and Services (ICMATS), 2014
 - T. Levermore - Literature research, manuscript production
 - A. Ordys, J. Deng - Conceptual advice, manuscript reviewing
- Levermore, T., Sahinkaya, M. N., Zweiri, Y., and Neaves, B. “Real-Time Velocity Optimization to Minimize Energy Use in Passenger Vehicles”, Energies 2017 [7].
 - T. Levermore - Literature research, manuscript production, algorithm development, in-vehicle system development, testing and deployment.
 - M. N. Sahinkaya, Y. Zweiri - Supervision, manuscript reviewing
 - B. Neaves - Supply of vehicle and test driver

1.4 Methodology

In order to test the hypotheses noted previously, an optimal control method is to be identified that can be made suitable for real time implementation. A dynamic programming algorithm is to be developed and in conjunction with a control oriented vehicle model is to be used to identify fuel savings that can be made by following an optimal velocity and gear selection profile. Initially

simple artificial road profiles such as fixed gradients are to be considered to verify the behaviour of the algorithm. Real road data is then to be used to assess benefits of algorithm under realistic conditions. Road profile data is to be extracted from existing sources as well as from journeys using onboard instrumentation to measure position and road elevation. The assessment of fuel saving is to be made in comparison to a number of baseline scenarios including maintaining a fixed speed, speed limit following and real driving data.

Data is to be recorded from a specific vehicle type during a range of journeys that covers motorways and main roads as well varying road layout and traffic conditions. The vehicle behaviour from these journeys is to be examined in conjunction with driveability research, in order that boundaries of acceptable driving experience can be established. Once quantified these boundaries are to be incorporated into the optimal control algorithm to ensure driveability is maintained. The results from the initial optimisation algorithm are to be compared to that of the optimisation which is bounded by driveability constraints, to test hypothesis (c).

1.5 Thesis Overview

To ensure ease of navigation a brief overview of the thesis structure is provided here. A literature review of research associated with the topic of this thesis is presented in chapter 2. An introduction to the pertinent economic and environmental issues relating to petroleum fuelled internal combustion engine vehicles is followed by details of the internal combustion engine and associated emissions. Legislation and other external pressures influencing the automotive industry are then summarised followed by an assessment of the factors that impact fuel consumption. The relevant publications in the field of optimisation and vehicle modelling are then covered to complete the background work related to this thesis.

Development of the optimisation method used in this thesis is the focus of chapter 3, with the application of Dynamic Programming (DP) described. The development and implementation of a DP algorithm in both a simulation environment and deployed in a vehicle is described in chapter 4. In the

final section of chapter 4 the focus shifts to applying the DP algorithm in a number of simulation scenarios and comparing the results with an alternative optimisation method that while only suitable for offline operation will be able to provide a performance baseline to assess the performance of the DP algorithm against.

The results of the DP algorithm implementation are discussed in chapter 5. The vehicle model using the algorithm is validated against real vehicle data to ensure the accuracy of results and then a sensitivity analysis is presented to identify how the algorithm is influenced by different variables and scenarios. Finally, real driving data is used to quantify the effect of implementing the DP algorithm in real journey scenarios. In chapter 6 the conclusions of the work are provided with reference to the initial hypotheses and findings from the project. Further work is considered in chapter 7 that would lead on from this project.

Note

This project is funded by Jaguar Land Rover and the knowledge and resources of the Diesel Powertrain Research department as well as other relevant departments were a vital part of the project.

Chapter 2

Background

With more than 900 million passenger vehicles currently in use around the world [8] and demand for personal transportation rising, particularly in China, the impact of these vehicles on society and the environment is increasing. A significant portion of this impact is due to the use of petroleum based fuels in the internal combustion engines (ICE) that have powered the overwhelming majority of vehicles for more than one hundred years. The finite nature of this fuel source, its importance to a wide variety of industries and its uneven global distribution has led to petroleum being the focus of much attention in fields as diverse as environmental science, economics, politics as well as engineering. Along with the economic benefit of reducing fuel consumption, it also impacts the amount of emissions produced from vehicles. In the United Kingdom 47.1% of new passenger vehicles sales were for diesel fuelled models [9] during the first quarter of 2016. Complete combustion of diesel under ideal conditions with pure oxygen produces carbon dioxide (CO_2), water and heat [10]. In reality combustion in compression ignition ICEs using diesel fuel produces carbon monoxide, unburned hydrocarbons, nitrogen oxides (NO_x), particulate matter amongst a number of other emission components. There are two main issues related to these emissions, the effect of harmful emissions on local air quality and the effect of CO_2 emissions on the atmosphere. Local air pollution consisting of particulate matter less than $2.5\ \mu\text{m}$ (PM2.5) and NO_2 has been linked to 436,000 and 68,000 premature deaths, respectively in 2013 across the EU-28

group of countries [5]. Diesel fuelled vehicles contribute to both of these types of pollution as well as particulate matter in the range $2.5\ \mu\text{m}$ to $10\ \mu\text{m}$ (PM10). Road transport accounted for 14% and 13% of PM10 and PM2.5 emissions, respectively, in the UK in 2015 [11].

The increasing levels of CO_2 in the atmosphere have been identified as a leading cause of global warming [1] due to the absorption of infra-red energy which is partially radiated back to warm the Earth but would otherwise have escaped the atmosphere. The contribution of the transport sector globally to CO_2 emissions is estimated to be 23% of which road transport was responsible for three quarters in 2013 [12]. In order to combat this, legislation to limit vehicle CO_2 emissions is in place in the nations that make up the majority of passenger vehicle sales [13–15]. In Europe, for example, the average of all a manufacturer’s registered new vehicle CO_2 emissions was required to be less than $130\ \text{g km}^{-1}$ by 2015 (with adjustments for heavier vehicles) and the target is to be reduced to $95\ \text{g km}^{-1}$ for 2020 [14]. Where a manufacturer’s average exceeds the limit a financial penalty as high as $\text{€}95$ per g km^{-1} per vehicle sold is imposed which gives manufacturers a great incentive to minimise CO_2 emissions. As the fuel consumption in a compression ignition engine is linked to CO_2 , [16]. This link is crucial to reducing vehicle CO_2 emissions as fuel consumption can be a deciding factor for consumer purchasing decisions and is a major on-going cost of vehicle ownership. Whether drivers are aiming to be more environmentally friendly or simply more frugal the result is a lower fuel consumption and CO_2 output.

2.1 Internal Combustion Engine

To reduce fuel consumption and or emissions from a vehicle with a conventional powertrain it is important to understand the mechanisms that drive fuel use and emission formation. For the purposes of this project we are focussed solely on vehicles powered by compression-ignition engines fuelled by diesel, however, as discussed later, the project lends itself to being extended to petrol, hybrid or electric vehicles. The compression-ignition engine is considered to have a greater fuel conversion efficiency relative to an equivalent spark-ignition engine [2] so is the more suitable combustion process for

a vehicle where fuel economy is a primary concern.

Efficiency

The fundamental purpose of the engine in a vehicle is to provide propulsion, however not all the energy available in the fuel can be utilised for this purpose. The power as measured at the engine output shaft is known as the brake power (bp) and the power developed by fuel combustion in the engine cylinders is known as indicated power (ip). The difference between bp and ip is known as friction power (fp) and is a general term to cover all losses within the engine. Brake Specific Fuel Consumption (BSFC) is a measure of fuel consumed per unit of brake power developed. This is used to compare efficiency of internal combustion engines of different sizes and typically the units are $\text{g kW}^{-1} \text{h}^{-1}$.

2.2 Environmental Legislation

In order to meet internationally agreed Greenhouse Gas (GHG) and global warming targets and with the contribution of passenger transportation to GHG outputs, countries around the world have implemented national legislation to limit the GHG emissions of new vehicles. In the European Union regulations are in place [14] that have limited manufacturer fleet CO_2 averages to less than 130 g km^{-1} in 2015 and 95 g km^{-1} in 2021. This figure is based on the New European Driving Cycle (NEDC) tests. In the United States, the Environmental Protection Agency (EPA) has established the Corporate Average Fuel Economy (CAFE) standards that require a fleet average across light duty vehicles (less than 4500 kg) of 163 g/mile CO_2 by 2025 [17]. Japan has implemented standards of manufacturer fleet CO_2 averages of less than 115 g km^{-1} by 2020 [18]. These standards are presented in comparable units in Table 2.1 as reproduced from [18]. With the widespread implementation of such standards manufacturers have legal and financial incentives for reducing their vehicle fleet fuel consumption regardless of their global sales distribution.

	CO ₂ (g km ⁻¹)	
Standard	NEDC	CAFE
EU 2021	95	90
US 2025	120	100.9
Japan 2020	105	100
China 2021	117	115.9

Table 2.1: CO₂ Emissions Standards Comparison [18]

2.3 Fuel Consumption Reduction

There are a number of approaches manufacturers can take to reduce fuel consumption from first principles, reducing the vehicle mass [19], aerodynamic drag [20] and rolling resistance. More innovative concepts to reduce fuel consumption have been implemented in the last decade such as stop-start systems [21] that turn off the engine when the vehicle is stationary, reducing the losses associated with idling which are reported to account for 40 billion litres of fuel used in the United States [22]. Another concept that continues to change the automotive sector is hybridisation of the powertrain with electric motors [23] to adjust the load on the ICE such that it is used more efficiently or not at all. The approaches to fuel consumption reduction can be grouped in two types, the first as described above involving changes to the mechanical design of the vehicle. The second type of approach involves modifying driver behaviour by such means as navigation systems that select the most economical route, car sharing and training or guidance to improve the style of driving [24].

While mechanical design changes such as mass reduction and hybridisation can lead to significant energy savings [23] they can only be introduced as part of the normal cycle of vehicle replacement which was every 16 years and increasing in the United States in 1990 [25] and also they require substantial investment by the manufacturer and often by the consumer as well. Behaviour modification, on the other hand, can have an immediate impact on the existing vehicle fleet as the vehicles themselves are not changing but rather the way in which they are used. In [26] such behavioural modifica-

tion involve three types of decision, firstly, a strategic choice of what vehicle type to use or purchase, secondly, route selection and finally, operational decisions that make up driving behaviour. The first decision is out of the scope of this project but the other two will be considered in more detail.

Behaviour modification by the use of map data to identify the most economical route for a journey is known as eco-routing. In [27] an average fuel reduction of 8.2% could be made by choosing a more economic route for a given journey and for 46% of the journeys in the study such an alternative route was available. An economic route in this case was considered to minimise the occurrence of junctions controlled by traffic lights, traffic calming measure and high speed limits. The eco-routing concept has grown in use since the wide spread deployment of Global Positioning System (GPS) satellite navigation systems in vehicles, with more than 25 million In-Vehicle System (IVS) sold in 2013 [28]. The first developments in this field required an accurate vehicle fuel consumption model combined with road information to simulate the fuel used for each possible route and thus allow an informed selection to be made [29]. One road property that is used to make such a selection is the road gradient and its effect is investigated in [30]. It was found that an increase in fuel consumption of 18% is seen when the road gradient increases from 0% to 1% for a velocity of 75 km h^{-1} . These figures were based on a light duty vehicle of average weight according to US sales data. While quantifying the difference in fuel consumption at different velocities and gradients is an important finding the reality of driving is that there are always constraints on the vehicle velocity and gradients are a fundamental feature of road networks. An obvious constraint is the presence of legal speed limits assigned to the overwhelming majority of roads, however a more dynamic restriction on velocity is the presence of other vehicles. In [27] the availability of traffic data is noted as a potential aid to fuel-saving in high traffic density areas. Historic traffic data, real time traffic data and traffic flow metrics can be used for this purpose.

While the most economical route for a journey can be calculated using these methods there are still fuel savings that can be made during the journey itself in the way the vehicle is driven.

2.3.1 Eco-driving

It has been shown that fuel efficiency is maximised between 60 km h^{-1} and 80 km h^{-1} [31] for vehicles produced in the mid-nineties however this is not optimal when considering journey time and traffic speeds on roads in Europe of up to 130 km h^{-1} . Driving efficiently therefore requires a more complex solution than simply slower driving and conveying this information in an understandable format to drivers requires training. Driver training with the goal of encouraging economical driving habits has been investigated in [24] with the finding that 10% fuel savings were achieved across a number of programmes. Similar studies show 15% fuel reduction [32] and [33] finds a 5.8% reduction in fuel consumption for 10 drivers following an eco-driving training course. A variety of cars were used in the test and the change in fuel consumption varied by driver as 20% or participants achieved no fuel saving. Pollutant emissions are considered in addition to fuel consumption in the investigation of eco-driving in [34].

Aside from training, it is noted in [35] that there can be three opportunities to encourage eco-driving - before, during and after a journey. An overview of manufacturer eco-driving systems is presented in [36] covering during and after a journey. The driving behaviour during the journey will be the focus of this work.

The work presented in [37] focuses on modifying driver behaviour during driving solely in relation to acceleration by modifying the accelerator pedals of four cars to increase resistance when the driver attempts to accelerate forcefully. The vehicles were used for delivering post along a number of set routes and only one of these routes showed improvement in fuel consumption with the modified pedal, leading the author to conclude that rate of acceleration is not the sole factor in fuel consumption reduction. Despite this, Nissan have implemented a similar feature in Japan, as investigated in [38].

As driving styles vary and sensitivity of fuel efficiency to driving style depends on vehicle power-to-weight ratio [39], there is not a universal solution to efficient driving. However following general rules were presented in an EcoWILL publication [40] for providers of driver training

1. Anticipate Traffic flow
2. Maintain a steady speed at low engine speed
3. Shift up early

The ability to anticipate traffic flow depends on the road layout and visibility amongst other factors and therefore is unreasonable to be expected consistently of a driver, however as noted earlier the use of real time traffic data is becoming more prevalent in modern vehicles. Relying on a connected system to anticipate traffic for the driver allows more economic driving to be undertaken with a reduced mental workload on the driver. Currently in Europe it is a requirement for new manual transmission passenger vehicles to include a gear shift indicator (GSI) [41] to aid the driver in applying the last two rules, however this indication is based on current engine speed and makes no consideration for road conditions or upcoming situations. Despite this drawback it is calculated in [42] that CO₂ reductions of between 3.5% and 4.5% could be achieved with GSI on the NEDC. The NEDC is the test cycle used to assess the fuel economy and emissions of passenger vehicles being approved for use in Europe. Both of the last two rules can be followed more easily with knowledge about the road situation, including speed limits and road gradient which again can be made available from a connected system. The rules also lend themselves to a degree of automation, for example, adaptive cruise control and automatic transmission. If the ability to control speed and gear selection is made available to an on board controller then the question is raised of how the control setpoints are generated and what is the goal of such setpoints? If the goal was to save fuel it must be balanced with a requirement to reach the journey's destination in a reasonable time, such a balance can be treated as an optimisation problem.

2.3.2 Coasting

As noted previously, a significant portion of energy loss in an internal combustion engine is due to overcoming friction in the engine. During driving when the transmission is engaged and couples a rotating axle with the engine output shaft, if no fuel is injected then the energy required to rotate

the engine must come from the kinetic energy of the vehicle, thus reducing the velocity; this is known as engine braking. If it is not desirable to reduce speed then it is possible to disengage the transmission and remove the effect of engine braking, allowing the vehicle to coast with only gradient, rolling resistance and drag to overcome. A number of manufacturers have developed such coasting abilities [43]. A decision has to be made with the engine disengaged from kinetic energy of the vehicle as to whether to supply fuel to run the engine at idle speed or to turn off the engine completely and restart when required. Historically an array of vital components were powered directly from the engine including belt driven power steering hydraulic fluid supply pumps and so turning off the engine was not a possibility. With the increasing electrification of vehicle components such as power steering this issue can be avoided and with the deployment of stop/start technology the reliability of restarting the engine is greatly increased. In [44] a compromise is proposed that uses the torque converter in an automatic transmission to partially decouple the engine during coasting but allow sufficient torque to be supplied from the wheels to maintain the engine at idle speed but with fewer losses than if the engine was running at the higher speed when completely coupled. The decision of when to initiate a coasting operation is not one that can be added to the driver's mental workload and so along with the driving behaviours above is a decision that is best taken by a system designed to optimise the vehicle behaviour to reduce fuel consumption. Incorporating coasting strategies into the framework of eco-driving described in subsection 2.3.1 further highlights the requirement for optimisation methods to be applied to the control of the vehicle velocity and gear selection.

2.4 Optimisation in Automotive Control

Optimal control is the selection of system parameters such that a quantity indicating performance is either minimised or maximised [45]. Three types of optimisation are noted in [46] as relevant to the task of designing vehicle propulsion systems. In a hierarchical order these are

- selection of the best powertrain component combination, such as a conventional, hybrid electric or fully electric powertrain

- optimisation of the component parameters for the chosen powertrain structure, such as the rating of the electric motors in a fully electric vehicle [47]
- optimisation of the supervisory control system to best utilise the components in the powertrain

It is this third level of optimisation that will be the focus of the following section.

While the topic of vehicle fuel consumption has been detailed in the previous section and has a long history, the idea of using optimal control for such an application is more recent. In 1977, Schwarzkopf et al. [48] published an algorithm based on the Pontryagin maximum principle [49] to minimise fuel consumption over a number of artificial road profiles of fixed up and downslopes of 10%. This method requires that the engine fuel consumption performance can be described by a polynomial function which limits the quality of results as the complexity of modern engine efficiency is lost in this approximation. The use of Pontryagin's maximum principle for fuel optimal velocity control is extended in [50] and compared with maintaining a fixed velocity.

Another early and influential publication was produced by Hooker et al. [51] in 1983. This work investigated control of vehicle speed for fuel economy in three scenarios; accelerating to a cruising speed, driving with an average speed over roads with gradients and driving between stop signs. The technique used was based on the seminal work of Bellman [52] on Dynamic Programming (DP) and presented two algorithms, one controlling acceleration and another that controls both acceleration and gearshifting. In [53] this work was extended by the same author to consider fifteen different vehicle models and it is shown that the optimal control for one vehicle may differ considerably from that of another vehicle on the same road. To identify this previously would have required multiple test runs with each vehicle under identical conditions which presents a number of challenges. The benefit of the vehicle simulation proposed in the above mentioned work is clear when considered against this previous testing procedure. The conclusions drawn from this work were that the potential for fuel saving by economical driving

was much greater with gradients and stop / start traffic than for cruising at a fixed average speed on a flat road section.

2.4.1 Heavy Goods Vehicles

More recently the focus of such work has moved to Heavy Goods Vehicles (HGVs) as fuel makes up a significant portion of a haulage company's operating costs [54] and the weight of HGVs means that a suboptimal speed over gradients has a much greater impact on fuel consumption than for a passenger vehicle. These are the motivating factors for focussing on HGVs. In [55] the optimal control strategy for a truck traversing a series of simple linear road profiles is considered and analytical solutions are found which can be used to compare with alternative optimisation methods, such as DP. In [56] the non linearity of a modern engine's specific fuel consumption (SFC) map is studied to identify the impact of this on the optimisation results. Two objective functions are evaluated, the first considering a time constraint and the second using a weighting factor between fuel and time. The research on HGVs focuses on highway cruising speeds with a limited window of possible speeds and accelerations, which make up a major portion of HGV journey. In [57] only speeds in the range 79 km h^{-1} to 89 km h^{-1} are considered. This restricts the scope of possible velocity profiles and therefore the complexity of the algorithm however for a passenger vehicle a much larger range is necessary. Model predictive control is applied in [58] to control both cruise control speed and gear selection in a HGV. The length of the horizon over which the prediction takes place varies is proportional to velocity resulting in a fixed time period over which the prediction operates.

2.4.2 Hybrid Electric Vehicles

A related area of research that is more focussed on passenger vehicles has been the application of optimisation methods to Hybrid Electric Vehicles (HEV) control policies where torque can be provided by either an electric motor, an ICE or a combination of both. An overview of this topic can be found in [59] which lists the optimisation methods used by various research groups. A common approach is what is known as equivalent consumption

minimization strategies (ECMS) and was first presented in [60]. Assuming a charge sustaining policy where the battery charge is the same at the end of a journey as at the start, the ECMS considers a deficient battery charge as the equivalent of a future fuel use to recharge the battery either from the engine directly or from regenerative braking to reduce a velocity that was produced by the engine. Similarly an excess battery charge is the equivalent of a future fuel saving as it can be used to reduce the work required of the ICE. By considering both elements of the powertrain in terms of fuel consumption it allows a comparison of different control actions.

This concept is used in [61–63] however the application of ECMS in optimisation differs. In [61] a DP algorithm is considered but deemed unsuitable for real time application due to knowledge of the future driving situation being required. A different approach is taken in [63], where the problem is considered as a single-objective linear optimisation problem, the single objective being minimising the equivalent energy consumption rate. This shows that even where the ECMS concept is common, there is a variety of approaches taken to implementing the concept in optimisation. Various optimisation methods are compared in [64] to highlight their suitability for use in HEV control strategies. A table summarising the conclusions of this publication on global optimisation methods is reproduced in Table 2.2. Two types of DP are considered in the analysis. The first assumes that the upcoming situation is entirely known and is referred to as deterministic DP, whereas for the second type the future situation is not known but probabilities can be applied to a number of possibilities and the optimisation can incorporate these probabilities; this second approach is known as stochastic DP. The only method considered to be suitable for real-time application in Table 2.2 is stochastic DP as the other methods rely on a complete velocity profile to be followed so that the HEV can be controlled optimally.

2.4.3 Stochastic Dynamic Programming

In a real-time application the velocity profile is unknown and there will be uncertainty over future power demands as this will vary with driving style and road situation. One solution as proposed in [65] is to use determinis-

	Computational Load	Robustness	Real-time
Linear Programming	+ +	-	- -
Optimal Control	- -	- -	-
Dynamic Programming	-	-	-
Stochastic DP	- -	-	+
Genetic Algorithm	- -	+	-

Table 2.2: Global optimisation method comparison with respect to computation load, robustness and real-time applicability. For a hybrid vehicle optimisation system and each particular category, suitable and very suitable methods are indicated by + and + +, respectively, and similarly, - and - -, indicate unsuitable and very unsuitable methods.

tic DP for a specific drive cycle and use the results from this to produce a heuristic rule based strategy. It was shown that this strategy produced results that were between 50% and 70% of the performance of the deterministic DP algorithm. The performance of such a system on unseen driving cycles is unpredictable and real driving can differ greatly from standard driving cycles. A solution to this problem is improved driver modelling and prediction. In [66] the driver’s control requests, represented by a power demand, are modelled by a Markov chain. At a specific wheel speed a probability distribution represents the likelihood of the next demanded speed, the distribution is dependent only on the current speed and not on previous speed, this being a property of a Markov chain [67]. The probability distributions are calculated based on data from a number of drive cycles, however the performance of the algorithm in unrelated situations cannot be predicted. This distribution represents a stochastic variable which cannot be predicted with absolute confidence but can be modelled from prior data. Rather than the driver demand being the stochastic variable, in [68] a lead vehicle velocity in a traffic scenario is considered as such. A slightly different approach is applied in [69] where rather than the future velocity being a random variable, a deterministic approach is taken to produce a complete velocity profile for the upcoming road section. A stochastic element is used however for the expected fuel consumption of a given state, due to road

conditions and drivers varying from the prescribed velocity profile.

In addition to varying by journey, driver demands also vary greatly between drivers; they are also hard to predict and can even vary for a single driver at different times. Despite the attention on the driver, a far greater influence on the future velocity profile comes from the road environment, primarily speed limit and local traffic speed. The current and upcoming speed limit requires no prediction, only access to a database and knowledge of the route ahead [70] and while traffic speed is difficult to predict in advance, if it is measured for the route ahead then a prediction can be made with much greater confidence. Such road information may not have been envisaged as being available when the authors of [66] published their work but the current state of the art allows that the first assumption of knowledge of the upcoming situation can be made with a high level of confidence. Another approach then, is to apply this knowledge of the upcoming road situation that will be available in connected vehicles of the near future. Assuming the driver obeys the legal speed limits and conforms with traffic flow speeds, then due to the availability of road data on both there is a limitation on the uncertainty of future driver demands. Alternatively, the task of regulating vehicle speed can be automated. Such a concept has been present in vehicles since the first implementation of a cruise control system and removes the need to predict driver behaviour. This also enables several of the optimisation methods deemed unsuitable for real-time applications in Table 2.2 to be applicable. The removal of the uncertainty caused by driver behaviour provides an opportunity to apply the deterministic implementation of DP.

2.4.4 Deterministic Dynamic Programming

The assumption of automated velocity control allows the stochastic element of the future velocity to be replaced by a velocity profile that can be completely determined by the optimisation algorithm.

This approach is taken in [57] for HGVs with the deterministic DP algorithm producing a fuel saving of 3.5% with no loss in time. It was noted also that the number of gear shift operations reduced by 42% due to anticipation of gradient changes due to road information. In [71] DP is applied to a con-

ventional powertrain vehicle and a three dimensional search space is used with time, velocity and distance as variables. Knowing two of the three variables allows you to calculate the third so this redundancy is removed in a following publication by the same authors [72]. This publication also includes an assessment of the impact of traffic on the optimisation by considering a time margin between of two and four seconds to the vehicle ahead. The effect reported was to reduce the fuel consumption saving from 34% to 28% and 15% for two and four seconds time margin, respectively when compared to unrestricted optimisation. The work in [73] does not consider traffic and gear shifting is according to a fixed schedule but introduces an iterative approach to DP where the algorithm is initially computed with a velocity resolution of 2 mph that reduces the computational load as there are fewer velocity choices than a 1 mph resolution would produce. Once the initial velocity profile is produced a narrow search space is constructed around it which is used for a second iteration of the DP algorithm using a higher resolution. It is highlighted in the publication that the fuel consumption following the profile calculated using 2 mph (0.89 m s^{-1}) is only 0.35% higher than that of the profile calculated using 1 mph (0.45 m s^{-1}) so the benefit is minimal for this second iteration. The horizon over which the DP operates is divided into 250 m sections. This iterative approach is expanded in several publications by Wahl [74], and [75] make use of DP where the algorithms that are applied to the HEV use an iterative approach where initially a coarse search grid is used to identify where the optimal solution lies. Subsequent iterations of the DP algorithm use a finer search space centred on the initial solution to improve the accuracy of the optimal solution. The benefit of using this iterative approach is that the computation time for the coarse grid is significantly lower than the equivalent search area with a finer grid size, but the following iterations can benefit from a finer grid without the associated computational load.

An issue for all the optimisation methods mentioned previously is the domain in which the horizon is considered, the movement of a vehicle can be described in the time domain and also in the spatial domain. The time domain was used in [51] and is often used in adaptive cruise control velocity optimisation such as [76] and particularly in urban scenarios where other

Author	Vehicle	Domain	Horizon	Velocity Interval
Hellström [57]	HGV	Spatial	1500 m 50 m steps	0.2 km h ⁻¹ (0.06 m s ⁻¹)
Wollaeger [73]	Conv	Spatial	250 m steps	1mph (0.45 m s ⁻¹)
Mensing [86]	Conv	Temporal	300 m 5 m steps	0.2 m s ⁻¹
Wahl [74]	HEV	Spatial	3000 m	-
Luu [79]	Conv	Spatial	1000 m 10 m steps	-
Gausemeier [84]	Conv	Spatial	2500 m 50 m steps	1 km h ⁻¹ (0.28 m s ⁻¹)
Dib [87]	EV	Spatial	-	-
Levermore	Conv	Spatial	1500 m 50 m steps	0.1 m s ⁻¹

Table 2.3: Dynamic Programming vehicle velocity optimisation applications

traffic is considered [77]. Optimisation of vehicle speed to synchronise with traffic light timing also often takes place in the time domain [78].

The spatial domain is intrinsically linked to the road topography and so is well suited for vehicle velocity optimisation on stretches of road with varying gradient [57, 74, 79–81]. The drawback of this approach is the inability to consider stationary periods [82], which are likely to occur in urban environments. The decision of what domain to operate in dictates how the final state is specified, in the spatial domain the horizon length and therefore the final distance are fixed however the final time can be unconstrained as in [82, 83] or minimised as in [57, 84]. In the case of [85] a range of final times are acceptable based on traffic light timings. A summary of DP algorithms used to implement vehicle velocity control in a variety of vehicles is presented in Table 2.3 along with the values used in this work.

2.4.5 Genetic Algorithm

The theory of a Genetic Algorithm (GA) was developed in the 1960s and formally presented in [88]. Taking inspiration from the natural world and evolution's "survival of the fittest", a GA aims to mimic this process to achieve an optimum or fittest solution. A GA requires two parts [89] which are similar to the DP algorithm requirements

- a pseudo genetic representation of solutions known as a "genotype" or a "chromosome"
- a function to calculate the "fitness" of solution, equivalent to the cost function used in DP.

A population of chromosomes are assessed using the fitness function and a variety of operations can be undertaken to produce the next generation of chromosomes from the initial population. The conclusion of the GA is decided by a range of criteria, such as a maximum number of populations to produce or reaching a threshold fitness value.

The application of GA to optimal control problems was described in [90] and a simple example problem is presented in which the distance travelled by a push-cart in a given time, minus the effort required to push the cart, is to be maximised. This is similar to the problem of optimal velocity control of a vehicle which has been tackled with DP above, however the GA method is noted as being better suited than DP to problems with a moderate number of dimensions. This GA method is applied in [91, 92] to optimise the energy management in a HEV with the fitness function being the weighted sum of fuel consumption and ICE emissions carbon monoxide, hydrocarbons and oxides of nitrogen. The switching points for turning off the ICE, the amount of torque used to charge the batteries and the upper and lower limits of battery charge are the parameters that make up the chromosome representation to be optimised. This work is developed in [93] with the powertrain parameters expanded to include peak power of ICE, Motor rated power, battery capacity, final drive ratio with the same fitness function. A different strategy is taken in [94] as the total energy consumption on a journey is the single variable in the fitness function and the genetic representation contains a

control command at every discrete distance in the road section. The control command is limited to one of three states, maximum acceleration, maintaining velocity and coasting with no propulsion force. GA is used in [95] to select the topology of a battery electric vehicle to optimise cost, efficiency and performance for a range of greater than 175 km of the NEDC.

The use of GA to optimise vehicle velocity is proposed in [96] albeit with a train used as the vehicle investigated as opposed to a car. The cost function comprised of the electrical energy consumption as well as the time taken for a given section of track of known gradient and speed limit. No mention is made of the computation time for the GA as the results were computed offline, not implemented in a time critical application. A similar approach is taken in [97], with a constraint on computation cut-off time of 30 seconds applied when running on a high performance personal computer. It is noted that despite the similarities in the optimisation problem being solved for either car or train, the GA approach is more commonly applied to the train problem. The regular scheduling of operation for a given train lessens the impact of the long computation times associated with GA and results can be calculated offline which would be unsuitable for the more unpredictable journey requirements of a car.

2.5 Vehicle Modelling

Fundamental to the design of any optimisation method is the accurate modelling of the system that is to be controlled. Vehicle modelling is a field of research that has a number of applications each concerned with different aspects of vehicle behaviour [98]. It is necessary to identify exactly what is required from a vehicle model and when the drawbacks of increasing the level of detail outweigh the benefits. Consider a vehicle moving in three dimensions relative to a road section, movement in the vertical plane does not contribute significantly to fuel consumption other than through the increase of the rolling resistance losses, as discussed later, but is vital for Noise, Vibration and Harshness (NVH) analysis. As noted in [99], to simplify a vehicle model the vertical dynamics are often separated from the model under analysis and can be disregarded. Modelling of vehicle lateral motion is

of particular interest to handling and safety research in lane departure [100] and at the limits of vehicle performance. Under normal driving conditions the lateral motion of a vehicle is not considered to contribute significantly to fuel consumption in [101], however in [102] a formula is presented that longitudinal drag due to cornering increases exponentially with vehicle velocity as well as increasing with slip angle and inverse of curve radius.

2.5.1 Longitudinal Model

There are two main approaches to vehicle longitudinal simulation [103], the first, is known as a backward-facing vehicle model or quasi-static inverse model [104] and in which the entire velocity profile to be simulated is known. From this using a model of the wheel dimensions, the wheel speed can be calculated, followed by driveshaft, transmission, and finally, engine speed. In parallel with this the torque required to follow the velocity profile is calculated for each component in the powertrain. In [105] this concept is presented with a simple formula for calculating torque using acceleration, the square of velocity and lumped coefficients to represent resistive forces such as aerodynamic drag. Using static efficiency tables populated by test data the fuel consumption for a given engine speed and torque can be found. The advantages of the backward-facing approach is the speed of calculation which due to the use of lookup tables can be relatively quick. The drawbacks can be attributed to the inflexible nature of a model based on velocity profiles which cannot account for variation in driving style or driver ability. Dynamic effects such as those resulting from exhaust gas recirculation (EGR) and turbo charger behaviour are also not considered in the model as the efficiency tables are based on static test measurements.

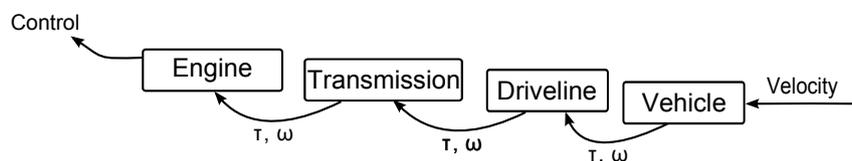


Figure 2.1: Vehicle model backward-facing calculation sequence with torque, τ and rotational speed ω used to link the elements of the model.

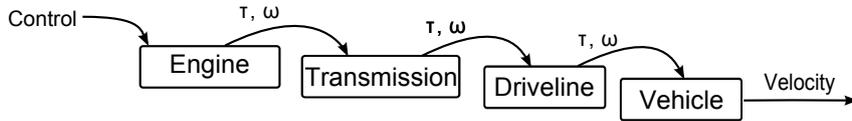


Figure 2.2: Vehicle model forward-facing calculation sequence with torque, τ and rotational speed ω used to link the elements of the model.

The second approach is the forward-facing model where the sequence of calculations works in the opposite direction. At the input to the model is a speed controller for the vehicle, representing either a driver or a cruise control system that provides a control input to the engine model (and brakes) which generates a torque and rotational speed as an input to the transmission model, followed by the driveshaft, the wheels and finally the vehicle speed can be calculated. The forward-facing approach commonly relies on differential equations to describe the behaviour of system component frequencies high enough to capture the dynamics of components such as air path variations. This high frequency leads to an increased computation time compared to the first approach, with benefits for example for pollutant formation simulation that relies on high frequency modelling [46]. Fuel consumption is less influenced by these high frequency variations and so results can be comparable for both approaches.

2.6 Performance and driveability

The perception of the driving experience is very subjective but progress has been made in recent years to quantify this under a title of driveability. This topic can be divided into two areas, first that of NVH dealing with the chassis, suspension and other mechanical design implications. The second area is that of the dynamic behaviour of a vehicle powertrain during various driving scenarios and the impact of this on driveability. Approaches to objectively measure driveability are developed independently in [106] and [107] with the aim of automating vehicle development using software that can provide a driveability rating based on the objective metrics developed. In [107] the

focus is on tip-in and tip-out, where the driver presses the accelerator pedal sharply and similarly releases the pedal. A correlation is shown between the subjective rating of a vehicle performance and acceleration overshoot experienced during a tip-in manoeuvre, which can be measured directly. More detail is provided in [106] considering the derivative of acceleration, known as jerk during the same situation. A neural network approach is proposed to predict driveability assessment based on vehicle and engine measurements. As both publications are from the commercial sector the details are limited.

A more detailed investigation is presented in [108] with assessment of three vehicles in three scenarios, starting from rest, heavy traffic conditions at low speeds and overall performance. The correlation between jerk and performance perception is highlighted as being stronger than that between acceleration and performance perception. This work is expanded in [109] with seven vehicles, five continuously variable transmission vehicle and two with automatic transmissions.

The behaviour of automatic transmissions is considered in respect to driveability in [110] during the development of an hybrid vehicle optimal energy management system. Shifting too frequently is noted as undesirable behaviour as well the varying shift timings. The authors implement a limit on the total number of gear shift occurrences and similarly engine restarts for the optimal control system. In [111] the development of these limits is detailed with a number of factors being approximated by the number of gear shifts and engine starts.

Performance and driveability are both considered in [112] where performance is objectively measured in a number of ways including acceleration times, top speed and gradeability which is the maximum gradient that a vehicle can maintain a given speed. Driveability is objectively defined by vibration and noise measurements as well as tip-in, tip-out response and jerk. In [113] only tip-in, tip-out and gearshift are considered in relationship to powertrain driveability. The perception of driveability is taken one step further in [114] where not just the longitudinal acceleration of the vehicle chassis is considered but also that of the driver's headrest as ultimately the human perception of driveability is what is of interest.

2.7 Electronic Horizon

In order to implement an optimisation algorithm that subscribes to the eco-driving rules noted in subsection 2.3.1 there is a requirement for knowledge of the upcoming road section. This requirement is shared with a number of automotive applications that are in development in the field of vehicle safety such as Adaptive Speed Recommendation [115] based on road information such as speed limit, road curvature, junctions and Adaptive Front-lighting [116] which directs the front headlight beam according to the upcoming road geometry to better illuminate the way. A system to provide information on the upcoming road section is commonly referred to as an electronic horizon.

Due to the number and variety of applications that would benefit from an electronic horizon system it is necessary for all relevant parties to collaborate on the design of such a system. This is the goal of the Advanced Driver Assistance System Interface Specification (ADASIS) forum [117], which brings together major vehicle manufacturers such as BMW, Daimler, Ford and Jaguar Land Rover, mapping companies such as TomTom and Here, as well as automotive original equipment manufacturers (OEMs) such as Bosch and Continental. The specification ADASIS v2 Protocol [118], produced by this forum, dictates the type and format of the signals available from an ADAS Horizon Provider (Av2HP) and thus allows any equipment connected via the Controller Area Network (CAN) bus to request and interpret the data to reconstruct a road model for a particular application. Each application is able to request specific items from the Av2HP and a Horizon Reconstructor (Av2HR) is made specifically for that application using only the relevant data. Of relevance to vehicle speed optimisation, road slope data is listed in the specification as well as the speed limit, which can either be the signed legal speed limit, or variable speed limits depending on the source of information. By adhering to the specification it ensures that any suppliers Av2HP can communicate with any other suppliers Av2HR thus reducing the costs associated with developing new applications and allowing specialists in areas such as navigation to focus solely on their area of interest. For example if the vehicle does not have a predefined destination then all the associated ADAS applications are reliant on the navigation system

to predict a most probable path that the vehicle will take, a problem that can be tackled with a number of approaches each of which if satisfactory will not impact the operation of the ADAS application.

While the ADASIS protocol provides isolation between the two components of the system, the horizon provider and the application, it is important to investigate the quality of the information provided. For instance, intelligent cruise control that relies on speed limit information being up to date and accurate much more than a system that simply informs the driver of speed limit information. Similarly an eco-feedback system that relies on road gradient information to optimise vehicle speed will require much greater accuracy in gradient data than for a conventional navigation system, and much greater again if the application is HGVs [119] rather than passenger vehicles due to the significantly increased vehicle weight. In [120] this is investigated with a commercial ADAS map database compared to a Digital Elevation Model (DEM) of the United States for both elevation and gradient accuracy.

2.8 Summary

This chapter has provided background information to the motivation for this project, the reduction of fuel consumption and CO₂ production from internal combustion engine passenger vehicles. Approaches to achieve this are briefly described including the optimisation of vehicle control which is then described in further detail. In order to implement practical optimisation systems an accurate and efficient vehicle model is required and this field is then described. Finally the concept of driveability is tackled with details of the current state of research in this area. The application of this background knowledge in the implementation of a vehicle velocity optimisation algorithm is covered in the following chapter.

Chapter 3

Optimal Control Algorithm

The development of in-vehicle systems to provide road data, engine data and driver information has brought many benefits to automotive innovation such as satellite navigation and eco-driving feedback systems. To combine the wealth of information that will be available in a vehicle in the near future with specific component control to achieve for instance lower fuel consumption requires a high level algorithm. Accurate vehicle models aid in the prediction of how different control strategies will influence the fuel consumption and the selection of the most appropriate control policy. This is considered an application of Optimal Control, defined in [121] as the minimising (or maximising) of a performance measure. In order to minimise both journey time and fuel consumption, the vehicle speed and by virtue of gear selection, the engine speed, are required to be optimised. A number of optimal control algorithms have been discussed in chapter 2 and common to all of them is the importance of formulating the problem suitably for the chosen optimisation method. This process is described in the following section. The optimisation's effectiveness is directly linked to the accuracy of the vehicle model applied in the algorithm and the vehicle model is detailed in section 3.3. The implementation of the algorithm in software code is covered in chapter 4 along with an analysis of how different configurations of the optimisation algorithm influence the output and the algorithm is compared to alternative optimisation methods to verify the optimality of the results.

3.1 Optimisation Problem Formulation

In [122] a dynamic programming (DP) problem is presented as having two main features, a cost function that increases cumulatively and a system that can be described by a discrete time model. The problem under investigation here centres on the multi-stage decision process of how best to traverse an infinite road network for a journey of unknown distance while minimising both fuel consumption and journey time. The system then is a vehicle model that calculates both fuel and time at discrete intervals. The state of the system at each interval of the discretisation is described by x . As the goal of optimisation is to maximise or minimise a function, known as the cost function, this approach can be applied to this problem with the cost function representing a combination of the fuel consumed and journey time. In this case the cost function is to be minimised and the optimal control strategy is to be found and applied to u , the control variable. As the road elevation changes as a function of distance then it is first proposed in [123] and subsequently considered in [80–82,124] to consider the formulae in terms of road position as opposed to the more common approach of using time. The cost function, J , is then formulated here with respect to distance, s .

$$\min_{u(s)} J(u(s)) \quad (3.1)$$

The cost function can be described by the following equation as modified from [121] in the time domain to the spatial domain

$$J(u(s)) = G(x(s_f)) + \int_0^{s_f} H(x(s), u(s), s) ds. \quad (3.2)$$

where H is the cost at each position, s , along the road, integrated with the respect to position, s to give the sum of all the costs. This is added to $G(x(s_f))$, a terminal cost associated with the state, x , of the system at the final distance considered, s_f . The distance between s and s_f is referred to as the horizon over which the optimisation is applied. The terminal cost is used to approximate the cost incurred beyond s_f for the remaining journey of unknown distance. This approximation is required to balance the impact of decisions made in the current horizon on subsequent horizons. Without which short term benefits in the current horizon may be selected at the expense of subsequent horizons.

In order to apply the methods of dynamic programming it is required for the continuous system of interest to be approximated with a discrete-time representation. This takes the form

$$x_{k+1} = f_k(x_k, u_k, z_k), \quad k = 0, 1, \dots, N - 1 \quad (3.3)$$

where N is the number of distance intervals that make up the road section under consideration and x_k is the state vector which is defined as

$$x_k = \begin{bmatrix} s \\ v \\ g \end{bmatrix} \quad (3.4)$$

where s is position (m), v is velocity (m s^{-1}) and g is the gear selection, a discrete number with no units. The control vector, u_k is defined as

$$u_k = \begin{bmatrix} u_v \\ u_g \end{bmatrix} \quad (3.5)$$

where u_v is the control variable for vehicle velocity (m s^{-1}) and u_g is the discrete control variable representing the selected gear number. The disturbance, z_k is defined as

$$z_k = \begin{bmatrix} \text{slope} \\ \text{wind} \end{bmatrix} \quad (3.6)$$

The continuous cost function (3.2) can be approximated in the discrete domain as

$$J = g(x_N) + \sum_{k=0}^{N-1} h(x_k, u_k, z_k) \quad (3.7)$$

where $g(x_N)$ is the terminal cost of the final state, x_N .

3.1.1 Constraints

State Constraints

To ensure that solutions are physically realisable some restrictions must be applied to the state space, as described in (3.4). Distances that can be considered are limited by

$$s_0 < s_k \leq s_{max} \quad (3.8)$$

where s_0 and s_{max} are the minimum and maximum distances allowed in the current horizon. The gears that can be considered are limited by

$$g_{min} \leq g_k \leq g_{max} \quad (3.9)$$

where g_{gmin} and g_{gmax} are the lowest and highest gears, respectively and the gears are integer values of the set

$$g_k \in G = \{0, 1, 2, \dots, g_{max}\} \quad (3.10)$$

The velocities considered for each gear are limited by

$$v_{min}(g) \leq v_k \leq v_{max}(g) \quad (3.11)$$

where $v_{min}(g)$ and $v_{max}(g)$ are the minimum and maximum velocities allowed in a given gear, g .

Control Constraints

Along with state constraints, it is required to restrict the control variables to those that are physically realisable with the following constraints

$$v_{min}(g) \leq u_v(s) \leq v_{max}(g, s) \quad (3.12)$$

where v_{min} is the minimum velocity (m s^{-1}) in a given gear and v_{max} is the maximum velocity at a given position in a given gear. The selection of gears is constrained as follows

$$g_{min} \leq u_g \leq g_{max} \quad (3.13)$$

where g_{min} and g_{max} are the lowest and highest selectable gears.

3.1.2 Cost Function

The cost function is the core of the DP algorithm and is developed from [123, 124] with the addition of a variable normalisation factor.

$$J = \begin{bmatrix} \frac{\lambda}{\mu_1} & \frac{1-\lambda}{\mu_2} & \zeta \end{bmatrix} \begin{bmatrix} J_t \\ J_f \\ J_c \end{bmatrix} \quad (3.14)$$

where J_t , J_f and J_c are the costs associated with time, fuel and comfort, respectively. λ is a weighting factor to adjust the cost function to favour either fuel or journey time and μ_1 , μ_2 are scaling factors to normalise each part of the cost function. ζ is the weighting factor for the comfort cost. The cost functions components, J_t , J_f and J_c , are calculated as follows

$$J_t = \frac{\Delta s}{v_{avg}} \quad (3.15)$$

$$J_f = \dot{m}_f \left(\frac{\Delta s}{v_{avg}} \right) \quad (3.16)$$

where \dot{m}_f is the mass fuel flow in g s^{-1} and is multiplied by time, as calculated in (3.15) to give total fuel used in g).

$$J_c = \Delta\tau \left(\frac{v_{avg}}{\Delta s} \right) \quad (3.17)$$

where $\Delta\tau$ is the change in torque in N m. This is used to penalise large changes in the derivative of torque with respect to time which could lead to sharp acceleration and impact driver comfort and satisfaction.

Normalisation

As the cost function is a combination of time and fuel costs which are not in directly comparable units, it is important to normalise these functions to ensure that a change in weighting factor λ has an even effect on both of these cost function elements. In order to do this, the values of μ_1 and μ_2 are to be carefully selected. A minimum value for time and fuel economy is decided and the μ_1 and μ_2 values are selected to ensure the fuel and time cost components range from 0 to 1. The minimum time that would be acceptable varies depending on the speed limits on the upcoming road sections and so the normalisation factors need to be recalculated at each repetition of the algorithm using the road data. The minimum time is calculated assuming the speed limit is followed as closely as the vehicle performance will allow. The minimum acceptable fuel economy is also calculated in this manner. In Figure 3.1 the effect of varying λ is shown with desired behaviour for both the fuel and time values. As the λ value rises the increase in fuel used is proportional to the decrease in time, which is to be expected as increasing

λ adjusts the cost function to penalise time costs more heavily leading to shorter travel times.

3.1.3 Assumptions

When considering the formulation of the optimisation problem it is important to highlight the assumptions that have been made in order to limit its complexity. With the discretisation of the system in section 3.1 it is beneficial to assume the behaviour of the vehicle and road section within one discrete interval can be simplified. Considering a gear shift operation which occurs over the course of two discrete intervals, after the first interval determining the state of the vehicle during a gear shift is a complex task as noted in [100]. As the average time taken to complete a gearshift can be ascertained and the maximum permissible velocity of the vehicle is known, the distance covered during a single gearshift operation can be calculated. It is assumed that a discrete distance interval will be selected that is greater than this gearshift distance and therefore gearshifts will be completed within one discrete distance interval. The road gradient and speed limit between intervals are assumed to be static so as to ensure that the vehicle model is only required to be evaluated once per interval. Aside from discretisation, it is noted that the frame of reference for distance travelled makes some assumptions. Conventionally the distance between two pairs of longitude and latitude coordinates assumes a straight line in two dimensions, neglecting the effect of any elevation variation occurring while traversing between the two points. Using on-board sensors such as wheel speed sensors, it is possible to provide a more accurate representation of the actual distance covered by the vehicle between the two positions in three dimensions. By making these assumptions, the responsibility for ensuring the accuracy of the system model lies with the choice of discretisation interval which is considered in chapter 5.

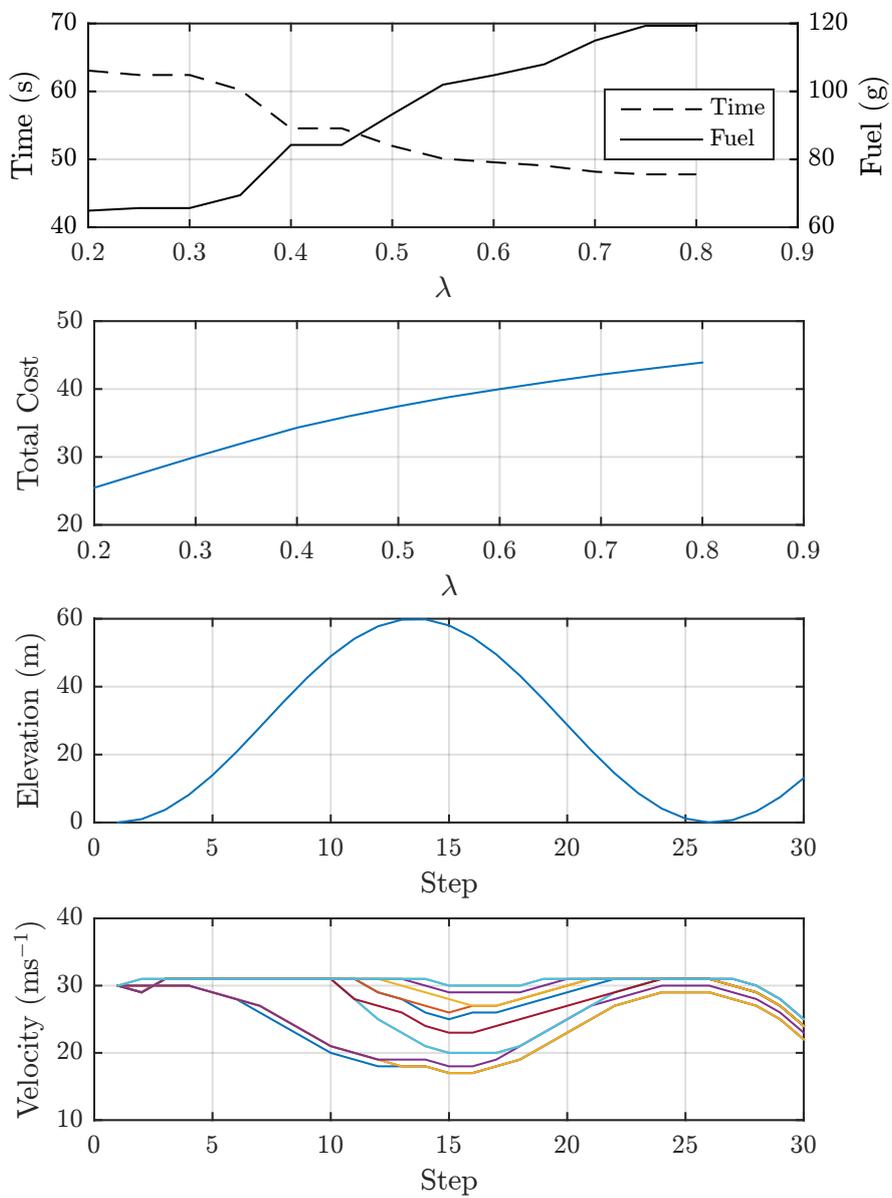


Figure 3.1: Fuel and time values varying as a function of λ (top), total cost as a function of λ (second) for road elevation profile (third) and associated velocity profiles (bottom)

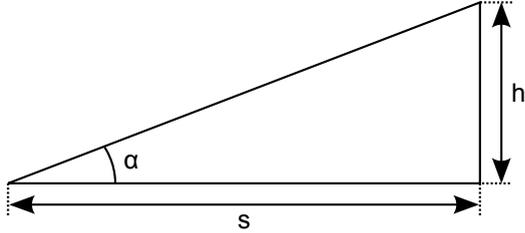


Figure 3.2: Road gradient specification

3.2 Road Model

The road gradient is calculated with reference to Figure 3.2 as

$$\alpha = \arctan\left(\frac{h}{\Delta s}\right) \quad (3.18)$$

where α is the angle of road inclination (rad), h is the change in elevation (m) for the distance interval considered and Δs is the horizontal distance (m). This road model is applicable for paved roads only. The discretisation interval is assumed to be greater than the wheelbase of the vehicle and the gradient static within an interval thus negating the problems of rapidly changing gradient noted in [125] for off-road situations.

3.3 Vehicle Model

The accuracy of the model of the system under study, in this case the vehicle, is fundamental to the quality of the DP algorithm results. The development of the model is described in this section.

3.3.1 Model Type

As noted in chapter 2 there are two types of vehicle longitudinal models. In the DP algorithm a backward-facing model is to be used as, at the algorithm's core, is the repeated calculation of the cost function for different velocity profiles and so the use of velocity as the model input is the most efficient approach. Rather than a complete velocity profile being analysed,

individual steps are investigated in isolation with a start and end velocity considered and step size sufficiently small that a constant acceleration can be assumed between these two velocities.

3.3.2 Vehicle Dynamics

In order to calculate the vehicle movement, the force transmitted to the road by the wheels is compared with the other forces acting on the vehicle to find the net propulsion force required to achieve a given change in velocity. These forces are used as follows

$$m_v \frac{dv}{dt} = F_t(u) - (F_a(t) + F_r(t) + F_g(t) + F_d(t)) \quad (3.19)$$

where m_v is the vehicle mass including inertial equivalent mass (kg), $v(t)$ is the vehicle velocity, F_t is the tractive force (N) applied by the driven wheel tyres that is a function of the control inputs that form u . F_a is the aerodynamic drag force of the vehicle (N), F_r is the rolling resistance (N), F_g is the gravitational force acting on the vehicle (N), F_d is the force of additional disturbances that are not modelled in detail (N). These disturbances include alternator load required to power ancillary items such as lighting, climate control as well as disturbances due to changes in environmental conditions such as wind speed, air pressure or road surface. This formula, with disturbances neglected is presented visually as a Sankey diagram in Figure 3.3. For a given acceleration the required tractive force needs to be found by making it the subject of the formula and calculating the remaining forces.

As proposed in [123] it is beneficial to relate formulae to position rather than time. As the road state changes with position and this can be measured directly rather than estimated based on velocity and time there are advantages to presenting the model with position as an input variable. The drawback with this is that standard equations with time as a variable need to be converted to consider position instead. The relationship noted in [123] using the chain rule is used here

$$\frac{dv}{dt} = \frac{ds}{dt} \frac{dv}{ds} = v \frac{dv}{ds} \quad (3.20)$$

This results in an equivalent to (3.19), with disturbances neglected, as fol-

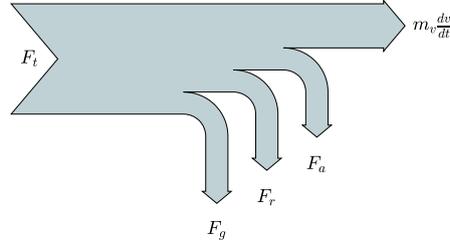


Figure 3.3: Forces acting on a vehicle to resist the tractive force, F_t in producing an acceleration.

lows

$$m_v v(s) \frac{dv}{ds} = F_t(u) - (F_a(v, v_w) + F_r(\alpha(s)) + F_g(\alpha(s))) \quad (3.21)$$

Each component is described in the following sections.

Aerodynamics

As detailed in [46] the complex behaviour of air resistance on a moving vehicle can be approximated by a simple object with a cross sectional area and drag coefficient travelling at a given speed through air of a fixed density. This resistance is calculated as follows

$$F_a(v, v_w) = \frac{1}{2} \cdot \rho_a \cdot A_f \cdot c_d \cdot (v + v_w)^2 \quad (3.22)$$

where ρ_a is the density of ambient air (kg m^{-3}), A_f is the frontal area of the vehicle (m^2), c_d is the vehicle coefficient of drag and v is the velocity of the vehicle (m s^{-1}) and v_w is the headwind speed (m s^{-1}) representing wind opposing the vehicles direction of travel. With the exception of the velocity, all the other elements of (3.22) are fixed according to the vehicle, so the F_a is a function of velocity, v . The coefficient of drag, c_d is assumed to be a constant for the range of scenarios considered in this project.

Rolling Resistance

The deformation of the tyres due to the load of the vehicle produces what can be represented as a resistive force acting against the direction of vehicle movement [100].

$$F_r(s) = c_r(v, p, \dots) \cdot m_v \cdot g \cdot \cos(\alpha(s)), \quad v > 0, \quad (3.23)$$

where c_r is the rolling friction coefficient (dimensionless), p is the tyre pressure and α is the road gradient rad [46]. The effect of tyre pressure and velocity on the rolling friction coefficient is to be neglected here in order to strike a balance between accuracy and computation time. With the exception of α , all the other elements of (3.23) are considered fixed for a specific vehicle and α is dependent on the current position, so F_r is a function of position, s .

Force due to gradient

When a mass is located on an incline there are force components perpendicular and parallel to the inclined slope that combine to equal the force due to gravity. The parallel component is of relevance for calculating the forces acting on the vehicle as

$$F_g(s) = m_v \cdot g \cdot \sin(\alpha(s)). \quad (3.24)$$

As m_v is assumed fixed for a given vehicle, only α will vary and as α is dependent on the current position, F_g is then a function of position, s . The road gradient is as described in section 3.2.

Resistive Forces

In order to accelerate the vehicle the tractive force F_t must be greater than the sum of all the resistive forces in Equation 3.21. The contribution of each component to the total resistive force and their variation is highlighted in Figure 3.4. At low speeds with zero gradient the rolling resistance, F_r has the biggest contribution of all the resistive forces. When any of the vehicle speed, headwind speed or road gradient increase the resulting forces begin

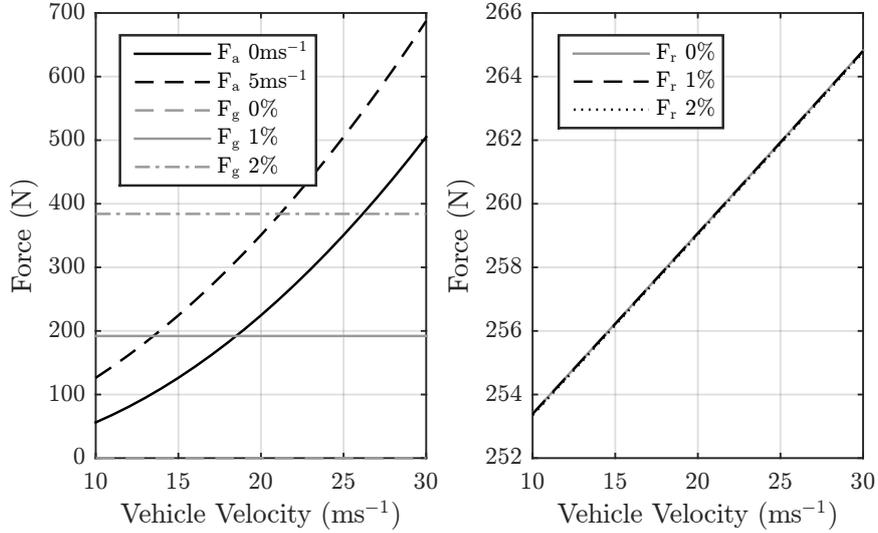


Figure 3.4: The resistive forces, F_a , F_r , and F_g are compared at vehicle speeds from 10 ms^{-1} to 30 ms^{-1} , headwind speeds of 0 ms^{-1} and 5 ms^{-1} and road gradients of 0%, 1% and 2%.

to dominate the resistive force and the rolling resistance becomes less significant. It can also be seen in Figure 3.4 that at the road gradients occurring under normal conditions the rolling resistance does not vary significantly due to road gradient. The sensitivity of each of these forces to variation in their arguments is investigated in chapter 4.

3.3.3 Engine Model

With the tractive force, F_t , required for a given transition described in Equation 3.21, the engine speed and torque are required in order to find the fuel consumption from the BSFC map used in the engine model. For a transition between v_0 and v_1 the acceleration can be calculated using Equation 3.20 and used in Equation 3.21 along with Equation 3.22, Equation 3.23 and Equation 3.24 to find the required traction force at the wheels, F_t . Using this force and the wheel radius, r_w , the torque equivalent is

$$\tau_w = r_w F_t \quad (3.25)$$

This traction torque is related to the torque available from the engine via the final drive and transmission ratios. The following formula describes this

$$\tau_{eb} = \frac{\tau_w}{R_{tr}R_{dr}}, \quad (3.26)$$

where τ_{eb} is the engine brake torque (N m), τ_w is the torque at the wheels (N m) and R_{tr} and R_{dr} are the ratios of the transmission and driveshaft, respectively. The useable torque available from the engine is referred to by convention as the engine brake torque as opposed to the engine indicated torque which is the theoretical torque produced by the engine not considering friction and other losses. The engine speed is also calculated using the transmission ratios

$$\omega_e = \frac{60\omega_w R_{tr}R_{dr}}{2\pi}, \quad (3.27)$$

where ω_e is the engine rotational speed (rpm) and ω_w is the rotational speed of the wheels (rad s^{-1}).

Torque Curve

A compression ignition engine has an upper limit on the torque it can produce that varies with engine speed. In order to reflect this limitation the engine torque and speed must be compared to the upper limit and transitions that require engine operating points above the limit are disregarded as infeasible. The torque curve used in the engine model is shown in appendix A.

Engine efficiency

The efficiency of a compression ignition engine varies over a range of conditions, most notably engine speed. The maximum efficiency of the engine considered in this vehicle model is 39%. As the losses can be attributed to a number of factors as seen in chapter 2 to model this behaviour in detail requires much greater complexity and fidelity than can be afforded in the optimisation algorithm. These complex modelling issues can be avoided by using empirical data of the measured fuel flow at a number of engine brake torques and speeds to create a table that can be quickly accessed to estimate engine efficiency.

3.3.4 Fuel Consumption

To model the fuel consumption of the vehicle there are three options, as described in [46], the average operating point approach, the quasistatic approach and the dynamic approach. The first approach estimates the fuel consumption based on that resulting from the engine run at an operating point that is the average of a range. This approach does not allow for variation of fuel consumption by optimising the engine operating point and thus cannot be used here. The dynamic approach requires mathematical models to be developed that accurately reflect the behaviour of the engine typically in the form of ordinary differential equations. The quasistatic approach involves calculating the fuel consumption at intervals over the entire profile. At each interval the engine torque and speed is assumed constant. This approach is well suited for use in a DP application due to its reliance on calculations occurring at intervals which in this case are inherent in the DP algorithm. This method does not consider the influence of engine dynamics but rather considers the fuel consumption at each operating point from a table at a number of engine brake mean effective pressures (BMEP) and speeds. The table is a discrete approximation of the following equation, modified from [2]

$$sfc = \frac{3.6 \times 10^6 \dot{m}}{\tau_{eb} \omega_e} \quad (3.28)$$

where sfc is the specific fuel consumption ($\text{g kW}^{-1} \text{h}^{-1}$), \dot{m} is the mass flow of fuel (g s^{-1}). The total fuel consumption for the horizon can then be calculated as follows

$$M = \int_0^{s_f} \frac{1}{v} \dot{m}(\omega_e, P_{me}) ds \quad (3.29)$$

with the discrete equivalent

$$M = \sum_{k=1}^N \frac{1}{v_k} \dot{m}(\omega_{e_k}, P_{me_k}) \Delta s \quad (3.30)$$

As the engine BSFC data is provided at discrete intervals of BMEP and engine speed that may not correspond exactly to the values required by the model it is necessary to interpolate the table data. The two-dimensional interpolation method used is bilinear interpolation.

$$f(x, y) = a_{00} + a_{10}x + a_{01}y + a_{11}xy \quad (3.31)$$

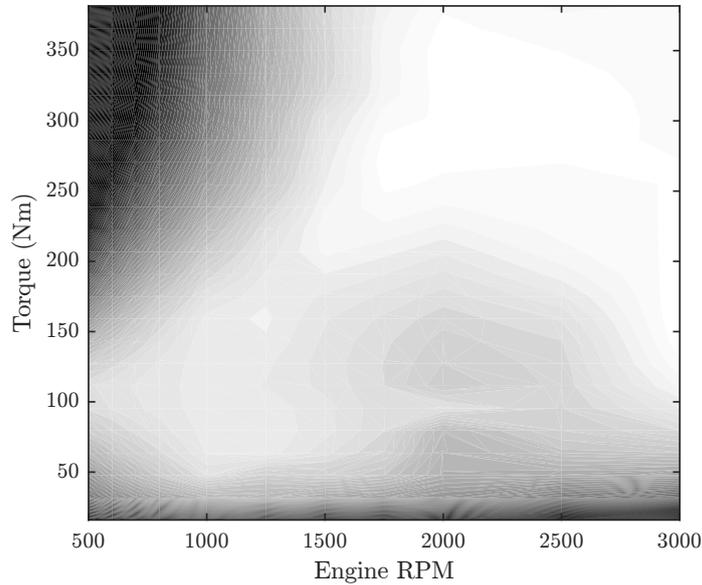


Figure 3.5: Brake Specific Fuel Consumption with engine coolant 90 °C where darker colour indicates lower efficiency

$$\begin{aligned}
 a_{00} &= f(0, 0) \\
 a_{10} &= f(1, 0) - f(0, 0) \\
 a_{01} &= f(0, 1) - f(0, 0) \\
 a_{11} &= f(1, 1) + f(0, 0) - (f(1, 0) + f(0, 1))
 \end{aligned}$$

where x is the engine speed required and y is the BMEP value required. The BSFC data used is shown in Figure 3.5 as measured when the engine coolant temperature was 90 °C and provided by the vehicle manufacturer. The temperature of this coolant is directly related to that of the engine and therefore of the engine lubrication fluids. The viscosity of the lubrication fluids reduces with increasing temperatures which leads to a reduction in engine friction as the engine warms up. Above the optimum operating temperature the lubricants begin to degrade and the engine friction increases.

Separate BSFC maps are provided for engine coolant temperatures of 30 °C, and 50 °C, and an investigation into the effect of different BSFC maps is detailed in chapter 5.

3.3.5 Time

In order to calculate the time for each step we use the equation

$$v = \frac{ds}{dt} \quad (3.32)$$

The time taken then becomes

$$t = \int_0^{s_f} \frac{1}{v} ds. \quad (3.33)$$

which in discrete form is represented by

$$t = \sum_{k=1}^N \frac{1}{v_k} \Delta s \quad (3.34)$$

The vehicle model as described produces both the fuel and time values that are required for the calculation of the cost function (3.14). In order for the vehicle model to be implemented in the DP algorithm it is required to be discretised.

3.3.6 Discretisation

The continuous system described in (3.21) is required to be discretised, or numerically approximated and there are many methods for the numerical approximation of ordinary differential equations such as this. The well used Euler method is a first order method of numerical approximation and there are two approaches that can be taken, using either a forward or a backward method. Each is investigated to identify which is more suitable for this application. The forward approach is presented as

$$\frac{v_{i+1} - v_i}{\Delta s} = \frac{1}{m_v v_i} (F_t(u_i) - F_d(v_i, s_i)) \quad (3.35)$$

where F_d is used to represent the combination of all the resistive forces in (3.21). The backward approach is presented as

$$\frac{v_{i+1} - v_i}{\Delta s} = \frac{1}{m_v v_{i+1}} (F_t(u_i) - F_d(v_{i+1}, s_{i+1})) \quad (3.36)$$

To compare the two methods a test scenario is considered with a road section of zero gradient, a fixed start (v_0) and end speed (v_2) and an intermediate step of variable speed (v_1). The work done is calculated using the

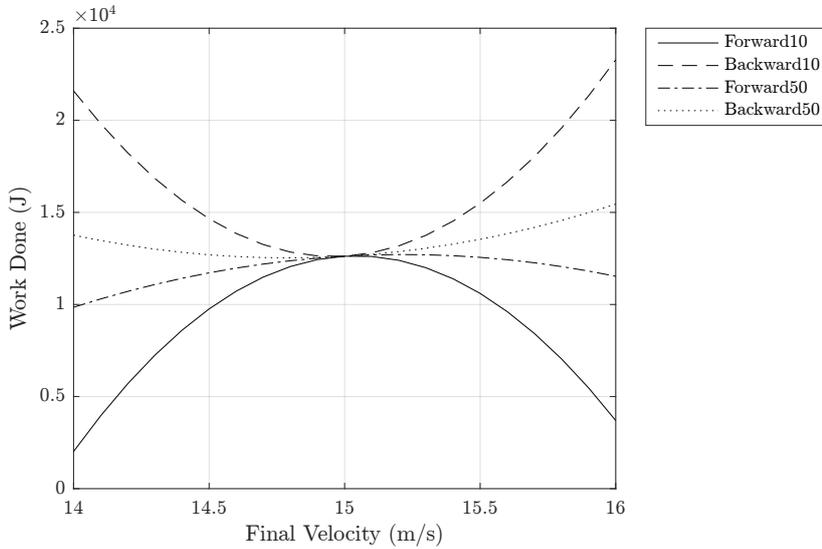


Figure 3.6: Euler forward and backward discretisation methods comparison.

force required to overcome the resistive forces and propel the vehicle over a fixed discretisation interval distance, Δs . Rolling resistance is neglected for this scenario. The work done as calculated by Euler forward and backward discretisation methods is shown in Figure 3.6, with a two step section considered with identical start and end velocity (in this case 15 m s^{-1}) and a variable intermediate velocity. A discretisation interval of 10 m is shown for two curves and for an interval of 50 m for two curves, as labelled. It can be seen that using the backward Euler method estimates the work required increases when the velocity, v_1 , is increased compared to the constant speed. However the forward method estimates that the work required decreases as the intermediate velocity increases, which does not reflect reality. The discretisation interval also influences the results as shown by the work done as estimated with an interval of 50 m compared to 10 m. Even for the backwards method, at the larger interval the work done is briefly below that of constant speed as v_1 decreases from 15 m s^{-1} to 14.6 m s^{-1} which again does not reflect reality. Despite this issue at higher interval sizes the output of the backwards method is consistently more representative of reality and so is applied in this case as opposed to the forward method or more complex discretisation methods.

Discretisation Interval

The process of converting the continuous system to a discrete system requires consideration of the interval period used in the discretisation. For a fixed horizon length, reducing the interval distance will require an increased number of calculations. In [123] a distance interval of 20 m is used in simulations for roads with rolling gradients and 10 m is used in simulations of cruising scenarios. The interval used in [124] is 50 m with 30 steps giving a total horizon of 1500 m. In [81] where a cloud-based optimisation is considered, the discretisation interval is dependent on the speed limit of the given section. For sections limited to 30 mph or below the interval is 50 m and for speed limits above this the interval is 150 m. While this approach can reduce the total number of search space calculations for a given horizon distance, the complexity of implementing the varying interval sizes is a significant drawback. Basing the interval distance on the speed limit neglects to consider situations where the vehicle speed varies significantly from the legal speed limit, for instance during heavy traffic or at traffic lights. In these scenarios it may be desirable to maintain a reduced interval distance regardless of the legal speed limit. Also in [81], the gradient is discretised with intervals of 0.5° . This stands out from other publications as elsewhere the gradient calculation is assumed to occur in advance of the optimisation and cost function computations involve trigonometric functions that do not benefit from a discretised interval in gradient. The computation time of a trigonometric function is constant regardless of the angle of gradient involved, thus providing no benefit to discretising the gradient in this way.

Velocity

Aside from gear selection which is inherently discrete the other dimension in the search space that is to be discretised is that of the vehicle velocity. In order to minimise the size of the search space, a sufficient discretisation interval is to be applied to the available velocity choices. The effect on an optimisation output of varying the velocity interval is examined in [73] where it is proposed that a resolution of 2 mph produces fuel consumption results less than 1% higher than at 1 mph. This result relates to a small sample

size and further investigation of this issue is covered in chapter 5.

3.3.7 Assumptions

The computation time of the algorithm is a function of the dimensions of the search space and the computation time of each calculation of the vehicle model. Restrictions of the dimensions of the search space are detailed in subsection 3.4.1. To limit the computation time of the vehicle model a number of assumptions are made that reduce its complexity while maintaining sufficient accuracy.

Vehicle mass and distribution

Within the volume of the vehicle there are many components that are able to move with varying degrees of independence from each other which contribute to the movement of the vehicle body. As noted in [126], where there is a requirement to reduce the complexity of a vehicle model it is common to consider all the components as one entity, referred to as a lumped mass. This convention is implemented in the vehicle model with a single point mass, including the wheels, representing the centre of gravity at the origin of the vehicle reference frame.

Prior to the start of a journey the vehicle mass is not always the same and the effect of this is studied in section 4.8.1. However the mass of the vehicle is assumed to remain constant for the duration of a given journey, as although the use of fuel will reduce the weight, no occupants or luggage should change. A typical fuel tank of 50l of diesel has a mass of 42.5 kg. When considering vehicle acceleration, the mass of the vehicle includes the equivalent mass due to the inertia of rotational elements in the drivetrain and the wheels [126]. This equivalent mass is approximated with the addition of a fixed mass to the overall vehicle mass. The equivalent mass due to inertia decreases as the gear ratios decrease at higher gears where there is less variation in gear ratio and thus equivalent mass, allowing a fixed approximation to be used.

Fixed acceleration

When considering two adjacent interval steps, s_1 and s_2 , where a gear transition does not occur, the velocity profile between the two steps is considered to be a linear gradient due to a fixed acceleration. This ensures that no further subdivision of the vehicle behaviour during the step interval is required thus restricting the complexity of the model. As discussed previously the key to this assumption is the correct implementation of the discretisation process to ensure minimal deviation from the vehicles true behaviour.

Wheels and tyres

The modelling of vehicle wheels and in particular tyres as noted in chapter 2 is a rich subject with many models developed with specific purposes in mind, such as calculating the lateral acceleration limits. This level of detail required for those models is unnecessary for the application studied here and so a rigid tyre body with fixed dimensions, pressure and coefficient of friction is used. The rolling resistance, as described by (3.23), assumes a lumped mass for the vehicle with weight consistently distributed evenly between the front and rear axles, thus ensuring that variation between rolling resistance of front and rear tyres, as described in [100], can be neglected in favour of one single all encompassing term. Longitudinal slip is the difference between vehicle velocity and the translational equivalent of the wheel rotational velocity and leads to additional losses in the propulsion system. The longitudinal slip varies under high acceleration and deceleration forces and is neglected here due to the calculation overhead and the absence of such extreme driving behaviour that the longitudinal slip will vary significantly.

Efficiencies

The engine efficiency is considered in subsection 3.3.3 but the efficiency of each component in the powertrain also impacts the fuel consumption to varying degrees and the efficiency can vary with rotational speed and temperature.

The transmission of power through the gear box incurs losses from a variety of effects such as mesh friction of the gears, bearing friction and

pumping losses in the torque converter. The overall efficiency of the transmission in each gear is assumed to be fixed at 97% for the purposes of minimising computation time without compromising model accuracy. The torque converter used in the automatic transmission can be mechanically locked to prevent the losses associated with the transfer of power through the fluid of the torque converter. The lockup clutch is applied whenever the torque converter impeller and turbine are rotating synchronously which is assumed to be any time other than starting from stationary and during a gear shift.

Environmental variables

The vehicle is affected by a number of environmental variables some of which are considered in the model and some of which are neglected. The aerodynamic drag force as described in (3.22) neglects the influence of wind speed and direction. In [127] it is estimated that wind conditions can reduce fuel economy by 2-3% compared to test cycles undertaken in test laboratories. Even with a uniform distribution of wind directions that produce a net zero wind direction over the course of a journey, a headwind will have a greater effect than a tailwind due to the squared velocity value in (3.22) and result in a reduced fuel economy. Due to the absence of wind measurement equipment in a production vehicle it is considered impractical to estimate this in detail, however average wind conditions for a given location can be used. Wind speed data is recorded for meteorological purposes however air flows local to the vehicle can vary greatly and rapidly compared to the wind speed data measurement.

Changes in atmospheric pressure and temperature will influence aerodynamics, as the density of air is directly related to both by the ideal gas law. While the air density can be estimated from the vehicle pressure and temperature sensors, it is assumed that it will not vary during the course of a specific journey. Analysis of the fuel consumption sensitivity to air density is seen in chapter 5. Such atmospheric changes will also have an effect on engine operation, however the engine efficiency data recorded to produce the BSFC map were collected under standard atmospheric conditions and

such conditions are assumed.

The material and condition of the road surface will also have an impact on rolling resistance due to the coefficient of friction which represents the interface between road and tyre. As road surface condition measurement is not commonly available on production vehicles such data is not considered here. Therefore, it is assumed that all journeys will take place on conventional highway road surfaces and the weather conditions will be such that the coefficient of friction is not affected.

The road gradient is considered to change over an interval much greater than the length of the vehicle ensuring that the vehicle body is always assumed to be parallel to the road surface. This assumes that all road surfaces are smooth, well maintained and designed according to [128] to avoid sharp changes in gradient. Superelevation due to road camber is not considered in the road.

With these assumptions in place the vehicle model can be used within the framework of the optimisation algorithm to produce a cost optimal velocity and gear profile for the upcoming road section.

3.4 Forward Dynamic Programming

With a suitably discretised optimisation problem state of the system can be calculated at each discrete interval. The Dynamic Programming implementation can approach this either in reverse chronological order from a final velocity to a known initial velocity, or in the opposite direction from the initial velocity to multiple possible end velocities of different total costs. In order for the DP algorithm to work optimally in real-time it is not possible to know what the final velocity should be before undertaking the optimisation and so the forward DP algorithm is used in this application. From the current initial velocity the possible future velocities are considered at each fixed interval in the upcoming road, thus building a grid of possibilities known as the search space.

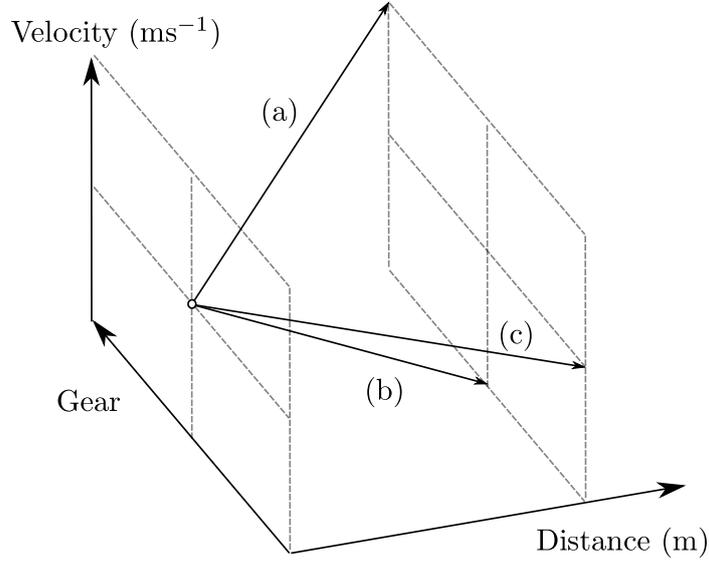


Figure 3.7: Dynamic Programming search space with transitions (a) increasing velocity and selecting a higher gear, (b) decreasing velocity and (c) maintaining velocity and selecting a lower gear.

3.4.1 Search Space

The cost function is to be evaluated at each node of the search space, where each node represents a unique state of the system as defined in (3.4). The search space is represented diagrammatically in Figure 3.7 with three transitions illustrating velocity and gear changes, where transition a is to a higher velocity while shifting to a higher gear, b is to a lower velocity while remaining in the same gear and c is maintaining velocity while shifting to a lower gear. Using the constraints noted in subsection 3.1.1 the limits of the search space can be defined in all dimensions. The search space size is directly linked to the algorithm computation time and any reduction in size that can take place prior to transition calculations proceeding will ensure a reduction in the time to produce a result. The complexity of the optimisation algorithm (O) can be calculated from the dimensions of the search space

$$O(N_s \cdot N_v^2 \cdot N_g^2) \quad (3.37)$$

where N_s is the number of distance intervals between s_0 and s_{max} , N_v is the number of velocity intervals between v_{min} and v_{max} and N_g is the number of

gears between g_{min} and g_{max} . As can be seen the number of distance intervals has a linear relationship with complexity as opposed to the velocity and gear intervals which both have an exponential relationship with complexity. The gear interval spacing is fixed by the discrete nature of the automatic transmission, however the relationship described in (3.37) assumes that all gears N_g are available from each of the gears which in reality is not feasible. To ensure feasible solutions while also reducing the search space complexity a restriction is imposed on the gear selection allowing a maximum difference of two between consecutive gears with the exception of coasting in gear zero. In a similar manner, the road speed limit can be used to reduce the upper boundary of the velocity search space to only allow legal speeds. The physical performance limits of the vehicle can also be used, as the maximum possible acceleration from the point of origin will immediately exclude velocity nodes above that reached by the maximum acceleration. Similarly with the maximum possible deceleration as dictated by the braking force, provides a lower boundary on the velocity search space. The engine speed boundaries can also be used to limit the search space for each gear as the engine speed limits relate to vehicle speed limits via the transmission and driveline ratios and wheel radius, this being an implementation of the constraint noted in (3.11). The search space is made up of nodes that contain the following information about that particular state

- lowest cost to reach node
- previous node on lowest cost path (i.e. velocity and gear) to allow the optimal profile to be reconstructed at the conclusion of the algorithm.

The process of executing the DP algorithm can be separated into two parts. Firstly, the cost of each transition from the point of origin in sequence by distance interval to the end of the current horizon need to be calculated individually. Once this is completed, the lowest cost at the end of the current horizon is identified and the lowest cost path is traced back to the point of origin to complete the optimal path.

3.4.2 Forward Process

The first stage of the DP algorithm calculates the cost from the point of origin to all of the available first step points. In order to calculate this, the following information is required by the vehicle model

- road gradient and speed limit between the point of origin and step one
- current velocity
- current gear selection
- current torque

Using this information each gear is taken in turn and the upper and lower velocity limits are calculated with the range dictated by the engine speed limits, vehicle acceleration limits and speed limits. Where the current gear and aim gear are different and a shift is required the calculation for vehicle acceleration limits include a shift time which reduces the maximum possible velocity. Following this, the possible future states are known and the search space boundaries are set so the transition costs can be calculated for each possible state. If a gear shift is not required the transition cost is calculated as in subsection 3.4.3 otherwise the cost incurred when a gear shift occurs is calculated as in subsection 3.4.4.

3.4.3 Transition Cost

In order to calculate the cost for a transition between two nodes as shown in the left of Figure 3.8 with a given start and end velocity, v_0 and v_1 , respectively and gear selection maintained, the associated acceleration with respect to distance is found using (3.20). This is then used in (3.21) and (3.25) to find firstly the tractive force, then the torque necessary to achieve the required acceleration. The current gear and final drive ratios are used to calculate the required engine torque and engine speed using (3.26) and (3.27) respectively. These engine values are used in conjunction with the BSFC table as described in subsection 3.3.4 to provide a corresponding fuel flow rate. Finally the transition time is calculated using (3.34) and together with the fuel flow rate can be used to calculate the total fuel used and therefore

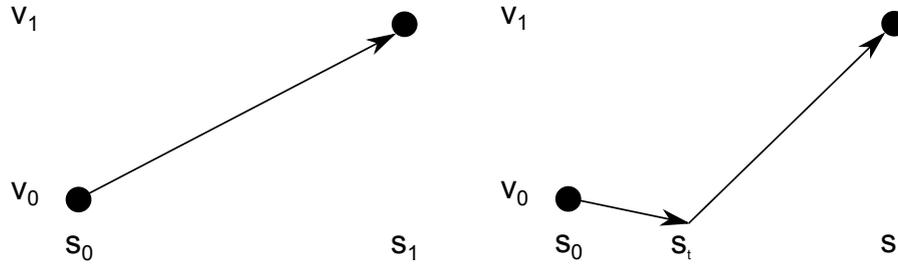


Figure 3.8: Transition cost example with gear shift (right) and without (left). The distance necessary to complete the gear shift operation is noted as s_t followed by the remaining distance over which the final velocity, v_1 is to be reached.

the overall cost function can be calculated using (3.14).

The cost of this transition is added to the lowest cost to reach the start node to give the total cost to reach the end node. If this is lower than the current lowest cost assigned to the end node then the new lowest cost is assigned to the node along with the relevant data associated with the transition including the coordinates of the start node.

3.4.4 Transition Cost with Gear Change

A similar, but expanded, subroutine is required to calculate the transition cost when a gear change occurs as shown in the right of Figure 3.8. The start and end velocity are still provided but the calculation is broken into two sections. Firstly the period during shifting is considered and then the behaviour once the new gear is engaged. During the initial period of the shift when the engine is disengaged, no positive torque can be provided at the wheels and so (3.21) can be calculated with F_t equal to zero. The subsequent acceleration (which could be negative) is used to find the velocity at the point of engine re-engagement. From this velocity at re-engagement to the specified end velocity the required acceleration for this final portion of the transition can be calculated. The fuel and time for each portion of the transition are calculated as in subsection 3.4.3. The costs for each of

these portions are then combined to give the overall transition cost.

3.4.5 Terminal Cost

On completion of the transition cost calculations and allocation to the relevant search space nodes the final costs can be assessed. As the optimisation is only required to act upon the horizon provided there is no consideration for what happens after the horizon and as such it is likely that the optimal velocity profile will end with a sharp reduction in velocity at the end of the horizon. Due to the reduction occurring briefly at the end of the velocity profile the time element of the cost function is only minimally increased while the fuel penalty can be zero as opposed to the fuel required to maintain a steady velocity for instance. As the algorithm is to be executed repeatedly in a vehicle where the journey may be further than the length of the horizon then this velocity reduction will present an issue for the next horizon and some consideration has to be made for this issue. As described in section 3.1 the use of a terminal cost is a way to minimise this problem. In this case the terminal cost is calculated in the same manner as a single step transition as in subsection 3.4.3 with the start and end velocity equal and a flat road considered [76]. The cost for this single step is then multiplied by the number of potential future steps that are being considered for the terminal cost. The influence of the terminal cost, in particular the number of future steps considered, is investigated in depth in chapter 5.

3.4.6 Coasting

An additional benefit of incorporating gear shifting into the DP algorithm is that it allows the possibility of utilising a neutral gear to implement a coasting strategy where beneficial. As discussed in subsection 2.3.2 the benefits of coasting have been shown in a number of publications. The process of assessing the cost of starting a coasting operation is the same as that described in section 3.4.7 however gear zero is the destination gear after the transition and torque cannot be provided from the engine. The force generated by the transfer of gravitational potential energy to kinetic energy on a descending road slope is the only available option to increase or

maintain the vehicle speed against the drag forces in (3.21) therefore limiting the maximum velocity. For nodes below this velocity a sufficient amount of braking force can be assumed where required to ensure the transition reaches each desired velocity. The operating policy when the engine is disengaged is to either supply sufficient fuel to keep the engine running at idle speed, or to completely turn off the engine and restart it when required. This decision is not considered part of the optimisation process and is set during the optimisation initialisation. There is a safety implication of coasting, in the event that torque is required from the engine for an evasive manoeuvre there will be a delay while re-engaging the engine. This is not considered in the optimisation as safety is assumed to be managed by an independent process in a production vehicle.

3.4.7 Driveability

Torque variation limitations

In order to maintain driver comfort it is noted in section 2.6 that vehicle longitudinal acceleration and jerk are important considerations in this regard. Rapid changes in torque supplied by the engine are directly related to jerk behaviour that is detrimental to driver comfort and so are to be penalised through the use of the cost function factor noted in Equation 3.17.

Gear shifting

The transition calculation with a gear change can, under specific circumstances, produce a lower cost than an equivalent transition without a gear change due to the period of driveline disconnect being followed by a short period of higher torque applied to reach the end velocity which allows the engine to operate at a higher efficiency. This lower cost can lead to the optimal strategy including frequent gear changes which may not be acceptable for driver comfort as noted in section 2.6. In [110] the gear shifting is limited by the use of a cost function considering total number of gear change events over the journey. As the total number of gear events is being considered this will not prevent a small number of gear changes over a short period of time, which would still be undesirable for driver comfort so an alternative

is required. It is proposed to prevent gear changes in consecutive steps, although the effectiveness of this depends on the vehicle velocity. At higher velocities one step will be completed in a shorter time than at lower velocities so fixing a time period between consecutive gear shifts must consider the maximum velocity and thus shortest time. At lower velocities such a gear shift restriction would lead to large time intervals between allowed shifts.

The two additions to the optimal control algorithm noted above represent proof of hypothesis (b) which states

- (b) Driving experience, in relation to a vehicle's longitudinal performance, can be quantified and applied to an optimisation algorithm.

The results of these two additions are examined in chapter 5.

3.5 Summary

The optimisation of future vehicle velocity and gear selection profiles to minimise fuel, time and driver discomfort are presented as a cost function to be minimised by applying a dynamic programming algorithm. The models used in the algorithm to describe the road, vehicle longitudinal dynamics and engine performance are detailed along with the structure and sequential process of the algorithm.

Chapter 4

Implementation

While the approach of DP aims to represent problems in a format that is suited to evaluation by a computer program, the process of implementing such a program to solve the DP problem still involves several obstacles. Two situations were considered where the DP algorithm would have to be implemented, the first being a simulation environment to allow testing of various scenarios and the second being implementation on a hardware platform capable of running the algorithm in real time in a test vehicle. Matlab and Simulink are utilised across many academic and industrial fields, particularly the automotive industry [129] for modelling and simulation work and it was intended that this software would be used to simulate the DP algorithm for testing. To run the DP algorithm in a vehicle a hardware platform was required that was compact, portable, had a flexible operating system and enough processing power to run the necessary software and compatibility with hardware components required to make up the complete system. There are a number of products available that would suit this purpose, including the BeagleBone Black [130], Arduino/Genuino Mega [131] and Raspberry Pi 2B [132]. The Raspberry Pi 2B single board computer was chosen as the intended hardware target for the algorithm due to the quad core processor allowing parallel processing of different modules of the system, the lower cost and the abundance of documentation available for incorporating the GPS, CAN bus and 4G mobile connection accessories necessary for the system. The specification of the hardware is detailed in Table 4.1. The soft-

	Raspberry Pi 2 B	BeagleBone Black
Processor	900 MHz Quad-core	1 GHz Single Core
RAM	1 GB	512 MB
Connections	4 x USB HDMI Ethernet General Purpose IO	1 x USB HDMI Ethernet General Purpose IO
Power Source	Mini USB	Mini USB or 5V Jack
Cost	£30	£40

Table 4.1: Hardware comparison of miniature computers considered for project.

ware implementation of the algorithm therefore had to be flexible enough to perform adequately in both scenarios and the programming language used to execute the algorithm is a crucial factor in this.

4.1 Programming Language

Simulink is able to run code that has been produced in MATLAB[®], C, C++, or Fortran and compiled using the Matlab executable (MEX) compiler. Writing software in Matlab limits portability as only systems with Matlab can run this type of program. The C programming language has a long history of use in the automotive industry [133], giving it an advantage over the two remaining alternatives.

The Raspberry Pi 2B has a Debian [134] based operating system pre-installed called Raspbian [135] and this means that a wide variety of programming languages can be considered, however, again the advantages of C also apply here. While the core of the DP algorithm was to be written in C, an interface with external inputs such as GPS position was required. This interface was written as a Python script due to the ease of development and existing libraries for interacting with peripherals such as GPS receivers [136] and CAN Bus interfaces [137]. The DP algorithm C code was compiled using the GNU Compiler Collection (GCC) [138] which is pre-installed with

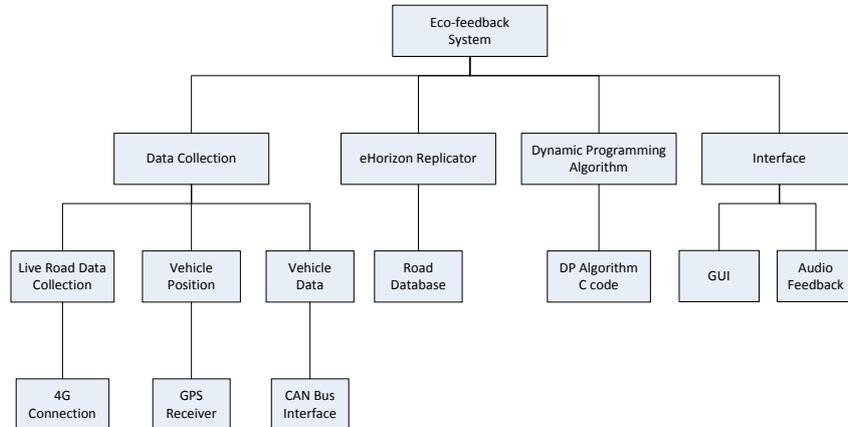


Figure 4.1: Functional decomposition of the developed real-time eco-guidance system with Interfaces (GUI and audio), DP algorithm, eHorizon replicator and data collection consisting of live road data, vehicle position and live vehicle data.

the Raspbian operating system.

4.2 Structure

The overall eco-guidance system software package developed considers the acquisition of current data, accessing a database of road information, the integration of the DP algorithm, data logging and the displaying of relevant data on a Graphical User Interface (GUI) along with audible feedback. In order to provide each of these tasks with sufficient computational power the multiple cores of the Raspberry Pi 2 were utilised with the multiprocessing standard library in Python [139]. The individual processes are shown in the overall structure in Figure 4.1.

4.3 Road Data

In order to test the DP algorithm in realistic scenarios and ultimately deploy a real-time in-vehicle optimisation system real road data would have to be available to the system. As noted in chapter 2 digital map providers

and automotive suppliers are able to provide solutions to access real road information based on current location and predicted route. In order to minimise cost and complexity for this project it was proposed to develop an equivalent road information provider with functionality limited to providing only the relevant information for the optimisation algorithm. This module is developed in two operations, firstly the collection and processing of road data to develop a database of test roads. Secondly, deploying this database in the python program and using the current location from the GPS device to extract the relevant road data to be presented as a road horizon for the optimisation algorithm.

4.3.1 Road Data Collection

To collect road data, two methods have been implemented to increase the number of roads available for simulation and testing. Firstly a bespoke road data logger was developed and implemented on the Raspberry Pi hardware detailed in Table 4.1. The second method used the data from a number of vehicles using commercial vehicle dataloggers which was available not only for driver analysis and fuel consumption comparisons but as the dataloggers are equipped with Global Positioning System (GPS) receivers, the location data can be used to reconstruct road profiles.

4.3.2 Bespoke Road Data Logger

To reconstruct the road profile a GPS receiver is used to measure the test vehicle's longitude, latitude and elevation as it traverses the road. The Raspberry Pi hardware is utilised along with a USB GPS receiver [140]. The GlobalSat WorldCom Corporation model BU-353-S4 as seen in Figure 4.6 is used with an existing Python library [136] that provides a conversion from a variety of GPS communication protocols to a standard, readable JavaScript Object Notation (JSON) format. This library presents the GPS data in a single object such that individual attributes such as latitude can be accessed separately. The required data is then extracted and logged for later post-processing.



Figure 4.2: GPS signal error

4.3.3 Road Data Processing

Due to the nature of receiving GPS signals in a moving vehicle under varying environmental conditions such systems can be prone to positioning errors. The bespoke road data logger is capable of recording GPS data with a frequency of 5 Hz which ensures that any errors in signal processing do not extend for large sections of road while any loss of signal is quickly rectified on return of the connection. The commercial data logger, while recording GPS data at a similar frequency (4 Hz), produced data that was more prone to errors such as that seen in Figure 4.2. Due to the presence of such errors the data must be filtered and a Savitzky-Golay filter [141] is used for this in line with the method presented in [120]. This low pass filter is used to remove noise and the Matlab script used for this filtering was presented in [142].

Distance Calculation

As the algorithm considers road elevation to be a function of road distance it is necessary to calculate the road distance from the recorded position information. Using longitude and latitude coordinate pairs the distance between two road data points can be calculated using the haversine formula [143] used in navigation. This formula is used to calculate the distance between two points on a sphere, which approximates the shape of the Earth and was implemented as a Matlab script.

Discretisation

The smoothed data is required to be discretised prior to reconstructing the road data for use by the DP algorithm to reduce computational overhead associated with the road database and extracting the required data from said database. An example of discretised road data is shown in Figure 4.3, in this instance 50 m intervals are used.

Gradient Calculation

With the elevation data smoothed and the road data discretised, the gradient can be calculated at each step interval in the road using the difference in both elevation and distance as in Figure 3.2. An example of the gradient calculated from real road data is given in Figure 4.3.

Additional Road Information

Along with elevation data it is necessary for the optimisation system to have access to legal speed limit data to ensure that the optimal velocity profile is within these limits. In order to add this information at each step interval in the road data, the longitude and latitude are used to query the speed limit on the given section of road. There are many commercial providers of speed limit map data [144] as well as crowd sourced services such as Open Street Maps (OSM) [145] however due to the cost of commercial providers and the complexity of extracting data from OSM, the free, limited usage routing service of Here [146] was used. This service allows a registered user to make 15,000 requests per month. From a Python script running the bespoke road data logger software, a request is sent to `http://route.st.nlp.nokia.com/routing/6.2/getlinkinfo.json?app_id={APP_ID}&app_code={APP_CODE}` for each pair of coordinates on the current road section. The response is parsed to extract the *SpeedLimit* data which contains a speed in m s^{-1} . The speed limit information is then appended to the road data.

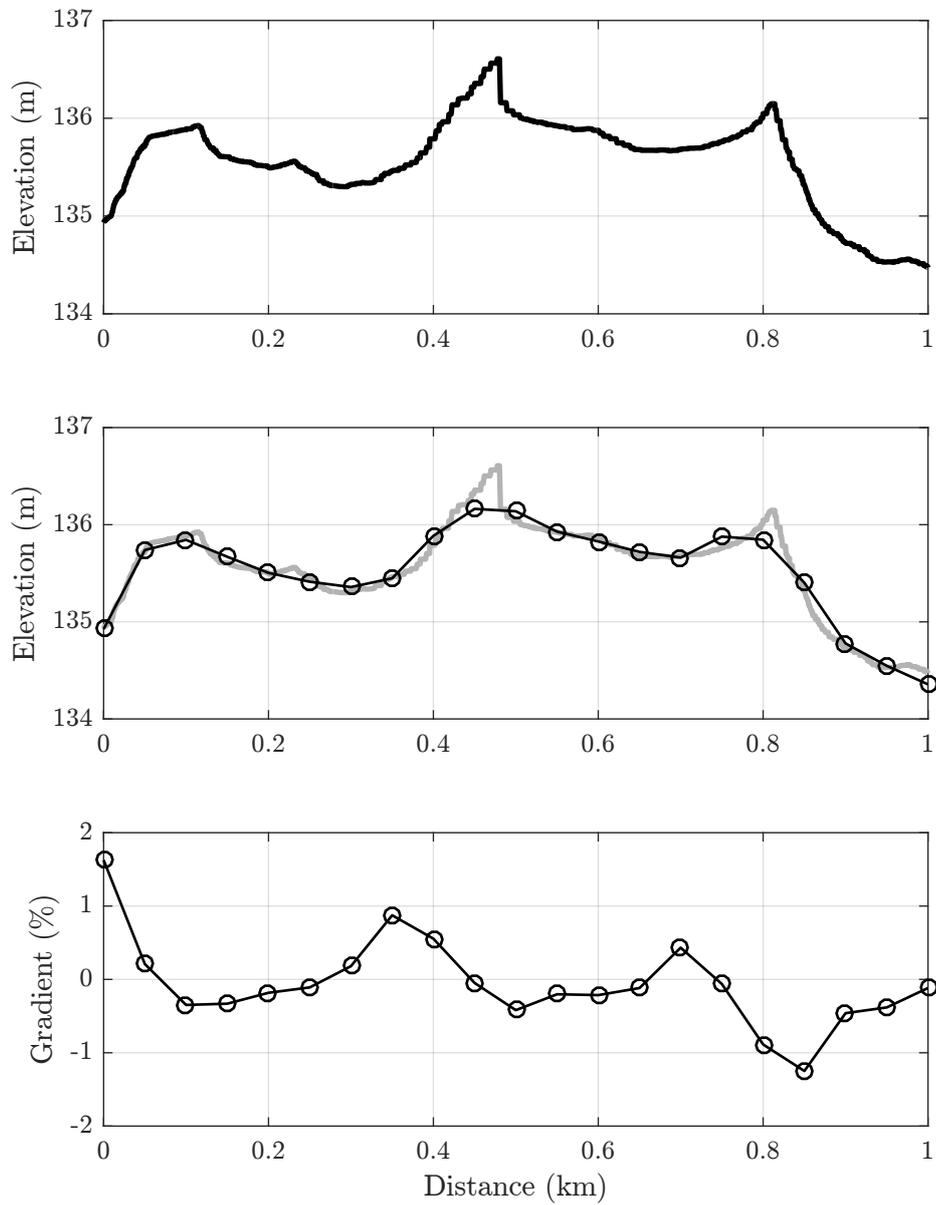


Figure 4.3: Road data processing beginning with raw data (top), filtered and discretised data (middle) and gradient (bottom).

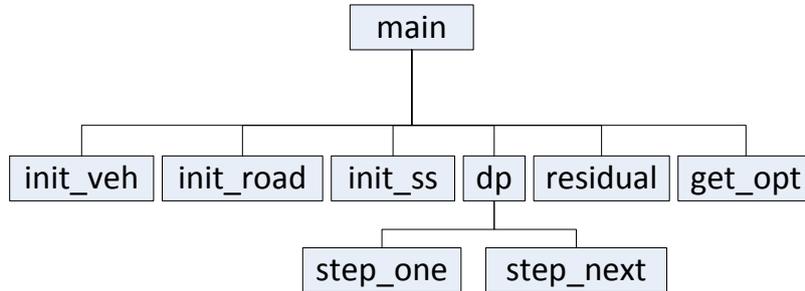


Figure 4.4: Dynamic Programming software structure

4.4 Algorithm Development

The development of the algorithm in C code is detailed here sequentially, in the order that the code is executed. Both the simulation work and in-vehicle implementation use the same C code at the core with variation only in the code that defines the interface with the higher level language, either Simulink or Python, respectively. The functions that make up the dynamic programming algorithm are represented in Figure 4.4 starting with the entry point for the program, the main function, taking in input arguments described in subsection 4.4.1. The vehicle and road data are initialised in *init_veh* and *init_road* as described in subsection 4.4.2. Similarly, the search space, *init_ss* in subsection 4.4.4. The core of the dynamic programming algorithm is contained in the *dp* function, which includes repeated use of the *step_next* function for each step in the search space, with the exception of the first step which is considered in *step_one* and described in subsection 4.4.5. On completion of cost calculations at each step a terminal, or residual, cost is applied in function *residual* described in subsection 4.4.6. Finally from the complete set of costs including terminal cost, the optimal trajectory is selected by the function *get_opt* which is described in subsection 4.4.7.

4.4.1 Input arguments

The starting point of the algorithm considers the input of essential data formatted for use in the algorithm. These input arguments and their data

Structure	Item	Type
Road Data	Distance (m)	Float
	Elevation (m)	Float
	Slope (%)	Float
	Speed limit (m s^{-1})	Float
Initial Conditions	Velocity (m s^{-1})	Float
	Gear	Int
	Torque (Nm)	Float
Algorithm Parameters	λ	Float
	μ_1	Float
	μ_2	Float
	ζ	Float
	Horizon length (steps)	Int
	Velocity interval (m s^{-1})	Float
	Wind Velocity (m s^{-1})	Float

Table 4.2: Algorithm input arguments with data types

type are listed in Table 4.2. Structure data types are used to organise related variables into an easily accessible group, such as road data which would include distance along the road section, elevation, slope and speed limit at each step. The variables in Table 4.2 are of data type float or integer depending on what they represent, for instance the horizon length is an integer value as it represents the number of steps used which must be a whole number, similarly with gear number. Typical values for the horizon length and distance intervals are thirty steps and 50 m, respectively, although the selection of these values is investigated in subsection 4.8.2.

4.4.2 Road and Vehicle Initialisation

As the parameters of the vehicle model do not change during the operation of the algorithm and are required in a number of functions the important parameters such as gear ratio, and wheel radius are set as external constants that can be read from anywhere in the program. These fixed vehicle param-

Item	Type
Gear ratios	Float
Front surface area (m ²)	Float
Air density (kg m ⁻³)	Float
Wheel radius (m)	Float
Vehicle mass (kg)	Float
Coefficient of friction	Float

Table 4.3: Vehicles constants with data types.

eters are listed along with data type in Table 4.3. To improve the efficiency of the program some fixed values are precalculated for each gear such as the conversion factor from engine torque to propulsion force at the wheels as described by Equation 3.25 and Equation 3.26.

The road data is allocated to a structure type variable of a fixed size of 4 sets of 200 fields that are populated with position, elevation, slope and speed limit data as provided from the road database. The amount of road data supplied is dependent on the horizon length provided as an input argument and only the required number from the 200 available fields are populated. The 4 fields at each position are float values, of 4 Bytes, resulting in a memory requirement of 3.2KB.

4.4.3 Memory Allocation

Due to the safety critical nature and the high reliability demanded of automotive control software, the Motor Industry Software Reliability Association (MISRA) produced a software development guide for C language automotive software [147]. A set of mandatory and advisory rules are contained within the guide. One key aspect of the MISRA standard for the DP algorithm is that dynamic allocation of memory is prohibited. Dynamic memory allocation is the process of managing memory during the lifetime of the program, allocating space as and when required and can lead to memory leak and unexpected behaviour. Static memory allocation, in contrast, requires that memory is allocated at compile-time and remains fixed for the duration of

Structure	Item	Type
Search space, ss	Cost	Float
	Gear	Int
	Path	Int
	Fuel	Float
	Velocity	Float

Table 4.4: Search space variables with data types.

the program. The implication for the DP algorithm is that the search space size and resolution have to be fixed at compile time and cannot be changed after.

4.4.4 Search Space Initialisation

The search space is configured as a three-dimensional array of structures based on gear selection, position and velocity. Each node in this search space is represented by a structure type containing the information as detailed in subsection 3.4.1 and Table 4.4. The maximum size of the array is fixed at compile-time as per MISRA guidelines but the amount of the array that is used depends on the horizon length as defined by runtime arguments. To give an example the search space for a gear that has a velocity range of 15 m s^{-1} , a velocity interval of 1 m s^{-1} and 30 steps in the horizon would require 9KB of memory. On the Raspberry Pi hardware both float and integer C types are 4 bytes.

The absolute velocity limits for each gear are fixed and the search space upper and lower indices for each gear are set according to the velocity discretisation interval. The search space nodes are then initialised with zero values in each field, with the exception of the cost field which is allocated a value of sufficient magnitude to prohibit selection as a minimum. This value is a practical representation of an infinite cost for infeasible regions which will be replaced by a calculated cost if the node is reachable. The velocity of each node is set according to the minimum and maximum velocity for the given gear and the velocity interval as defined by the runtime argument of the same name.

For each gear an upper and lower boundary for velocities is initialised, this is updated during the running of the algorithm. At each step in the algorithm, for a given velocity the upper and lower boundaries for possible velocities in each gear can be updated based on the maximum positive and negative torque translating to a maximum and minimum velocity that can be achieved from the given initial velocity.

4.4.5 Cost Calculation

With the search space initialised the nodes of the first step can be considered for cost calculation. As the initial conditions are fixed the first step only has to consider transitions from this single initial node. The calculation process is repeated for each gear and all feasible velocities within that gear search space. To find the limits of feasible velocity the gradient and speed limit for the current step are required to be extracted from the road data. The process by which the upper and lower limits are calculated differs depending on whether a gear shift is required. This sequence is shown in Figure 4.5.

Initially the lowest gear is selected and is to be incremented until the highest gear calculations are completed. If the gear selected for calculation is different from the initial gear in step zero of the search space then a gear shift is required (indicated in the left branch of Figure 4.5). If the gear selected for calculation is the same as that in step zero then the upper and lower velocity limits for step one are calculated. This calculation uses the vehicle model, gradient and speed limit to ascertain the physically realisable velocities. Starting from the lowest realisable velocity and incrementing until the highest, the transition cost is calculated and the search space updated.

The transition cost function uses the road slope, start and end velocity, gear, torque and step distance. The output of this function is a structure of the format described in Table 4.5.

4.4.6 Terminal Cost Calculation

Once each of the feasible nodes has been assessed an optimal path can be traced back from the lowest cost at the final step of the search space. Due to the rolling horizon concept that calls for repeated calculation of the optimal

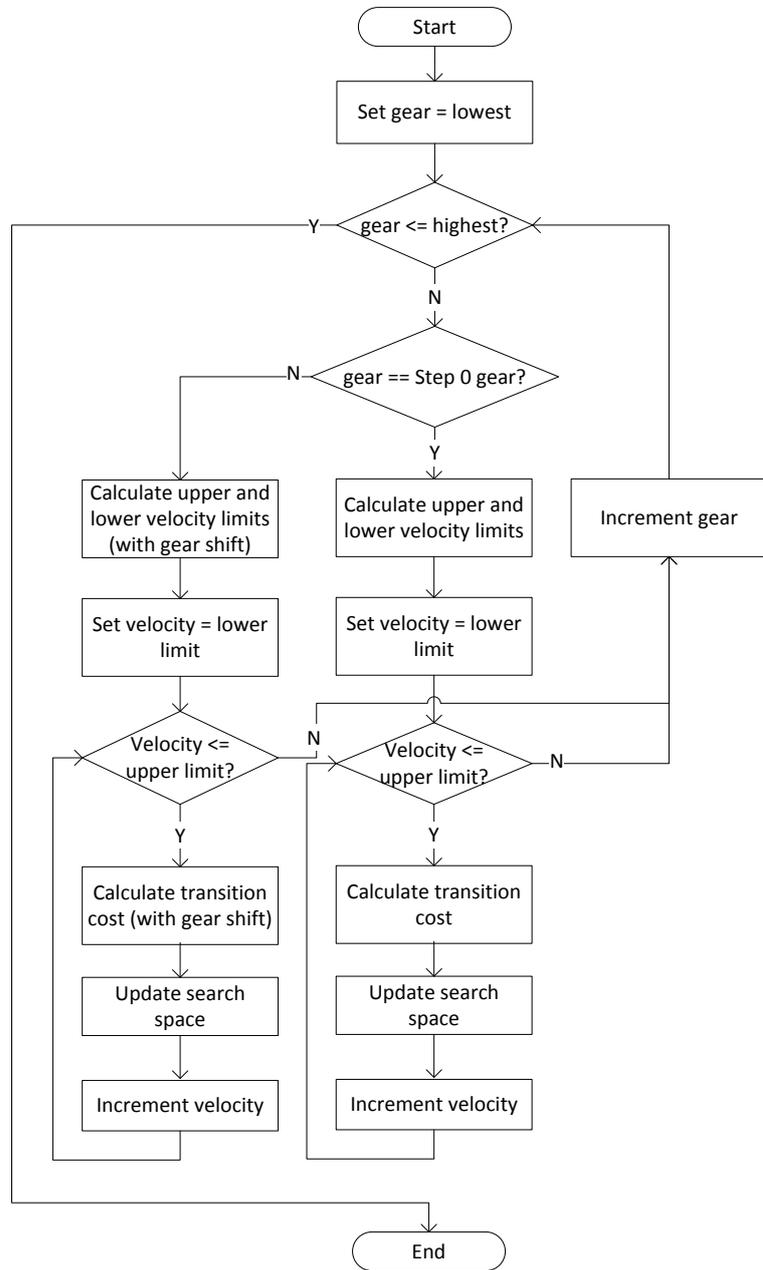


Figure 4.5: Calculation of cost for the initial step in the search space

Structure	Item	Type
Cost	Cost	Float
	Fuel	Float
	Time	Float

Table 4.5: Cost variable with data types.

control policy as the vehicle traverses the road, the vehicle state at the end of one horizon is the starting condition for the next horizon, as discussed in section 3.4. To balance the cost of the current horizon with the deferred cost applied to future horizons the terminal cost is calculated at each node of the final step.

4.4.7 Optimal Path Search

With the terminal costs allotted to each final node in the search space the minimum cost can be found by searching each final node on a per gear basis between the upper and lower velocity limits of the gear. The optimal velocity and gear profile can then traced back from this node through the search space. As each node contains an index of the velocity and gear from which the current node was reached this enables the optimal profile to be reconstructed.

4.5 Data Acquisition

To deploy the algorithm in a vehicle to provide real time guidance the system relies on a number of data sources that need to be considered for both datalogging and updating the vehicle and road models to reflect the current situation.

4.5.1 Controller Area Network

To ensure the DP algorithm provides useable and relevant results it is vital that the current status of the vehicle is known at the algorithm start. The current velocity and gear selection are necessary for the algorithm. This

data is being transmitted on the internal communication network of the vehicle, the Controller Area Network (CAN) bus. In order to communicate on this network a hardware interface is required between the Raspberry Pi and the vehicle On Board Diagnostics (OBD) port which has a connection to the CAN bus. Such an interface board is commercially available as the SK Pang PiCAN2 CAN Bus Interface Board [148], as seen in the right of Figure 4.6 and an existing Python library [137] can be utilised to access the data on the network. The CAN Bus consists of two signal wires, CAN HIGH and CAN LOW which are found on pin 6 and pin 14 respectively on the test vehicle OBD port and pin 3 and pin 5, respectively, on the PiCAN2 board. A custom interface cable was fabricated for this project to account for this.

Once the Raspberry Pi has been configured to communicate with the PiCAN board using the General Purpose IO connector (GPIO) and the device is set up as a network interface the Python Library is able to access the CAN bus. All control modules in the vehicle, of which there are commonly more than 100 [149], communicate on the CAN bus resulting in a large amount of data passing through the bus network. The CAN specification [150] details a standard format for messages that includes an 11-bit identifier to allow filtering so that only the messages required by the optimisation system are read from the CAN bus. Using manufacturer specific message identifiers the required data can be extracted in this way.

Once retrieved the required messages have to be converted so the relevant data is in a readable format in the Python script and is subsequently logged as well as used to update the current vehicle status. Often several pieces of related information are encoded in one CAN message to minimise traffic and the relevant information is to be decoded. Each message is converted to binary from hexadecimal and the location of the relevant data is provided as the starting bit position and the number of bits that make up the relevant data. This location data along with the message IDs are confidential manufacturer information.

The process of extracting the relevant data was developed using a Microsoft Windows based CAN bus logging and replaying software package, BUSMASTER [151] which as shown in Figure 4.7 is able to replay previ-

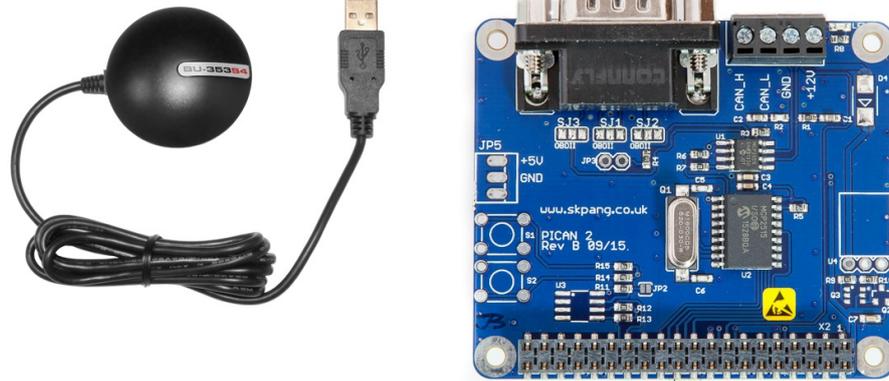


Figure 4.6: GPS and CAN interface board

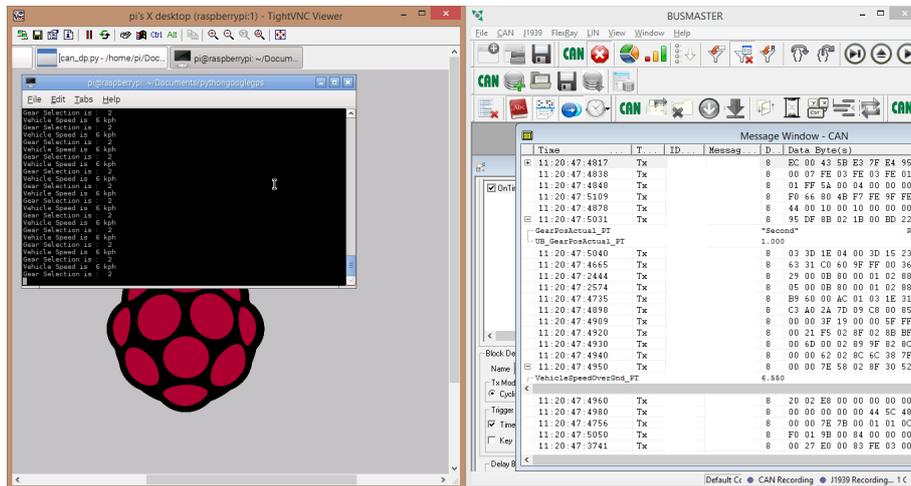


Figure 4.7: CAN Bus network simulation

Item	Units
Vehicle Speed	(km h ⁻¹)
Engine Speed	(RPM)
Engine Torque	(N m)
Fuel Consumption	(ml)
Gear Status	(-)
Engine Coolant Temperature	(°C)

Table 4.6: Vehicle data from received from CAN bus.

ously recorded CAN bus messages to simulate the network traffic of the particular vehicle from which CAN bus data has been recorded. Using a USB to CAN converter from IXXAT [152], the Windows PC running BUS-MASTER was configured to supply CAN data to the Raspberry Pi exactly as if it was connected to the real vehicle, which proved invaluable for testing the system without requiring constant access to a vehicle. The data required by the feedback system are listed in Table 4.6 and they are logged for later analysis while being used by the system to update the algorithm and GUI.

4.5.2 Vehicle Position

In order to locate the vehicle within the road section database a GPS receiver is used to measure the vehicle's longitude and latitude as described in section 4.3. With the vehicle coordinates known the road data is interrogated to identify the road section coordinates closest to the current vehicle position using a brute-force search approach.

4.5.3 Real Time Traffic Data

As the traffic conditions play an important role in the ability to follow an optimal velocity profile, a method for receiving real time traffic data in the vehicle is required. A ZTE MF823 4G USB Modem is used to establish a mobile internet connection which allows access to various internet services that provide traffic data. With the internet connection established, traffic data is acquired from the Here routing service [146] as described in section 4.3.3.

Along with the legal speed limit, traffic speed data, is provided where available and this is extracted from the *TrafficSpeed* data which contains a speed in m s^{-1} .

4.6 Algorithm Implementation

When all the current data required by the DP algorithm has been made available in the Python script it is passed as a number of arguments to DP C program for execution. The current vehicle speed, gear and engine torque are provided along with the weighting factors used in Equation 3.14. The current position in the road data is also provided along with a road identifier to allow retrieval of the relevant road section from a database of road data. This database is developed to replicate the data fields from the advanced driver assistance interface standard (ADASIS) that are relevant for the DP algorithm and would be provided by a commercial eHorizon system but without the cost of such a system. Using a combination of GPS data recorded from previously driven routes with speed limit data from the Here routing service mentioned above, the road database is constructed. Using the current vehicle longitude and latitude coordinate pair the distance to each data point of the road can be calculated using the haversine formula [143] used in navigation. The nearest point of road data to the current position is used as the starting point of horizon data.

With the horizon data and current vehicle status provided to the DP algorithm, the optimisation can begin. A forward DP algorithm is implemented that calculates the transition cost for each step in the horizon from the current position to the end of the horizon. In order to reduce the computational load, an assessment is made at each step of the physically realisable states possible in the next step and only feasible transition costs are calculated. The physically realisable states are assessed based on the maximum traction force available at the wheels that can be generated at the current engine speed and the maximum braking torque each of which are used in Equation 3.21 to find a maximum and minimum velocity, respectively. As the algorithm is to be used repeatedly with a receding horizon, the cost function includes a terminal cost to penalise velocity trajectories that bene-

fit the current horizon at the cost of future horizons, for instance by reducing the velocity drastically at the end of the horizon and thus requiring a high acceleration at the beginning of the following horizon.

4.7 User Interface

To test the algorithm in a vehicle it was required to provide an interface that would present information to both test personnel and the test driver. The operating status of the system components was required to be presented for the test personnel without interfering with the driver feedback system. Driving feedback had to be presented to the driver in the most straightforward way to ensure the feedback could be followed without providing too much distraction. A combination of audible and visual feedback was implemented allowing driver preference to enable either or both of the elements. Audible feedback was deemed to be less distracting to the driver compared with the visual feedback however the frequency of audible commands provides a problem to balance regular feedback with negatively impacting the driving experience, with constant computer generated voice commands. The visual feedback system was designed as part of this project, such that the relevant driver information is easily and quickly recognised without distracting the attention of the driver. The design of the visual feedback is for testing purposes only and due safety requirements would need to be considered in a production system.

On completion of the DP algorithm the initial portion of optimal velocity profile is compared to the current velocity and the driver feedback is developed from this to advise either maintaining, increasing or decreasing speed to follow the optimal profile. The driver feedback is provided by means of a graphical user interface (GUI) developed as part of this work using the Python library Kivy [153] and deployed on a 7 inch touchscreen display [154]. An icon representing each piece of advice is displayed on the GUI, as shown on the right of Figure 4.8, along with a voice command generated from a text to speech library [155] repeating this advice. Also shown on the GUI are the current speed and gear as well the advised speed for the upcoming step. The icons representing the advice based on current speed

are shown at the bottom of Figure 4.8, with advice to increase speed (left), maintain speed (middle) and reduce speed (right). Buttons are provided to select the road data used for testing and exit the GUI.

The design of GUIs is increasingly an area of detailed research under the umbrella term user experience and therefore to limit the scope of this project it does not include developing a GUI beyond a sufficient proof of concept as shown here.

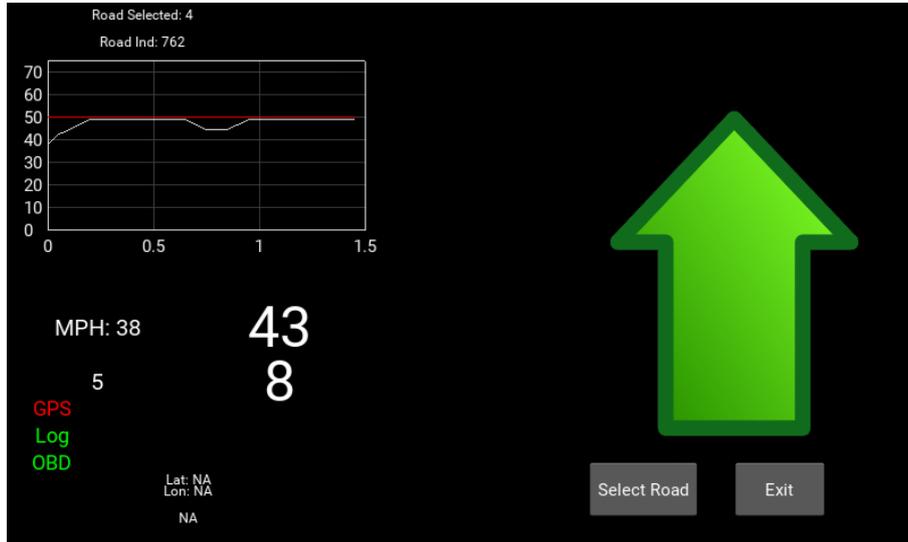
4.8 Algorithm and Model Configuration

A number of parameters are required to be set in the DP algorithm and vehicle model prior their implementation in either offline simulation or testing and in order to identify the correct values for such parameters an investigation is necessary into the influence of each of them.

4.8.1 Model Configuration

Vehicle Mass

As shown in Equation 3.21 the mass is an integral component in the calculation of the force required to maintain or change vehicle velocity. While the publicised kerb weight of the vehicle is used in the vehicle model by default it is likely that the mass of the vehicle will change depending on the amount of luggage and/or passengers. In order to assess the impact of such changes a simulation was undertaken of the vehicle combined resistive forces on a flat road while maintaining speeds from 10 m s^{-1} to 30 m s^{-1} . The results are shown in the first plot of Figure 4.9 with vehicle mass increasing from 0% to +20% of the manufacturer specified mass noted in appendix A. The maximum of this range equates to approximately 3 adult passengers (average 75 kg) and luggage in addition to the driver and corresponds to a 6.6% and 15.8% increase in resistive force at 10 m s^{-1} and 30 m s^{-1} , respectively. This rises to 14% and 19%, respectively when the road gradient is +5%, the increase at 10 m s^{-1} being more pronounced due the proportion of the resistive force due to the mass is greater at low speeds, as the other components of the resistive force are functions of vehicle speed.



(a) Driver feedback GUI



(b) Feedback icons

Figure 4.8: (a) Driver feedback GUI with guidance to increase speed based on current speed, in mph (38) and aim speed, in mph (43) with current gear (5) and guidance gear (8). Status of GPS, datalogger (Log) and vehicle data connection (OBD) shown for testing purposes. Optimal Velocity profile shown on plot with speed limit for current road selection and position. (b) Icons representing driver advice for increase speed (left), maintain speed (middle) and decrease speed (right).

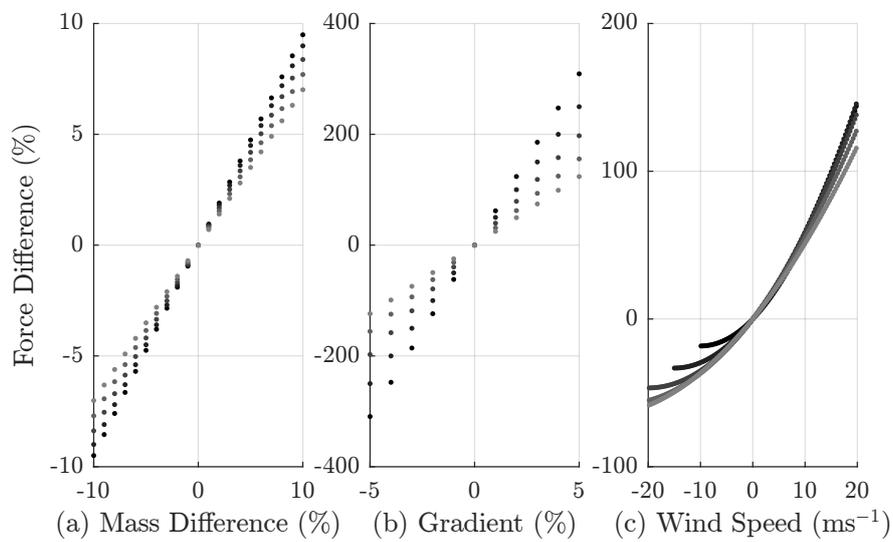


Figure 4.9: Vehicle model resistive force sensitivity to change in mass (left), gradient (middle) and wind speed (right). The difference in total force resulting from a change from default to maximum and minimum of each variable is plotted for vehicle speeds from 10 m s^{-1} to 30 m s^{-1} in steps of 5 m s^{-1} . The range of values are 0 to 20% for mass, +5% to -5% for gradient and maximum headwind to maximum tailwind of 30 m s^{-1} .

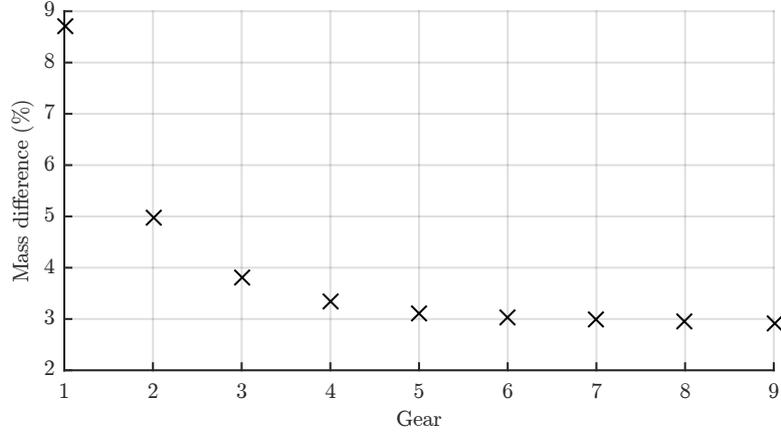


Figure 4.10: Inertial equivalent mass for each gear as a percentage increase on the static vehicle mass

As the mass used in Equation 3.21 includes the inertial effect of the engine, transmission and wheels and the inertia of the transmission varies depending on the gear selection, the impact of this change on the vehicle model is required to be investigated. The inertial equivalent mass is calculated as

$$m_{eq} = m_v + \frac{4I_w}{r^2} + \frac{I_{tc}I_{tr}R_{dr}^2}{r^2} + \frac{I_e(R_{dr}R_{tr})^2}{r^2} \quad (4.1)$$

where I_w , I_{tc} , I_{tr} , I_e are the inertias of each wheel, torque converter, transmission and engine, respectively. The impact of this inertial equivalent mass is investigated in Figure 4.10 with the additional effective mass due to the inertia of each gear ratio compared to the static mass of the vehicle. At higher gears where the gear ratio is lower the equivalent mass due to inertia is reduced due to the transmission ratio in the numerator of Equation 4.1 and this is apparent in Figure 4.10. As it is not efficient to drive in lower gears other than when accelerating from stationary and for all gears above three, the changes in equivalent mass are within 0.5% of each other the impact of changes in inertia due to gear are not considered further in the model, with the additional calculation load outweighing the vehicle model accuracy improvements.

Environmental Variables

As discussed in chapter 3 the environment local to the vehicle can influence the accuracy of the vehicle model. Primarily the aerodynamics are influenced by air density as well as wind speed and direction. In order to quantify the influence of these elements a number of simulations were undertaken. Using the standard vehicle model applied to a flat road scenario with a fixed speed the forces to be overcome in order to maintain this constant speed can be calculated as per Equation 3.21. The aerodynamic component of the force calculation can be varied to model various air density scenarios and the effect on the overall fuel consumption can be found. Air density was found by using the ideal gas law

$$\rho = \frac{p}{R_{spec}T} \quad (4.2)$$

where p is the absolute pressure of the air (Pa), R_{spec} is the specific gas constant ($\text{J kg}^{-1} \text{K}^{-1}$) and T is the absolute temperature (K). A range of air density conditions were considered to cover the extreme conditions that could be present for a vehicle sold globally. Air temperatures from -20°C to 60°C and standard air pressure at elevations from 0 m to 2000 m above sea level were selected to cover the range of possible conditions. The effect on fuel consumption of changes in air density are shown in Figure 4.11 for increasing vehicle speeds from 10 m s^{-1} to 30 m s^{-1} where the road section is flat.

Vehicle Model Sensitivity

To compare the influence of the vehicle model parameters of mass, wind speed and gradient a sensitivity analysis is presented in Figure 4.9. It is seen that the resistive force of the vehicle model is most sensitive to changes in road gradient and this effect increases with vehicle speed. The influence of the wind speed is shown as the variation from 0 m s^{-1} to a maximum headwind and tailwind speed of 20 m s^{-1} (45 mph) which is considered to be a gale, at the upper range of that likely to occur under normal conditions. In the United Kingdom the average wind speed in 2015 was 9.4 knots (4.8 m s^{-1}) [156]. At 20 m s^{-1} headwind, the force required to maintain a

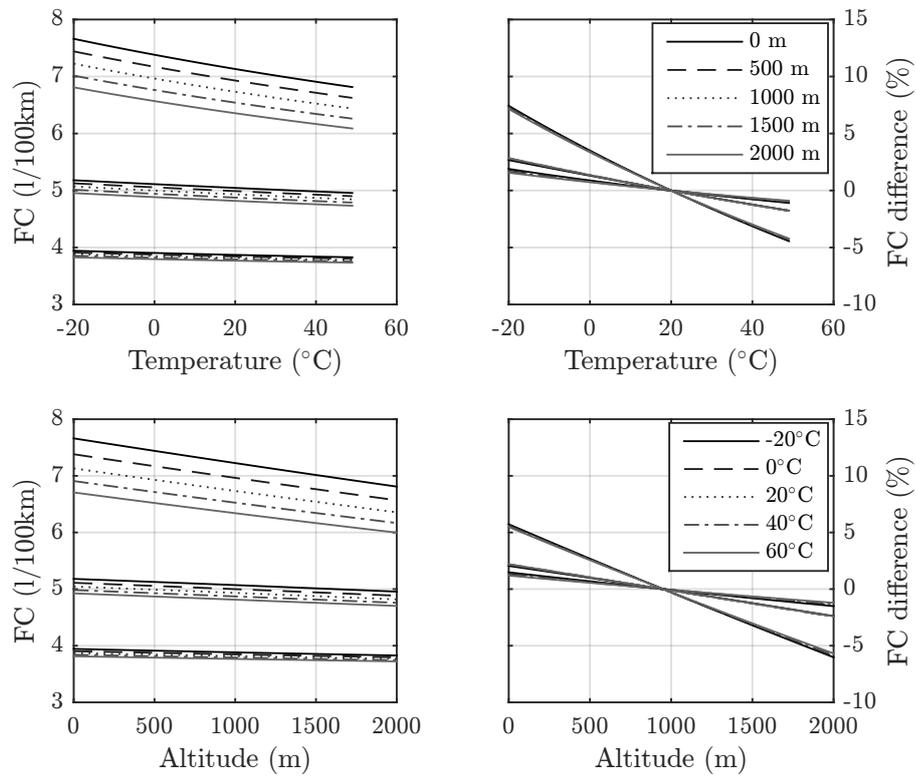


Figure 4.11: The effect on fuel consumption (FC) of aerodynamic variation due to air temperature (top) and air pressure and altitude (bottom), at varying vehicle speeds.

10 m s^{-1} vehicle speed is 145% higher than at zero headwind. To maintain a 30 m s^{-1} vehicle speed requires 115% more force than at zero headwind.

In the second plot of Figure 4.9 the impact of gradient changes are illustrated. At 10 m s^{-1} the force required to maintain velocity at +5% gradient is 123% that of the force required at 0% gradient. When the velocity being maintained is 30 m s^{-1} the difference is 309%. In the first plot of Figure 4.9 the vehicle mass is considered between 0% and +20% of the manufacturer specified mass noted in appendix A, but this has the least influence of the three parameters and is unlikely to change during the course of an individual journey.

This study identifies that the two most important factors in the vehicle model Equation 3.21 are the road gradient and wind speed. While map providers are increasingly aware of the necessity for improved road gradient data [119], the influence of wind speed is not so well investigated, despite its impact.

4.8.2 Algorithm Sensitivity Analysis

To test the robustness of the DP algorithm to changes in the algorithm parameters as well as external variables, a number of sensitivity analysis tests were undertaken. This approach involves running the algorithm multiple times with identical conditions except for one parameter or variable to see the effect that this alone has on the algorithm output. The three dimensions of the search space each have a range and a discretisation interval which can be adjusted depending on the performance requirements of the optimisation. The parameters of horizon length and discretisation interval impact the computation time and the optimisation results linearly as noted in subsection 3.4.1 and these parameters will be analysed first. A consideration that is integral to the horizon length is the frequency with which the algorithm is repeated, previously not considered in literature.

The range and discretisation of the velocity dimension is noted as exponentially related to algorithm complexity in subsection 3.4.1 and so requires careful consideration to balance optimality of results and computation time. The gear dimension is not considered to be a parameter as the range and

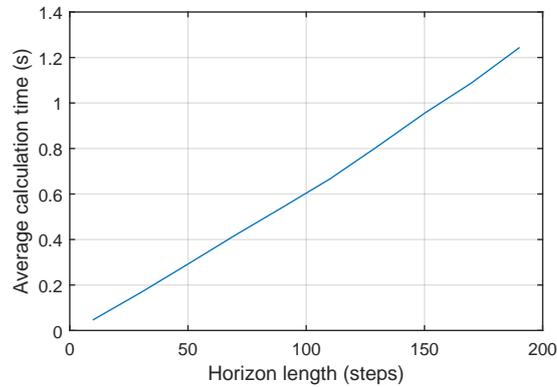


Figure 4.12: Calculation time as a function of horizon length with 1 m s^{-1} velocity discretisation interval

discretisation interval is fixed by the vehicle model. The sensitivity of the algorithm to the vehicle model parameters such as weight and BSFC map is then investigated.

Horizon Length

The horizon is the distance into the road ahead over which the optimisation is to operate. As noted in subsection 2.3.1 knowledge of the road ahead is an important factor in economical driving and with access to a database of an entire road network it is possible to be informed of the road data for an entire journey in advance, no matter the length. However, there are limitations on what is practical and with respect to computational load, extending a horizon will eventually lead to diminishing returns. The computational load of the algorithm increases as a function of search space size and as the horizon length is a fundamental dimension of the search space an increase in horizon length will lead to a greater computational load as shown in (3.37). This effect is shown in Figure 4.12 as the linear relationship between horizon length and average algorithm calculation time, as performed on an Intel i7-2600 CPU. There are two additional parameters that must be considered when selecting a horizon length, the frequency with which the algorithm is repeated and the formation of a terminal cost.

Frequency of Calculation

Due to the speed of calculation and the length of the horizon it is entirely feasible that the optimisation can be repeated several times over the course of one horizon and the results could be used to update the vehicle control. It is important that the calculation is completed in a reasonable time to account for changes in route or other changes to the horizon, however under normal operation it may be detrimental to update the control inputs too frequently. To test this hypothesis, multiple scenarios were simulated with varying frequencies of control strategy updating. The frequency of calculation is a function of distance intervals, for example the calculation could repeat every ten distance intervals. To compare the effect of varying the calculation frequency the total cost function Equation 3.14 is calculated for a number of calculation frequency intervals. In 4.13 four artificial road gradients with distance intervals of 50 m are considered with calculation frequency increasing from two to thirty in steps of two. The calculated cost is steady or decreasing in most cases as the recalculation interval is increased until intervals of between 20 and 24 steps after which the calculated cost increases. This trend is consistent across the four road gradients. The cost increases at higher recalculation frequencies are attributed to the optimisation only considering one horizon and not beyond. It may be optimal in one isolated horizon to conclude at a low vehicle velocity however this is detrimental to any horizons to be considered after, as in this case of two horizons of thirty steps. This highlights the need for a terminal cost as included in Equation 3.7.

When considered with a terminal cost the increase in cost at longer recalculation intervals is less pronounced as shown in 4.14 where a terminal cost calculated over 0.5 km is included in addition to a thirty step horizon. What is clear from these results is that at shorter recalculation intervals the cost is no better and in some cases worse than a slightly longer recalculation interval. This contrasts with the desire to produce an algorithm that can generate results as fast as possible, which is still necessary if an unexpected route change occurs and the optimal trajectory needs to be recalculated quickly. Outside of this exceptional situation the recalculation interval need

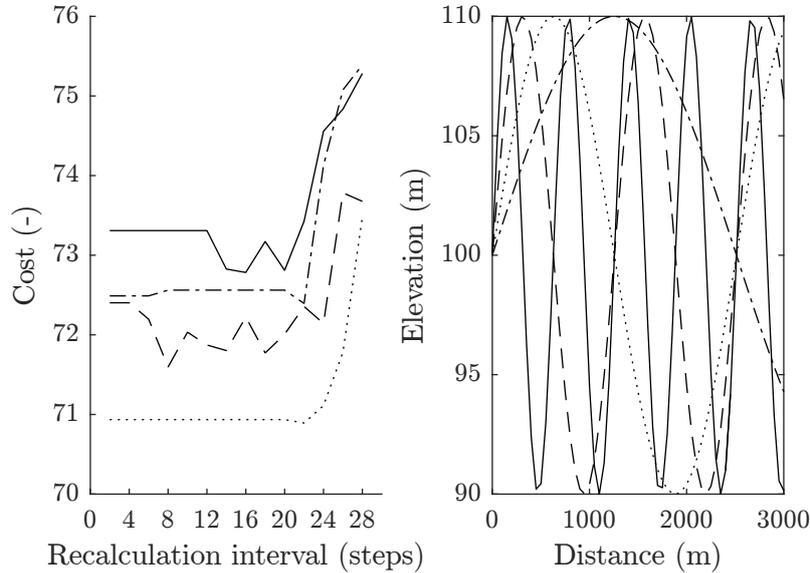


Figure 4.13: Total cost variation due to calculation frequency (left) for road profiles (right). A thirty step horizon is considered with no residual cost.

not be so short.

Velocity Quantisation

The size of quantisation interval with regard to the velocity variable has an exponential relationship with complexity as shown in (3.37) and illustrated in Figure 4.15.

In order to investigate the impact of increased algorithm complexity on the quality of results, a number of repetitions of the algorithm are undertaken with increasing velocity quantisation intervals and the computation time and overall cost function results are compared. A smaller quantisation interval allows a velocity profile to be generated that is closer to the absolute optimal profile, as shown with a reduction in the overall cost function as the interval reduces in Figure 4.16. One important consideration for the velocity quantisation is the application of the optimisation algorithm, for instance if the optimal profile is to be presented to a driver as real-time eco-guidance then the precision with which a driver could follow the optimal profile is limited by a number of factors including human physiology and the human

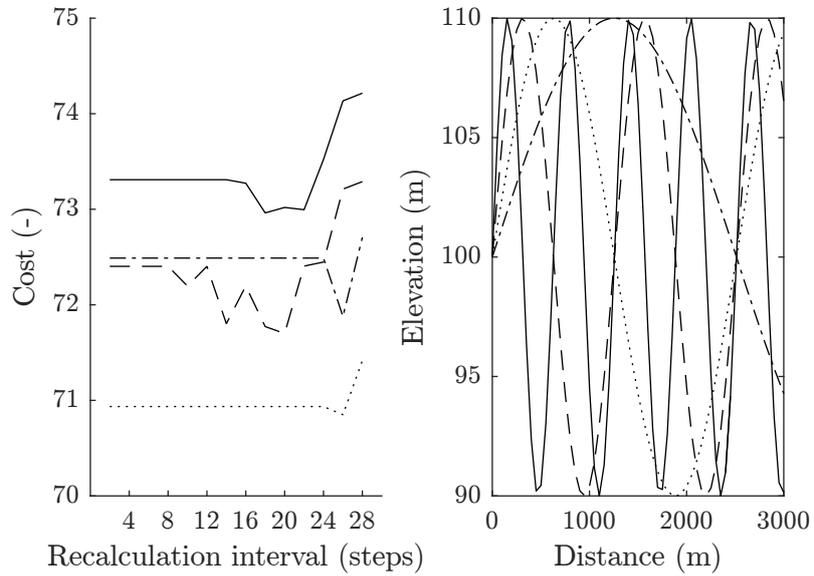


Figure 4.14: Total cost variation due to calculation frequency (left) for road profiles (right). A thirty step horizon is considered with an additional ten step residual cost.

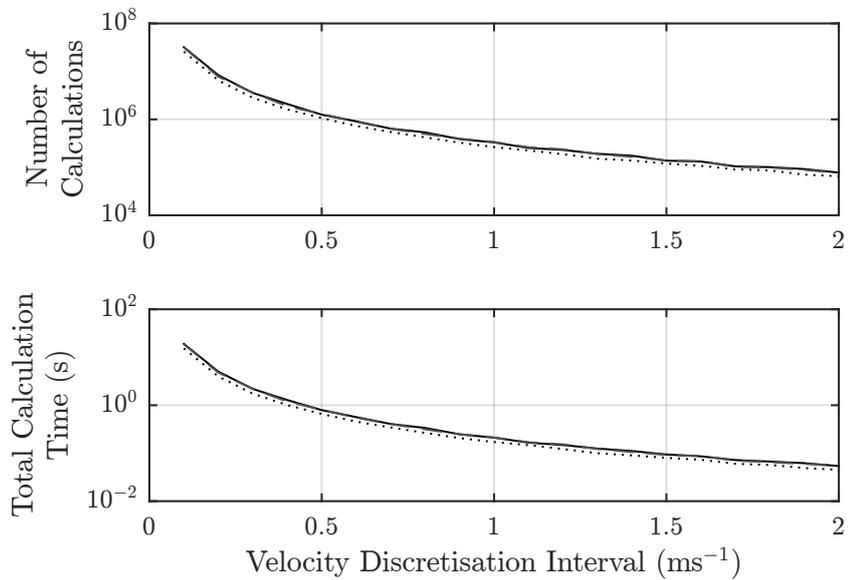


Figure 4.15: Computational load as a function of velocity discretisation interval, total number of calculations (top) for a fixed horizon and total calculation time (bottom)

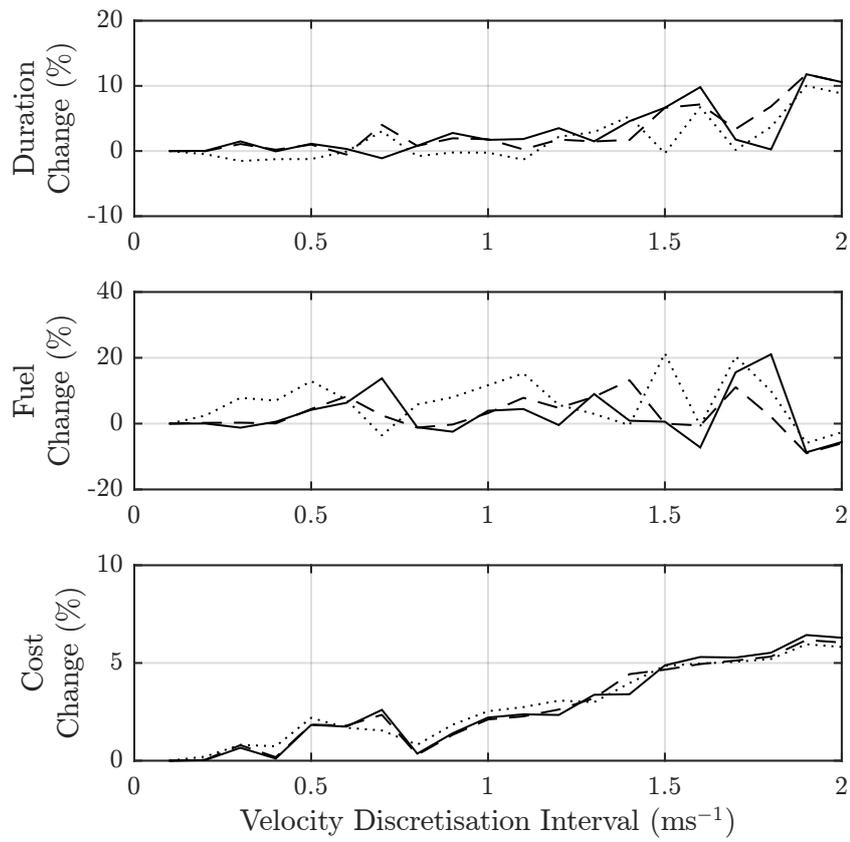


Figure 4.16: Horizon travel time change (top), fuel consumption change (middle) and total cost change (bottom) as a function of velocity discretisation interval.

machine interface. A fine velocity quantisation that is inevitably computationally expensive would be unnecessary as the driver could not exploit the higher resolution optimal profile because of the reasons mentioned. If however the optimal profile was to be used as a target velocity for an intelligent cruise control system which was capable of controlling at a much higher precision than the previous human driver example then the benefits of a fine velocity quantisation could be realised.

4.9 Algorithm Performance

The DP algorithm is capable of producing, for a given horizon length of road, a velocity and gear selection profile that is fuel and time optimal as per the cost function described in chapter 3. To verify that the algorithm has been faithfully implemented in the software code and the profile is optimal, the DP algorithm results are compared to those of an alternative off-line optimisation method which is described in the following section.

In order to assess the performance of the DP algorithm, a comparison between unaided real world driving and the DP algorithm results is not sufficient as many policies may lead to fuel consumption improvements compared to unaided driving despite being suboptimal. The improvement due to implementing the DP algorithm is to be assessed against an alternative optimisation method with a comparable implementation structure. While DP is a proven optimisation approach and individual profile costs can be calculated manually the scale of the search space means that the program developed to implement the DP algorithm is not easily tested step by step. Handling a large number of memory locations with repeated storage and retrieval operations ensures debugging of the software is an intensive process. Results from an alternative approach to the same problem will give some context to the DP results and highlight if the algorithm is not performing adequately. For comparing velocity and gear profile, fuel and time, a Genetic Algorithm was used as the alternative optimisation method.

To ensure a fair comparison the same cost function as in (3.1) is used for this alternative optimisation method. The velocity and distance discretisation intervals are fixed for both algorithm implementations.

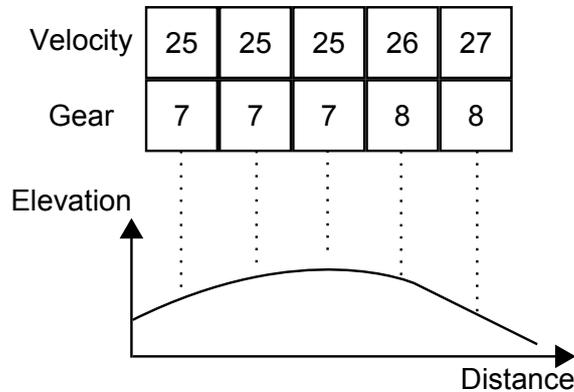


Figure 4.17: Genetic representation used in GA, consisting of a sequence of pairs of velocity values (m s^{-1}) and gear selections each corresponding to a position in the road data.

4.9.1 Genetic Algorithm

To include gear selection as an optimisation variable an approach is required that can handle the integer nature of the gear choice. A Genetic Algorithm (GA) can be used for this application, due to its flexibility with regard to the structure of problems it can be used to solve. As noted in subsection 2.4.5 a GA requires a genetic representation of the system to be optimised and a fitness function to assess each genetic representation. The genetic representation allows potential solutions to be presented to the GA for processing and in this case contains both the velocity and gear selection choices at each step of the considered horizon. These genetic representations are known as individuals in the GA. An initial population is created with a set number of individuals whose representations are randomly allocated. Each individual is then assessed using a fitness function and can be ranked amongst the other individuals in the population prior to creation of the next generation. The fitness function calculates the cost for each transition in an individual profile using the same method as the DP algorithm. The sum of these costs represents the fitness. The makeup of an individual for the optimisation problem explored here is shown in Figure 4.17 with a sequence of velocity and gear pairings that relate to a specific distance in the optimisation horizon.

The process by which subsequent generations are produced is the core of the GA and there are a number of possible operations and parameters to control the process.

4.9.2 Genetic Algorithm Parameters

The methods used to generate a further generation of a population are based on evolutionary concepts and are chosen for use depending on the desired behaviour of the algorithm. A crossover operation takes two parent individuals and produces a child individual based on a combination of the parent genes. There are a number of ways in which this combination can be generated. A single point crossover function takes a section of one parent up to a certain point and a section of another parent from the same point to the end. This process is illustrated in Figure 4.18 with the single point placed after the second element. Following the crossover operation a mutation can be applied to the child individuals to replicate the gene mutations that occur in evolution. With a given probability a child individual will be selected for mutation and each element of the individual can be subject to mutation. The probability of a particular element of the individual being mutate is fixed by a GA parameter. If selected for mutation the element is replaced by a random value in a fixed range of either velocity values or gear options depending on the element. A mutation is illustrated in Figure 4.18, with the third gear element mutating in child (d). Due to the random nature of mutation it is possible that the individuals produced by this process violate conditions of the system and so a feasibility test is required to ensure that the modified individuals conform to the system restrictions such as frequency of gear shifting.

The population size has a direct effect on the calculation time for each generation and so should be considered as a compromise between calculation time per generation and the total number of generations to achieve suitable results. A vital consideration then is what constitutes a suitable result and hence when to stop the process. A fixed number of generations or a simple time limit can be chosen as the stopping criteria, however due to the random nature of mutations and variable initial conditions there is no

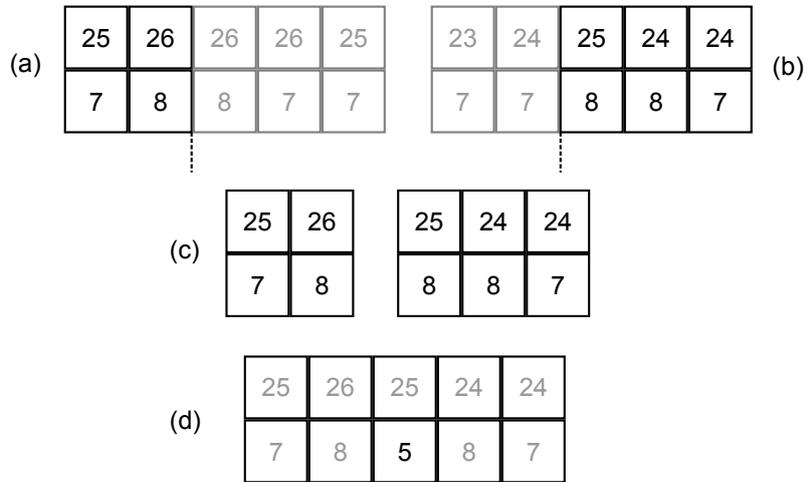


Figure 4.18: Genetic algorithm using single point crossover of parent (a) and parent (b) to produce child (c) and subsequently applying a mutation to the gear at the third step to produce child (d). Velocity shown as top row, gear selection as bottom row.

guarantee that a suitable result will be achieved for all problems in a fixed number of generations. A threshold for the fitness function can be employed so only when a result achieves the threshold does the algorithm terminate, however, this requires some prior knowledge of suitable solutions and will also limit the results to only achieve the threshold and no better. A stall test can be employed that looks at the change in fitness of the best individual in a generation and for how many generations the current best fit has been present.

Implementation

The GA is implemented using a Python script that utilises the Distributed Evolutionary Algorithms in Python (DEAP) package [157] with numerous modifications to suit the problem of this project. This open source framework provides functions to produce individuals and populations in a variety of formats that maintain compatibility with a suite of tools that implement operations such as a single point crossover. The fitness function is devel-

oped specifically for this project utilising a two dimensional array to store the genetic representation of the control policy and implementing the cost function calculation as in the DP algorithm. Similarly the mutation and feasibility functions are created specifically for this project. The feasibility function implements the constraints described by Equation 3.12 and Equation 3.13. The mutation function takes an individual and a probability of any single gene mutating and iterates through each gene in the given individual deciding if a mutation is to take place. This decision is based on the probability provided and a random number generator which on returning a decimal number in the range 0-1 that is less than the probability provided indicates that a mutation should take place. In this case a pair of random numbers in the feasible range for both velocity and gear are generated and replace the relevant genes in the individual.

Comparison

As described previously, the use of an alternative optimisation algorithm to compare with results of the DP algorithm is for verification purposes rather than as a viable alternative for real-time implementation. This is primarily due to the computation time of the GA algorithm, which due to the number of possible states under consideration will be much longer than DP. While the two algorithms both require the repeated calculation of fuel and time for given transitions, the DP only needs to calculate a specific transition once, whereas with each generation the GA could produce velocity and gear profiles that have been assessed previously as well as profiles that are physically impossible to achieve and are thus discarded. Other major disadvantages of GA over DP are due to the random nature of each generations creation which leads to unpredictable computation time required to reach a result that is comparable to DP, whereas a worst case computation time can be calculated for the DP algorithm and a result will always be produced within or below this time. Despite these disadvantages, purely for the purposes of verification of the DP algorithm results, a comparison of the results produced by the two algorithms are presented in section 5.5.

4.10 Summary

The development and implementation of the DP algorithm in both the simulation environment and hardware deployment have been detailed in this chapter. The individual software modules that make up the complete in-vehicle eco-guidance system are presented along with the interfaces between the modules. The challenges of implementing the DP algorithm are described along with the considerations of the algorithm parameters and how they influence the resulting optimal velocity and gear profile.

Chapter 5

Results

The algorithm being implemented in both a simulation environment and on in-vehicle hardware allows two types of investigation to take place. In the simulation environment the parameters of the algorithm and its sensitivity to varying conditions can be investigated in detail. The findings from these investigations can be applied to the in-vehicle implementation of the algorithm. The first section of this chapter details the validation of the vehicle model which is a necessary step prior to considering the algorithm results in any detail. The satisfactory implementation of the vehicle model is followed by the testing of the DP algorithm on artificial road profiles in order to verify that the algorithm is performing adequately under conditions with limited variation. Following this, testing of the algorithm with real road data is then presented along with a comparison to real driving data to identify the potential savings if the algorithm was deployed in place of or to assist a range of drivers. Consideration of the driving experience is then explored to see the effect this has on the impact of the algorithm.

5.1 Vehicle Model Validation

The success of the optimisation algorithm is reliant on the accuracy of the vehicle model it utilises. In order to assess the model accuracy, data from a real vehicle was used to compare to a number of parts of the vehicle model. The data used was recorded from a fleet of test vehicles fitted with

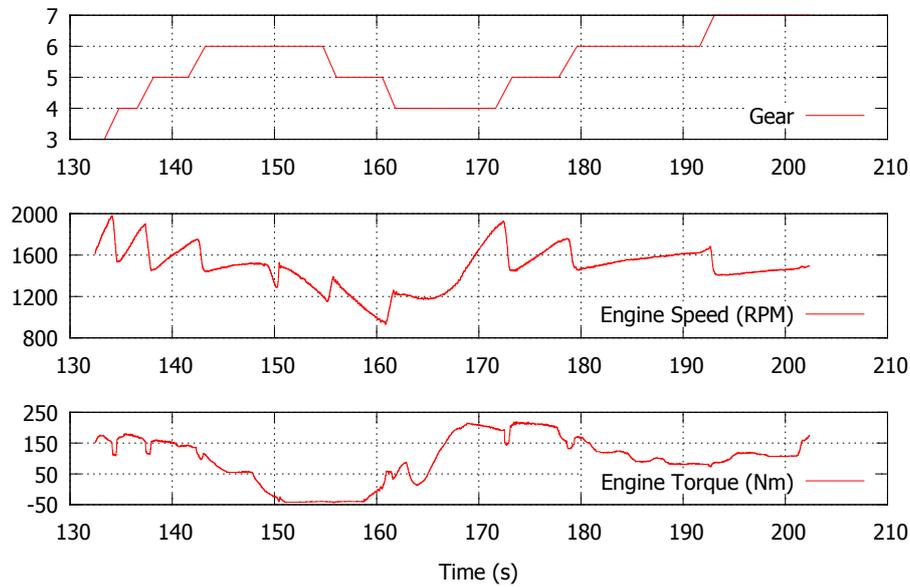


Figure 5.1: Engine speed (middle) and torque (bottom) behaviour during gearshift events (top)

commercial datalogging devices to record all the data that is transmitted on the CAN bus. The vehicle fleet data was provided by the project sponsor, Jaguar Land Rover. The first comparison was the engine speed to verify that the transmission model produced the correct engine speed when provided with the vehicle speed and gear selection from a real vehicle. Using the recorded gear selection, the gear ratio is known and can be used with the final drive ratio and wheel radius to calculate the engine speed from the vehicle speed as in (3.27). Due to the design of automatic transmissions (which the vehicle model was equipped with), the engine speed is decoupled from the transmission during a gear shift and as such the engine speed can no longer be estimated from the vehicle speed. This is illustrated in Figure 5.2 where the difference in calculated engine speed and recorded engine speed reaches an absolute error of 20% during two gear shift events. In order to resolve this problem while maintaining the accuracy of the model, a decoupled engine model is implemented during a gear shift that attempts to replicate this engine behaviour. It is observed that the engine speed increases during a downshift and decreases for an upshift, as shown in Figure 5.1.

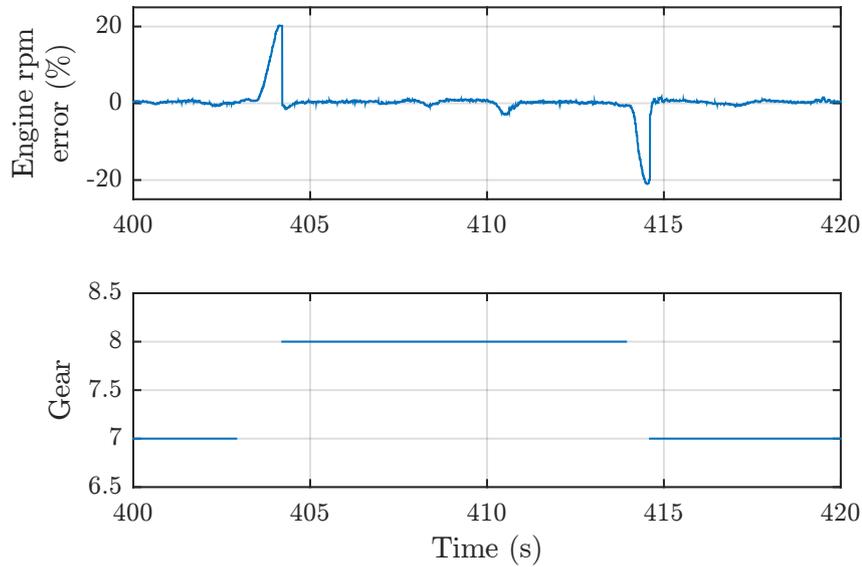


Figure 5.2: Engine RPM calculation error during gearshift from 7 to 8 and 8 to 7. Discontinuous sections are shown to indicate uncoupled status during a gearshift.

In order to control the engine speed in either decoupled scenario it is required that fuel is injected and this amount will vary based on the required post shift speed, the control of which is by the undisclosed algorithm deployed by the engine and transmission control unit (TCU). However in order to model this short time period for which the engine is decoupled an average fuel consumption is used.

Building on this calculated engine speed and using the recorded engine torque the fuel consumed is calculated using the BSFC map data as described in subsection 3.3.4 and compared to that recorded across a number of journeys. Due to the variations in engine efficiency that occur at different cylinder temperatures the BSFC data is measured under three sets of coolant temperature conditions, 30 °C, 60 °C and 90 °C. The engine coolant temperature is logged and after an initial warm-up period the coolant temperature is seen to settle around 90 °C under normal operating conditions as shown in Figure 5.3 and so only the 90 °C BSFC data is considered by default in the model. In order to model the engine fuel use accurately there

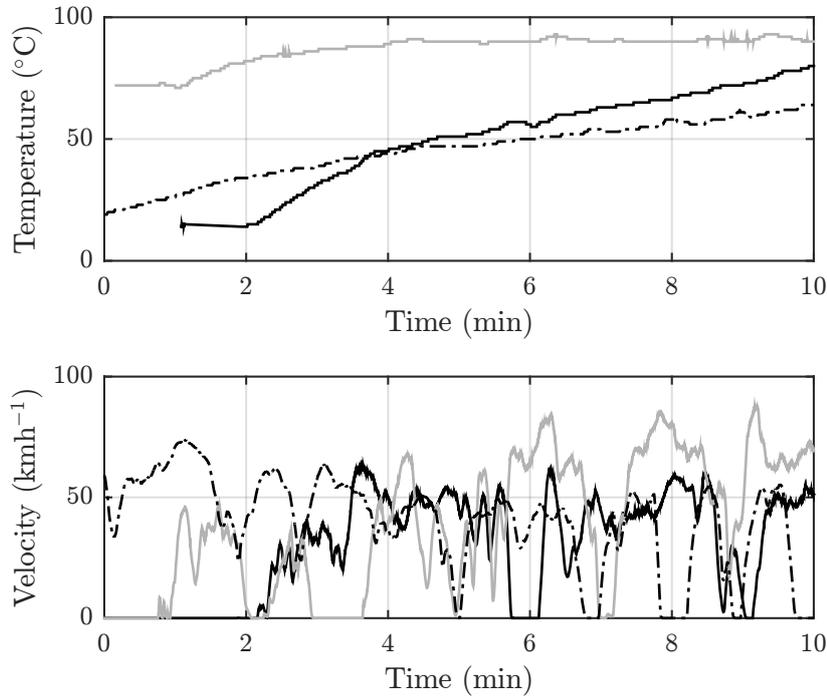


Figure 5.3: Engine coolant temperature over time (top) with three different velocity profiles (bottom)

are two situations that must be considered in addition to normal operation, the first is the idle situation when the engine is to continue rotating without any torque being required at the driveshaft and the second is the fuel cutoff situation where engine braking is desirable and so the resistive forces of the engine are used with no fuel being provided to the engine. Both of these situations require handling outside of the standard BSFC lookup operation.

The final step is to validate the vehicle model fuel consumption from only the vehicle speed and gear selection across a range of situations. Four test journeys are selected to compare the recorded vehicle data with the model fuel consumption produced when the recorded velocity and gear profile are followed on the same road section. The four journeys selected are detailed in section 5.3. As the fuel use is related to the torque required for propulsion, the resistive forces as noted in Equation 3.21 are required to be accurately modelled. The wind speed is shown in Figure 4.9 to be an important factor



R1 - M6 South from Wigan

R2 - M6 North from Cannock

Figure 5.4: Real road routes 1 and 2 as recorded in fleet vehicles.

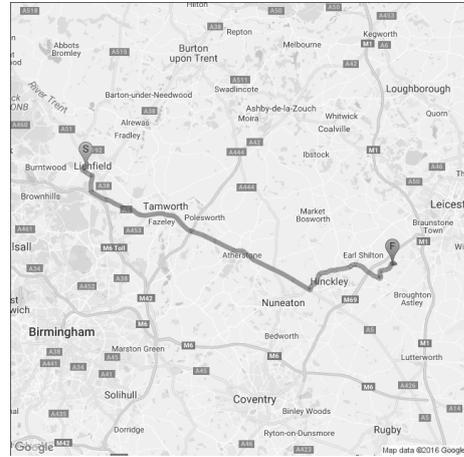
in the total resistive force seen by the vehicle and so wind data was required to be included in the vehicle despite not being recorded during the test journeys. Historical wind speed and direction data is available from the Met Office, an executive agency of the UK government, [158] for a number of observation points distributed across the country. Using an average of the observation data local to the current route, the wind speed influence on the vehicle fuel consumption can be estimated. For the roads R1-R4 headwind speeds of 1 m s^{-1} , -2 m s^{-1} , -5 m s^{-1} and -5 m s^{-1} were used, respectively, based on the component of the wind parallel to the average direction of travel for the route. As noted previously the average wind speed for the UK is 4.8 m s^{-1} [156]. The fuel use for sections of each road are shown in Figure 5.6 and the results listed in Table 5.1.

5.2 Artificial Road Profiles

To ensure that the algorithm is robust when applied to the variety of road scenarios that could be encountered in real driving, a range of road sections are used for testing. Initially artificial road profiles are created for this purpose before real road profiles are reconstructed from GPS data of real journeys.



R3 - A4177 Warwick to Solihull



R4 - A5 Lichfield to Leicester

Figure 5.5: Real road routes 3 and 4 as recorded in fleet vehicles.

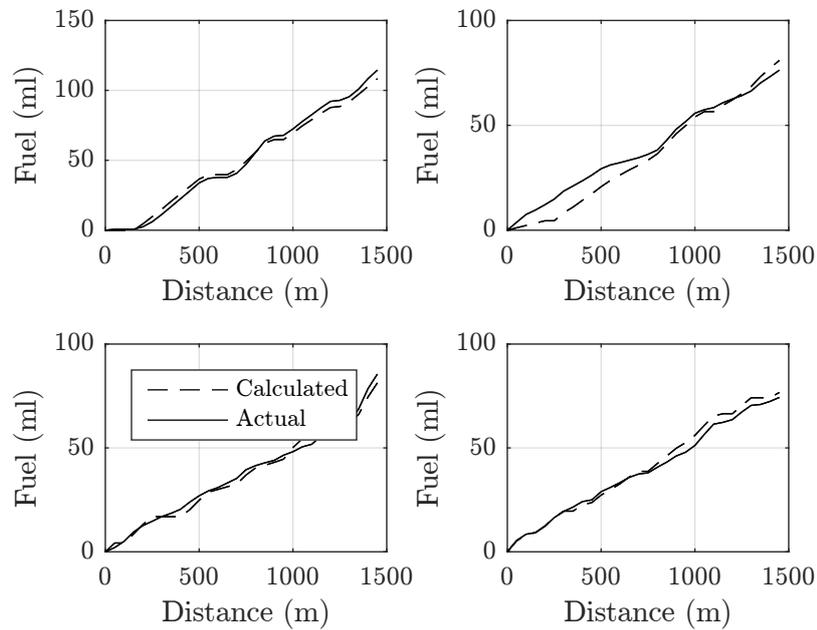


Figure 5.6: Fuel consumption for four road tests, R1 (top left), R2 (top right), R3 (bottom left) and R4 (bottom right)

Road	Simulated FC (1/100km)	Recorded FC (1/100km)	Difference (%)
R1	7.47	7.90	-5
R2	5.58	5.26	6
R3	5.61	5.90	-5
R4	5.291	5.12	3

Table 5.1: Vehicle model average fuel consumption (FC) validation for roads R1-R4 with simulated and recorded fuel consumption and the percentage difference.

The elevation of the set of artificial road profiles are shown in Figure 5.7. It is necessary to quantify road gradient variation so that a sufficient range of roads can be tested. To visualise the range of road gradients a normalised probability density function is shown in Figure 5.7 for the artificial roads.

By applying the vehicle model at a fixed velocity and gear to the road profiles in Figure 5.7 the fuel consumption and section time can be calculated. This fixed velocity strategy provides a baseline fuel consumption that is uninfluenced by driver variety and will be used to compare to the optimisation results. Each artificial road profile in Figure 5.7 is considered at fixed velocities between 20 and 30 m s⁻¹ in steps of 2 m s⁻¹ and the fuel consumptions and times are shown in Figure 5.8. The DP algorithm is applied to the same set of artificial road profiles with λ values between 0.2 and 0.8 to see the effect of biasing the cost function to fuel (0.2) and time (0.8). The results are also shown in Figure 5.9 and the relationship between λ value, fuel and time are highlighted by the Pareto curves formed between discrete λ values, shown in grey. The trends of the fixed velocity results are similar across the artificial roads with an almost linear relationship between fuel and time as any exploitation of the road topography for fuel saving is nullified by the fixed velocity requirement. At higher λ values which are heavily time biased the fuel consumption for the DP algorithm and high fixed velocity converge as the DP algorithm is so restricted by the time bias that few opportunities for fuel saving can be taken at the expense of time.

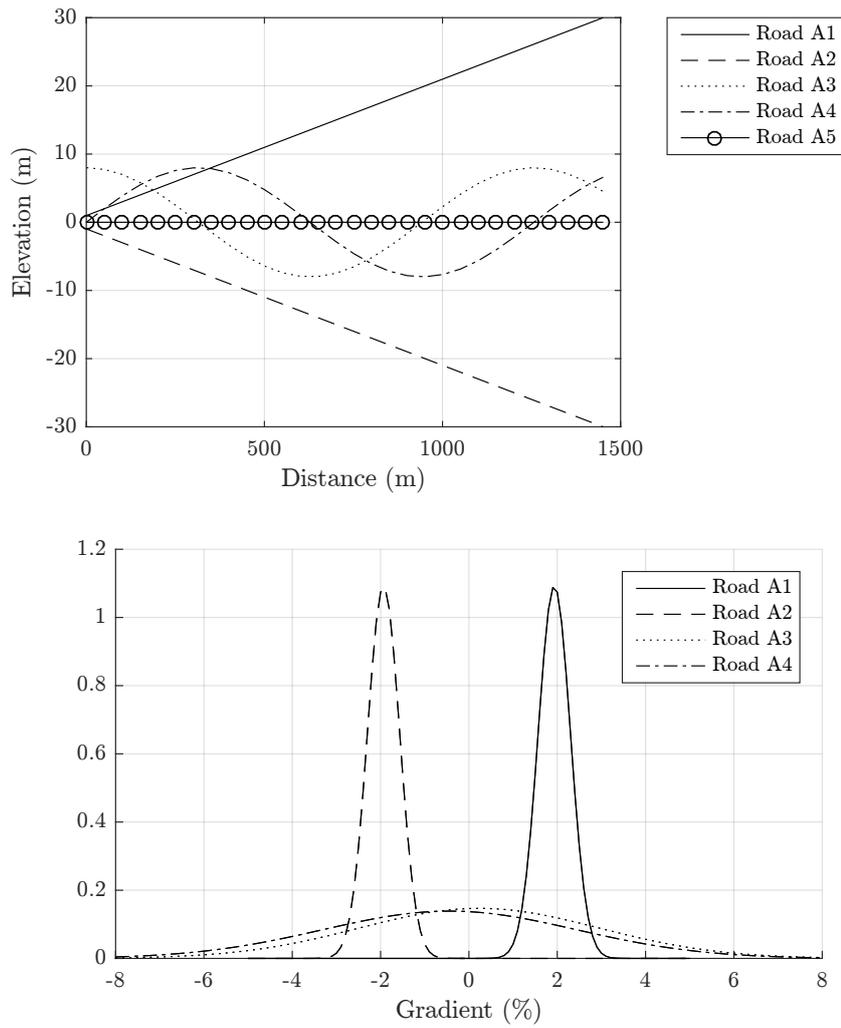


Figure 5.7: Artificial road elevations (top) and normal probability density function of gradients (bottom). Road A5 not shown as gradient is zero.

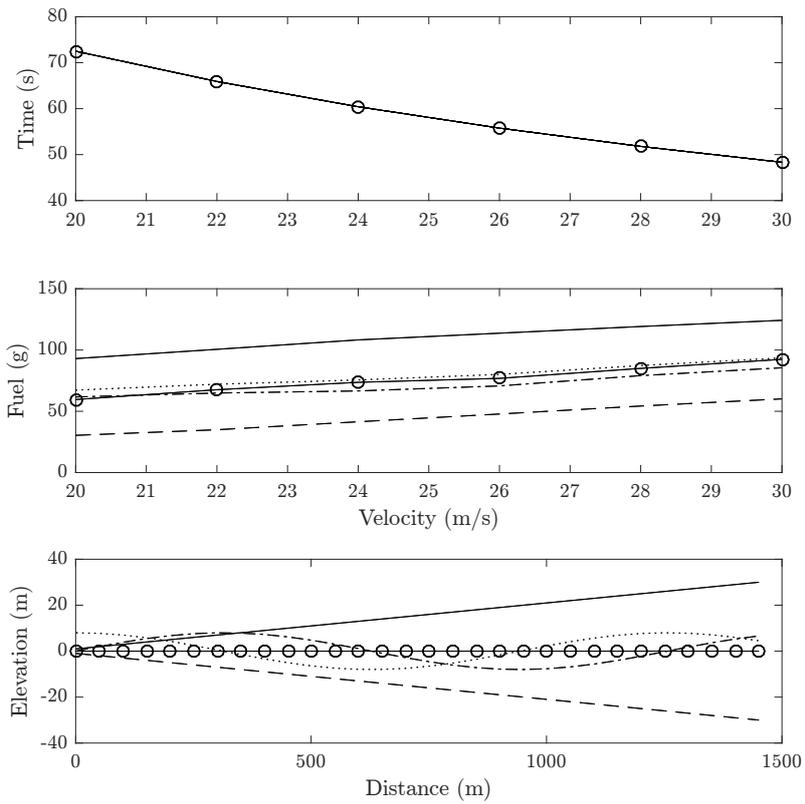


Figure 5.8: Artificial road elevations (bottom) with journey time (top) and fuel consumption (middle) as a function of vehicle velocity. The velocity is fixed for the duration of each road section along with the gear selection, which in this case is 7.

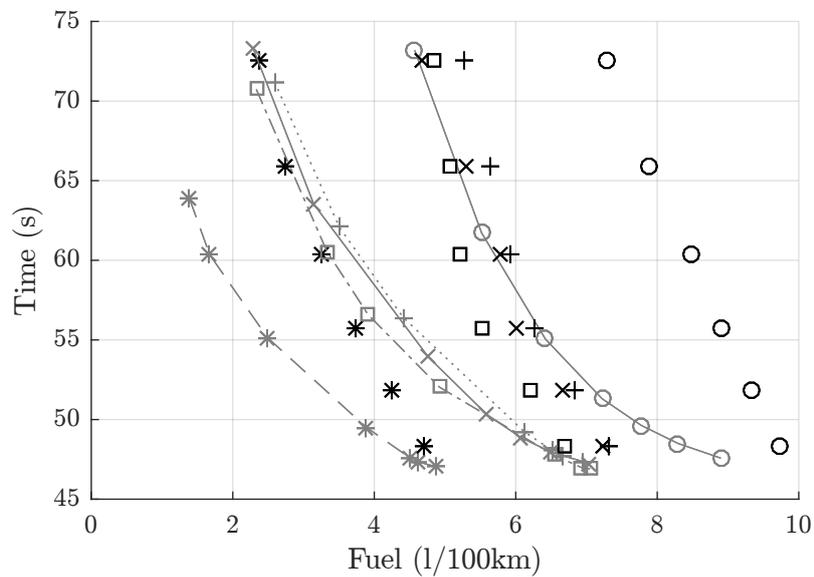


Figure 5.9: Algorithm and fixed velocity results for artificial road profiles with line type corresponding to roads A1-A5 as previously. Algorithm results for seven fuel/time weightings from $\lambda = 0.2$ to $\lambda = 0.8$ (grey markers) and fixed velocity results from 20 m s^{-1} to 30 m s^{-1} (black markers according to road)

5.3 Real Driving Data

To assess the impact of the optimisation algorithm in practice it is compared to actual driving data from recorded journeys. To achieve this a fleet vehicle database is used that contains data logged from journeys taken by a number of production vehicles of identical specification on long term testing. Using the GPS data recorded in the vehicle the route can be identified and the road reconstructed using longitude, latitude and elevation to give position and distance covered and legal speed limit data at defined intervals along the route. The GPS data includes an elevation measurement, however the accuracy of this is inconsistent and so has to be compared to an alternative source of elevation data to identify anomalies. Such elevation data can be acquired at specific locations using Google Maps Elevation API [159]. Following the processing of elevation data the gradient can be computed as per section 4.3. The routes of the journeys selected from the database are shown in Figure 5.4 and Figure 5.5 and represent a combination of motorway and A-road sections.

The variation in legal speed limit on these routes is shown in Figure 5.11 along with the variation in recorded driver speed for those specific journeys. It can be seen that road R2 which is almost entirely motorway is narrowly distributed around the speed limit for this type of road, 70 mph (112 km h^{-1}) whereas the other roads are more varied.

5.4 Algorithm Performance

In order to test hypothesis (a) detailed in section 1.1 which states

- (a) Fuel savings can be made by utilising optimal control methods to control vehicle speed and gear selection in real-time, based on instantaneous vehicle and road data,

it is necessary to compare the fuel consumption and time generated by following optimal profiles as produced by the DP algorithm to the equivalent generated by the current standard velocity control system found in vehicles, a fixed velocity cruise control system. Hypothesis (a) also requires that any

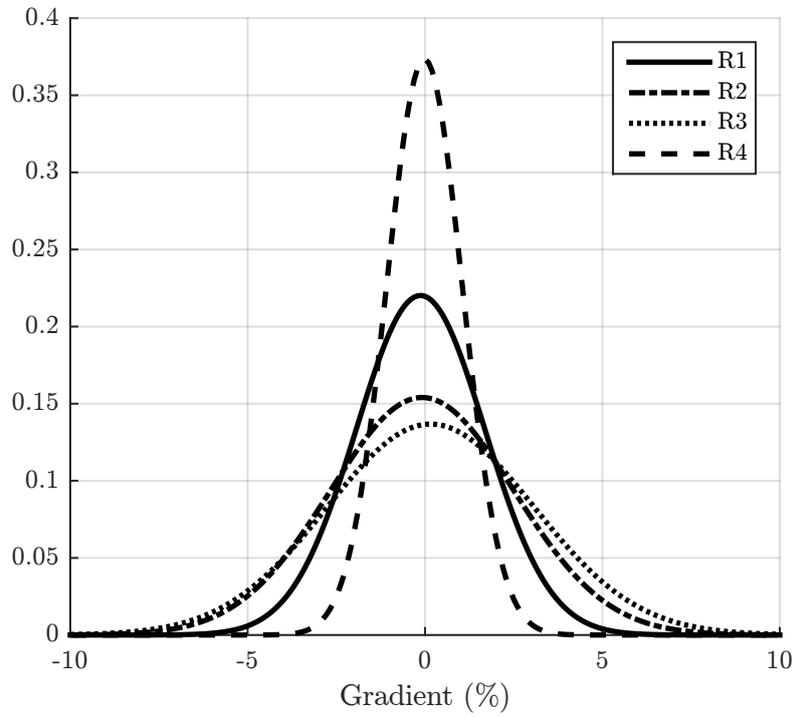


Figure 5.10: Road gradient normal probability density function for four real road profiles

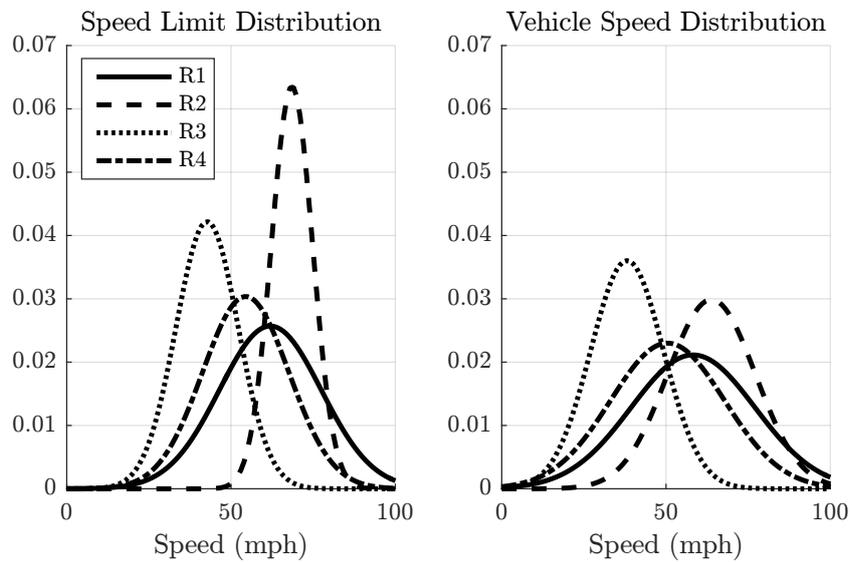


Figure 5.11: Road speed normal probability density function for four real road profiles with legal speed limit (left) and recorded driver speed (right).

fuel savings can be made using a real time system and as such the computation time of the DP algorithm must be reported on hardware that can be deployed in a test vehicle.

5.4.1 Fixed Velocity

The basic implementation of a cruise control system involved a feedback loop to maintain a fixed velocity regardless of gradient and road conditions [160]. This policy was implemented with a range of target velocities on the roads contained in the database to identify the performance of such a policy. The fuel consumption and time results are shown in Figure 5.12 at fixed velocities ranging from 20 m s^{-1} to 30 m s^{-1} on three of the artificial road profiles detailed in section 5.2. The DP algorithm is configured with a horizon of 1.5 km divided into thirty steps, a velocity interval of 1 m s^{-1} and the full range of the 9 speed transmission available where feasible. The cost function weightings are $\lambda = 0.5$ with normalisation factors $\mu_t = 1$ and $\mu_f = 4$.

The DP algorithm is implemented without constraints and with an initial velocity equal to each of the fixed velocities tested to examine the difference in results between the two control policies. As seen in Figure 5.12 the overall cost of the DP results is always lower than the equivalent fixed velocity policy, however there are certain situations where either the time or fuel is higher for the DP results. For instance on road A5, which is flat, maintaining a fixed 30 m s^{-1} leads to a lower time than the DP algorithm, however the fuel consumption in this case is halved resulting in a lower overall cost. Conversely on road A4 the fuel consumption is higher for the DP algorithm starting at 20 m s^{-1} however the time is 30% less than the fixed velocity. The cost reduction in each scenario is shown in Table 5.2 with an average cost of the fixed velocity profile 11.6% higher than that of the DP algorithm.

Real Road Profiles

The results of the fixed velocity policy and DP algorithm when applied to the four real road routes shown in Figure 5.4 and Figure 5.5 are listed in Table 5.3. The average cost reduction by the unconstrained DP algorithm

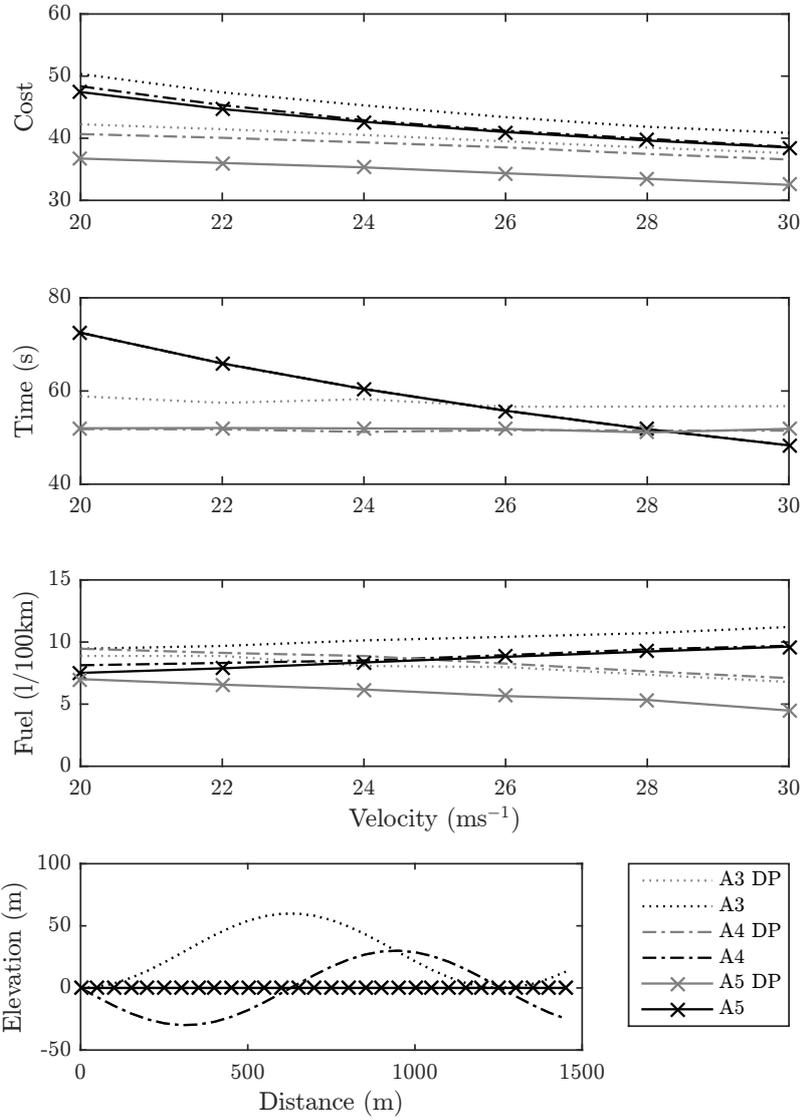


Figure 5.12: Fixed velocity and DP results roads A3-A5 with initial velocity from 20 m s⁻¹ to 30 m s⁻¹.

	A1			A2		
Velocity	Fixed	DP	%	Fixed	DP	%
20	47.9	39.3	17.9	40.0	33.2	17.0
22	45.5	38.5	15.5	37.3	32.5	12.9
24	43.7	37.7	13.8	35.4	31.7	10.4
26	42.1	36.9	12.4	33.9	30.9	8.8
28	40.8	36.0	11.8	32.7	30.0	8.1
30	39.7	35.0	11.8	31.7	29.1	8.1

	A3			A4			A5		
Velocity	Fixed	DP	%	Fixed	DP	%	Fixed	DP	%
20	44.7	36.7	17.9	44.0	36.1	17.9	43.7	36.2	17.1
22	42.0	35.9	14.5	41.1	35.4	13.9	41.4	35.4	14.4
24	39.7	35.1	11.4	38.5	34.6	10.2	39.4	34.7	12.0
26	37.9	34.3	9.6	36.7	33.8	8.1	37.5	33.9	9.6
28	36.8	33.4	9.2	35.8	32.9	8.0	36.5	33.0	9.6
30	35.9	32.4	9.6	34.9	32.0	8.2	35.7	32.0	10.3

Table 5.2: Total cost results from fixed velocity and unconstrained DP policies on artificial road profiles A1 to A5 with initial velocities from 20 m s^{-1} to 30 m s^{-1} . Total cost for horizon given for fixed velocity and DP and the percentage difference are shown for each velocity and road tested.

	R1			R2		
Velocity	Fixed	DP	%	Fixed	DP	%
20	44.1	36.2	17.9	43.4	35.9	17.3
22	41.7	35.4	15.0	41.0	35.1	14.6
24	39.7	34.6	12.7	39.2	34.3	12.5
26	38.2	33.8	11.5	37.3	33.5	10.2
28	37.2	32.9	11.6	36.0	32.6	9.3
30	36.3	32.0	12.0	35.2	31.7	10.0
	R3			R4		
Velocity	Fixed	DP	%	Fixed	DP	%
20	43.4	34.8	19.7	42.7	34.9	18.3
22	40.7	34.1	16.2	40.2	34.2	15.0
24	38.6	33.2	13.9	38.3	33.4	12.9
26	36.8	32.4	12.0	36.5	32.6	10.7
28	35.3	31.5	10.6	35.1	31.7	9.6
30	34.2	30.6	10.6	34.1	30.8	9.7

Table 5.3: Total cost results from fixed velocity and unconstrained DP policies on real road profiles R1 to R4 with initial velocities from 20 m s^{-1} to 30 m s^{-1} . The cost difference between the two policies in percent is shown in bold.

across the four real road profiles is 13.1%.

5.4.2 Speed Limit Following

An intelligent cruise control system that could be implemented to minimise journey time uses speed limit data to produce a velocity profile that follows as close as possible the legal speed limit. Due to the instantaneous transition from one speed limit to another which would be impossible for a vehicle to follow precisely, the transition between speed limits is accomplished with a gradient according to the limits of the vehicle performance. When the speed limit velocity profile is compared to the DP algorithm results as shown in Table 5.4 the difference in cost is much lower than in the fixed velocity

Road	Cost		
	Speed Limit	DP	
R1	35.7	33.9	5%
R2	36.5	36.1	1%
R3	43.1	43.6	-1%
R4	41.2	40.1	3%

Table 5.4: Total cost results from speed limit following and DP policies on real road profiles R1 to R4. The cost difference between the two policies, in percent, is shown in bold.

scenarios shown previously, in fact, for road section R3 the cost for the DP algorithm is actually higher than the speed limit following. As shown in Figure 5.11 the road sections vary in speed limits with road having the lowest average speed limit. For the section of R3 under examination here, the speed limit varies from 50 mph to 40 mph, while for the other road sections it is either 70 mph or 50 mph, so by following this lower speed limit the engine is being operated in an efficient speed range, thus limiting the potential fuel savings the DP algorithm could make in comparison. In addition to this the DP algorithm is prohibited from exceeding the speed limit and therefore cannot achieve any time savings compared to the speed limit following profile, in fact due to the conversion from mph to ms^{-1} producing a non integer value the speed limit following will always be capable of producing a better time. If the DP algorithm is unrestricted by the speed limit, the optimal velocity profile it provides is in fact much higher than the speed limit while simultaneously producing a lower overall cost, this can be attributed to the cost function heavily favouring a time reduction compared to an equivalent fuel reduction. This relationship between time and fuel is specified by a combination of the cost function weighting factor, λ and the normalisation factors μ_t and μ_f , which as noted in subsection 5.4.1 were $\lambda = 0.5$, $\mu_t = 1$ and $\mu_f = 4$. While the relationship produces suitable results at the higher speed road scenarios this example has highlighted its flaws at lower speeds.

		Cost		
Road	μ_f	Speed Limit	DP	
R1	2.1	46.9	41.4	12%
R2	1.9	49.4	42.8	13%
R3	0.7	70.0	66.3	5%
R4	1.0	65.8	61.0	7%

Table 5.5: Total cost results from speed limit following and normalised DP policies on real road profiles R1 to R4. The cost difference between the two policies, in percent, is shown in bold.

5.4.3 Normalisation

The ability to vary the optimisation cost function to favour fuel or time using just one weighting factor, λ to suit a particular driver’s requirements helps to ensure adoption of the eco guidance, however as highlighted above, depending on the speed limit of the road, the impact of the weighting factor varies. To ensure consistency across a range of road scenarios the cost function is required to be normalised based on the speed limit of the road. When the cost function is normalised for a specific road section prior to optimisation the results when compared to the speed limit following policy are much improved as seen in Table 5.5.

5.4.4 Computation Performance

While testing the optimisation algorithm it is necessary to consider its computation time, in order that it is suitable for deploying in a real-time application, as hypothesised in chapter 1. Due to the portability of the C code used to program the DP algorithm it can be run on a variety of platforms, with simulation work undertaken on a Windows based PC, however this is impractical for installation in a production vehicle and so the Raspberry Pi Miniature computer, as described previously, was used to test the algorithm prior to its installation in the test vehicle. Although the processing power of the Raspberry Pi is impressive for its size, it does not compare with an entry level PC and so the computation time of the algorithm will be much

lower than in the previous simulations.

The real road profiles described in section 5.3 are used to test the computation time of the DP algorithm on the Raspberry Pi with a horizon length of 30 steps, a velocity interval of 1 m s^{-1} and a starting velocity of 20 m s^{-1} to remain within the speed limits for all four real road profiles. The results are shown in Table 5.6 with the longest and shortest computation times of 2.13 s and 0.86 s for roads R1 and R3, respectively. This variation in computation time can be attributed to the higher speed limits in road R1 leading to an larger feasible search space requiring a greater number of calculations than the smaller search space of road R3 that is restricted by the upper limit on speed which is lower than in R1, as seen in Figure 5.11. Increasing the starting velocity to closer to the speed limit reduces the computation time marginally as there are fewer higher speeds initially available in the search space, however overall the effect is limited as the subsequent search space remains the same. The ultimate lower limits of the search space are dictated by the lowest speed in each gear that still results in a feasible engine speed, which is unaffected by the road conditions, therefore the key factor in search space size and thus computation time is the upper limit of velocity for the given road section.

As the DP algorithm complies with the legal speed limit, which in the UK is a maximum of 70 mph (31.3 m s^{-1}) the worst case computation time would occur for a road of maximum speed limit. The real-time requirements of the DP algorithm implementation are classed as firm real-time [161], as the consequence of not providing guidance in a suitable time are not catastrophic but limit the usefulness of the results. With a discrete distance interval of 50 m and a maximum speed of 31.3 m s^{-1} , for the result when calculated at the beginning of a distance interval to still be relevant when required at the start of the next interval it would need to be calculated in 1.6 s. This demand would only apply on initial calculation of a given horizon where no prior calculation had occurred, for instance during an unexpected re-routing. While it can be seen in Table 5.6 that the algorithm exceeds this threshold for two of the road sections, the relevance of the results would not be lost entirely during such an unexpected event as the result would be available within the next discretisation interval.

Road	Calculations	Time (s)
R1	350544	2.13
R2	308340	1.88
R3	127890	0.86
R4	140617	0.93

Table 5.6: Computation time of DP algorithm applied to real road profiles on Raspberry Pi 2 with total number of cost calculations and time taken, in seconds.

5.4.5 Driveability Consideration

As shown in subsection 5.4.1 potential fuel and time savings can be made on artificial road profiles with an average cost reduction of 11.6% and real road profiles of 13.1% however no consideration is made for the driving experience of following the calculated optimal velocity and gear profile produced by DP. Hypothesis (c) states

- (c) Fuel savings can be made as above, without compromising the driving experience and this can be verified across a range of real driving data.

In order to test this, it is necessary to compare these results to that of a modified algorithm that considers the driving experience. To ensure that any potential fuel and/or time savings can be realised by following the optimal profile, the driving experience must be acceptable to the driver and so constraints are applied to the DP algorithm to achieve this. Firstly the frequency of gear shift operations is restricted to occur less frequently than once every two distance intervals to minimise gear shift hunting behaviour that is detrimental to the driving experience. The other aspect of driving experience that provides a constraint on the algorithm is the rapid changing of torque supplied by the engine which leads to changes in acceleration that impact driver comfort. The cost function incorporates the component J_c to account for this as shown in Equation 3.14 and Equation 3.17. The cost reduction in the fixed velocity artificial road scenarios is shown in Table 5.7 with an average cost reduction of 9.3% compared to 11.6% when no consideration is given to driveability. This change of only 2.3% highlights that

	A1			A2		
Velocity	Fixed	DP	%	Fixed	DP	%
20	47.9	40.8	14.8	40.0	34.5	13.9
22	45.5	40.0	12.2	37.3	33.4	10.5
24	43.7	39.1	10.7	35.4	32.4	8.4
26	42.1	38.0	9.7	33.9	31.5	7.1
28	40.8	37.1	9.0	32.7	30.4	7.0
30	39.7	35.5	10.6	31.7	29.3	7.6

	A3			A4			A5		
Velocity	Fixed	DP	%	Fixed	DP	%	Fixed	DP	%
20	44.7	38.2	14.5	44.0	37.5	14.6	43.7	37.5	14.3
22	42.0	37.3	11.1	41.1	36.5	11.1	41.4	36.6	11.6
24	39.7	36.4	8.1	38.5	35.6	7.7	39.4	35.6	9.7
26	37.9	35.4	6.7	36.7	34.6	5.9	37.5	34.7	7.5
28	36.8	34.2	7.0	35.8	33.6	6.1	36.5	33.7	7.8
30	35.9	33.1	7.7	34.9	32.6	6.4	35.7	32.4	9.4

Table 5.7: Results from fixed velocity and constrained DP policies on artificial road profiles A1 to A5 with initial velocities from 20 m s^{-1} to 30 m s^{-1}

driveability constraints do not have a large effect on the velocity and gear selection profiles, this can be attributed to the fact that the goals of efficient driving and driveability overlap in that both aim for smooth changes in speed.

Real Road Profiles

The cost reduction in the four real road scenarios is shown in Table 5.8 with an average cost reduction of 10.4% compared to that of the unconstrained DP algorithm of 13.1% as noted in section 5.4.1. These results are comparable to those produced on the artificial roads with a difference between unconstrained and constrained DP of 2.7%.

	R1			R2		
Velocity	Fixed	DP	%	Fixed	DP	%
20	44.1	37.8	14.3	43.4	37.2	14.2
22	41.7	36.9	11.5	41.0	36.1	12.0
24	39.7	36.0	9.2	39.2	35.0	10.7
26	38.2	34.8	8.8	37.3	34.1	8.6
28	37.2	33.6	9.8	36.0	33.4	7.2
30	36.3	32.5	10.4	35.2	32.1	8.9
	R3			R4		
Velocity	Fixed	DP	%	Fixed	DP	%
20	43.4	36.6	15.6	42.7	36.2	15.2
22	40.7	35.7	12.2	40.2	35.3	12.1
24	38.6	34.7	10.1	38.3	34.4	10.4
26	36.8	33.6	8.8	36.5	33.4	8.6
28	35.3	32.7	7.1	35.1	32.4	7.6
30	34.2	31.7	7.3	34.1	31.2	8.4

Table 5.8: Results from fixed velocity and constrained DP policies on real road profiles R1 to R4 with initial velocities from 20 m s^{-1} to 30 m s^{-1}

Road	Fuel (l/100km)			Time (s)		
	Driver	DP	%	Driver	DP	%
R1	7.79	6.87	11.8	40.05	46.79	-16.8
R2	4.19	6.85	-63.6	68.70	53.61	22.0
R3	5.05	4.11	18.6	84.70	74.09	12.5
R4	4.16	3.98	4.5	69.15	65.91	4.7

Table 5.9: Real driver compared to DP results for roads R1-R4 with $\lambda = 0.5$, $\mu_t = 1$ and $\mu_f = 4$ and a residual cost considered for 1 km.

5.4.6 Real Driver Comparison

The fixed velocity strategy used previously allows the comparison of the DP algorithm to consistent reference values produced by an idealised cruise control system however this comparison does not consider the majority of driving undertaken without cruise control. To assess the performance of the DP algorithm under normal driving circumstances, real driving data was used as a comparison. For the four journeys described in Figure 5.4 and Figure 5.5, the DP algorithm is applied to a single horizon length.

The difference in fuel consumption and time between the recorded driver behaviour and the DP algorithm shows a wide range of results as seen in Table 5.9. A fuel/time weighting of $\lambda = 0.5$ was used for the algorithm to balance fuel and time equally in the cost function. For road R1 the section time is 16.8% higher for the DP algorithm than the driver however the velocity profile of the DP algorithm is restricted by the legal speed limit for the section of road, which in this instance was not always adhered to by the driver. Conversely in road R2 the section time is 22% longer while fuel consumption is 63.6% lower which indicates that the driver, while driving very efficiently in relation to fuel, this came at a great cost to the journey time. It is noted that driving considerably lower than the legal speed limit, such as highlighted in R2, may be due to traffic or road conditions rather than driver preference and where available information about such conditions should be used in the DP algorithm to ensure an unbiased comparison. The section of road considered from R2 is motorway approaching a junction



Figure 5.13: Road R3 speed limit reduction and upcoming curve responsible for speed reduction.

with no topographical features to explain a speed reduction. The impact of traffic information is investigated further in section 5.6 where the system is deployed in a real vehicle.

Roads R3 and R4 both produce fuel and time savings for the DP algorithm, however based solely on the legal speed limit there is no justification for the longer journey time of the driver. Further investigation of the sections of road in question identifies that in road R3 the speed limit reduces from 50 mph to 40 mph as the single lane road enters a rural village, however prior to this the upcoming curvature of the road may lead to a reduction in the speed of other vehicles propagating back to the vehicle being recorded. This road section is shown in Figure 5.13 as extracted from the Google streetview database [162] to highlight the features mentioned.

A reduction in the speed of the driver on the section of road R4 studied cannot be attributed to the road geometry so it is concluded that it is due to either traffic build up or suboptimal operation by the driver.

5.5 Genetic Algorithm Comparison

In the previous sections the DP algorithm has been shown to improve fuel consumption by producing an optimal velocity and gear profile, however

λ	0.5
μ_t	1.0
μ_f	4.0
Horizon interval (m)	50
Horizon steps	30
Velocity interval (m s^{-1})	1.0

Table 5.10: Optimisation algorithm settings

as noted in chapter 2, DP is not the only optimisation method that can be applied to this problem. To put the DP results in context with other optimisation methods identical scenarios are considered with both DP and Genetic Algorithm (GA) methods applied. For the four real road profiles presented in Figure 5.4 and Figure 5.5, the DP algorithm was applied with the settings detailed in Table 5.10. The GA was identically configured to ensure a fair comparison and the minimum costs produced by the GA for each road profile are shown in Figure 5.14 along with the DP results. The GA was run for 50,000 generations with results recorded every 100 generations resulting in an inverse relationship between the time taken to compute the given number of generations and the overall cost of the optimal profile produced. When deployed on an Intel i5-4200 CPU the GA computation times for 100 generations is 4 s and to take Road R3 as an example the cost after 50,000 generations is less than 2% lower than at 100 generations despite requiring more than 32 min to calculate. In comparison to the DP algorithm the result is more than 13% higher. It is therefore considered that although the GA method is theoretically capable of producing profiles of an equal cost to DP, the computation time of such a method rules it wholly unsuitable for real time in-vehicle optimisation due to the complexity of the problem. This work highlights the benefits of DP over GA in terms of performance as well as repeatability.

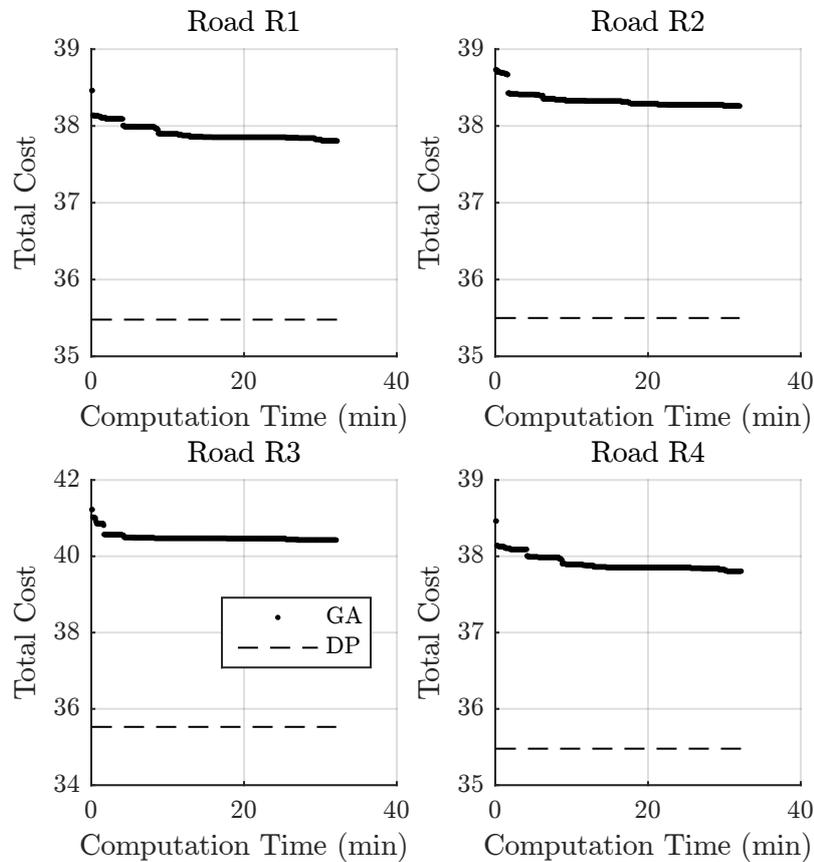


Figure 5.14: Genetic algorithm and DP results, with increasing numbers of GA generations producing lower cost profiles at the expense of longer computation times. DP results profile cost and time are shown for reference.

5.6 In-vehicle implementation

To test the practical aspects of calculating an optimal velocity and gear profile based on current vehicle state in real time in-situ in a vehicle, the optimisation system as described in chapter 4 was deployed in a midsize SUV. The GUI was presented on the 7" screen as described in chapter 4 and mounted centrally on the dashboard of the test vehicle as shown in Figure 5.15.

A real journey was to be undertaken on extra urban roads, the route of which is shown in Figure 5.16. Data was recorded from the vehicle dur-



Figure 5.15: Driver feedback system mounted on test vehicle dashboard with recommended speed and gear along with feedback icon to maintain speed.

ing the journey on extra urban roads with driver feedback provided with a weighting of $\lambda = 0.5$. The DP algorithm was continually updating however the implementation of the optimal velocity and gear profile relied on the driver being able to identify and apply the guidance provided given the traffic and road conditions. The test vehicle was equipped with an automatic transmission that allowed driver control of shifting within the limits of the baseline transmission control policy.

The data recorded during the test drive was retrospectively compared to simulated data generated by the vehicle model following the DP profile precisely, as seen in Figure 5.17. The fuel consumption and road section time for both the test drive and the DP simulations are shown in Table 5.11. The DP algorithm results were simulated with fuel and time weighting of $\lambda = 0.3$, 0.5 and 0.7.

The velocity profile produced by the driver is shown in Figure 5.17 along with the legal speed limit and real time traffic information while the three DP velocity profiles are also shown. The fuel consumption at $\lambda = 0.3$ is 6% lower than the test drive while also achieving a 13% reduction in journey time due to the higher average speed of the DP algorithm results. However the DP results for $\lambda = 0.5$ and $\lambda = 0.7$ show higher fuel consumption than the test drive while further reducing the journey time. The test drive velocity profile is always below the legal speed limit and includes a sharp reduction in speed prior to 19 km which does not correspond with the speed limit data, but follows more closely the traffic velocity profile. In the first



Figure 5.16: Test route on extra urban roads including multi lane freeway and single lane country roads.

set of results the DP algorithm only considers the legal speed limit and not the traffic speed, leading to profiles that may not be achievable in the given traffic conditions, hence the deviation in results.

Using the real-time traffic information recorded during the journey the effect of traffic on the DP algorithm can be investigated retrospectively. As seen in Figure 5.18 the driver velocity follows the traffic velocity rather than the speed limit however there is still deviation from the real-time traffic information around 16 km and 17.5 km. As such the DP algorithm is unfairly disadvantaged by the implementation of the traffic speed as an upper limit, for this reason a +10% upper bound is allowed above the traffic speed. The results are shown in lower section of Table 5.11 with improvements in fuel consumption for $\lambda = 0.3$ and $\lambda = 0.5$ coupled with smaller deteriorations in journey time. For all λ values the results are worse for time, which can be attributed to the upper speed limitation restricting the DP algorithm to lower speeds than the test drive velocity profile. The faster and more efficient policy is to exceed the traffic speed limit by more than the allocated 10%. Increasing the reliability and timeliness of the traffic information would

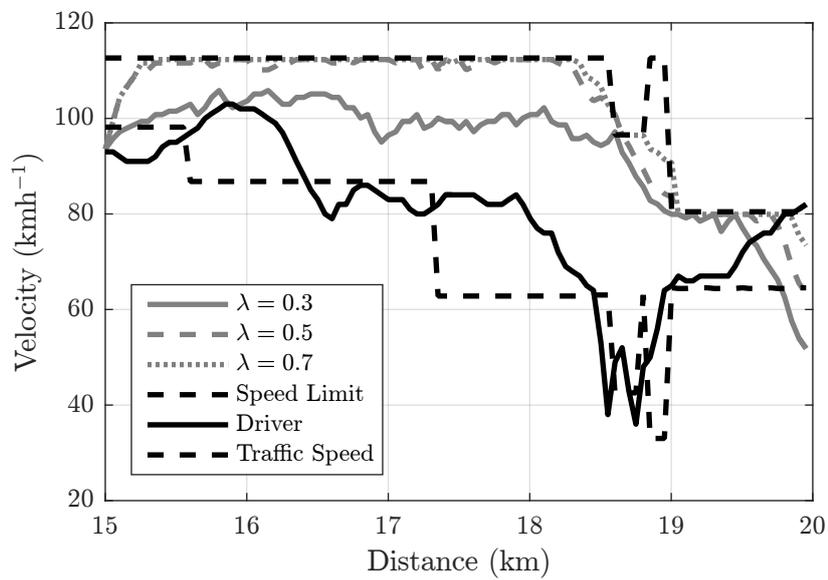


Figure 5.17: Velocity profiles from driver data (solid black) and DP optimal velocity (solid grey $\lambda = 0.3$, dashed grey, $\lambda = 0.5$ and dotted grey, $\lambda = 0.7$). Legal speed limit and real time traffic shown as dashed black line, upper and lower, respectively.

		Fuel Use (l/100km)	Difference (%)	Time (min)	Difference (%)
Test Drive (Recorded)		4.8		3.9	
Test Drive (Sim)		4.9	0	3.9	0
Without Traffic	DP, $\lambda = 0.3$	4.6	-6	3.4	-13
	DP, $\lambda = 0.5$	5.5	12	3.0	-22
	DP, $\lambda = 0.7$	5.9	20	3.0	-24
With Traffic	DP, $\lambda = 0.3$	4.4	-10	4.2	7
	DP, $\lambda = 0.5$	4.5	-8	4.1	6
	DP, $\lambda = 0.7$	4.9	0	4.1	5
Driver Window	DP, $\lambda = 0.3$	4.5	-8	3.9	0
	DP, $\lambda = 0.5$	5.0	2	3.7	-4
	DP, $\lambda = 0.7$	5.3	8	3.7	-4

Table 5.11: Fuel consumption and time from a 5 km section of the test drive. Test drive velocity and gear profile used to simulate fuel consumption and compared to results from DP algorithm without traffic information.

minimise this issue. The incorporation of vehicle sensors and/or vehicle to vehicle (V2V) communication would also improve the quality of traffic information. The gear selection of the optimization algorithm and the driver are shown in Figure 5.19 highlighting the foresight of the algorithm, as the traffic speed reduces at 18.5 km the algorithm initiates coasting in neutral gear. It is also observed prior to 17 km that the driver uses gear 8 while the algorithm remains in gear 9, this behaviour was observed occasionally during the test run due to the automatic transmission overriding the driver selection of gear 9. This can be explained by the restrictions of the TCU in the vehicle preventing selection of gear 9 below certain engine and vehicle speeds in order to maintain available power and prevent frequent hunting behaviour. The maintaining of available power is normally required to ensure that any upcoming power demand can be met, however with the DP horizon ensuring that the future power demand required to follow the optimal velocity profile is known, the decision to utilise gear 9 is justified.

By implementing a DP algorithm in a test vehicle to provide real time

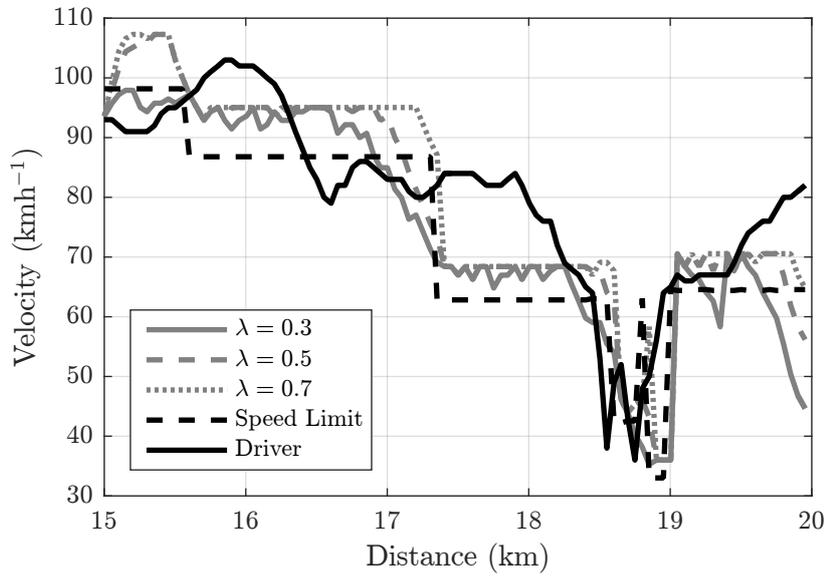


Figure 5.18: Velocity profiles from DP and driven data with traffic considered. Driver velocity (solid black), traffic speed (dashed black), algorithm profiles with $\lambda = 0.7$ (dotted), $\lambda = 0.5$ (dashed) and $\lambda = 0.3$ (solid grey).

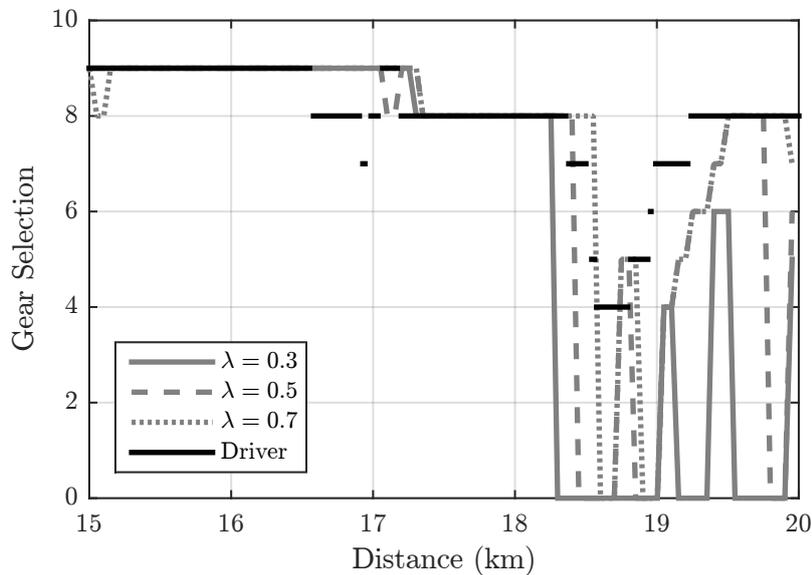


Figure 5.19: Gear profiles of optimization algorithm and driver selection where zero corresponds to neutral and the discontinuity in the driver gear selection is due to gearshift operations where no gear is selected.

guidance for velocity and gear selection it has been shown that such a system can be used to reduce fuel consumption. The development of the system as an eco-driving guidance tool or as an eco-cruise control system is made possible with this proof of concept.

5.7 Summary

To test the optimisation algorithm firstly the vehicle model was validated to ensure that any simulated results can be reasonably compared to real journey data. A range of both artificial and real road profiles were then presented for the algorithm to be tested with as well as real driving data from recorded journeys. To compare the DP algorithm results with real driving data depends on the behaviour of the driver during the period of data logging, as any efficiency improvements depend on how efficient the driving was initially. To overcome this inconsistency a baseline fuel and time are calculated for a policy of maintaining a fixed velocity and initially this is compared to the DP algorithm results. A speed limit following policy is then used as a minimum journey time comparison with the DP algorithm results. The driver's acceptance of the DP algorithm optimal velocity and gear profile depends on the driveability resulting from following the profile and so results were presented for which the DP algorithm was restricted in terms of torque produced and frequency of gearshifting in order to maintain driveability.

Chapter 6

Conclusion

A review of the hypotheses proposed in chapter 1 of this work is provided here with a summary of the findings in relation to each hypothesis.

- (a) Fuel savings can be made by utilising optimal control methods to control vehicle speed and gear selection in real-time, based on instantaneous vehicle and road data.
- (b) Driving experience, in relation to a vehicle's longitudinal performance, can be quantified and applied to an optimisation algorithm.
- (c) Fuel savings can be made as above, without compromising the driving experience and this can be verified across a range of real driving data.

6.1 Hypothesis (a)

The first part of hypothesis (a) that fuel savings can be made is investigated in simulation in both the literature reviewed in section 2.4 as well as in the results presented in chapter 5. A Dynamic Programming (DP) algorithm was developed and tested in simulation firstly against a fixed velocity policy replicating a standard cruise control system on both artificial and real road profiles with an average reduction in cost function of 15% for the DP results. The DP algorithm was then compared with a speed limit following policy on real road profiles with an average cost function reduction of 9%. In order

to verify that this method could be applied in real-time a computationally efficient DP algorithm was developed that was able to produce an optimal velocity and gear selection profile in a sufficiently short time to be utilised in a real-time in-vehicle deployment. The parameters of the DP algorithm, primarily the horizon length and velocity discretisation interval were investigated to identify the optimal balance of computation time with quality of results. To ensure that the algorithm was calculated based on current data, a road database was developed utilising a GPS module to provide road data specific to the current position. Similarly current vehicle data was incorporated with a bespoke Python CAN bus extraction module. A Raspberry Pi was chosen as the target hardware for deployment due to its flexibility and as it represents a reasonable approximation of the computational power available in a production vehicle control unit that could ultimately run such a eco-guidance system. While it is shown that there are situations where the algorithm will violate the firm real-time constraints applied, the results even in those exceptional circumstances will still be of relevance. It is considered that the ever increasing computational power available in a production vehicle will render this issue null and void.

6.2 Hypothesis (b)

Driveability is found in literature reviewed in section 2.6 to be a complex issue that involves numerous distinct elements that influence the driving experience. Despite this it has been shown that only a subset of these elements are directly linked to longitudinal performance and influenced by the control of vehicle velocity and gear selection and so only these were considered for further investigation in this project.

In a vehicle with automatic transmission, the performance of the gear shifting is frequently highlighted as a major contributing factor to the assessment of vehicle driveability [110,111,113]. The frequency of gear shifting is one aspect of this performance and can easily be recorded using data from the vehicle communication network. The response time of the gear selection control to changes in power demand is also a driveability issue which although harder to quantify should benefit from the use of a road data hori-

zon allowing future road conditions to be factored in to the gear selection to pre-empt power demand.

With regard to the vehicle velocity selection, longitudinal acceleration and deceleration play a crucial role in the driveability and so limits were set to ensure that any proposed velocity profile does not compromise driveability. These two main driveability constraints were implemented in the DP algorithm by penalising abrupt torque changes with a variable in the cost function and a hard limit on the frequency of gear shifting by storing the gear selection in each step of the gear profiles being compared. It was noted that the DP algorithm complied with these constraints to produce velocity and gear profiles that objectively would satisfy longitudinal driveability conditions.

6.3 Hypothesis (c)

By combining the work produced in testing the first two hypotheses it was possible to test hypothesis (c). The DP algorithm developed in this work was modified to incorporate constraints on the gear selection policy, velocity profile and its derivative to ensure driveability in the context of longitudinal control was maintained. By simulating the algorithm on firstly artificial road profiles and then real road profiles recorded from typical journeys it was shown that fuel savings could be made with varying levels of impact on the journey time according to the algorithm parameters. The fuel savings were achieved on a variety of road types with speed limits from the maximum of 70 mph down to 30 mph. It was shown that the introduction of driveability constraints, while having an impact on the optimal profiles produced, do not alter drastically the velocity and gear selection profiles generated as achieving fuel efficiency and driveability are often complementary goals. This result gives confidence that eco cruise control systems would be well received by drivers with driveability maintained.

Following the simulation work the algorithm was deployed in a test vehicle using Raspberry Pi hardware to test the real time performance of the vehicle in a production vehicle. The optimisation algorithm performed consistently in the vehicle however the limitations of the driver following the

guidance via visual and audible feedback as well as the influence of traffic conditions resulted in less than optimal performance. The limitations of the guidance system would be removed if an adaptive cruise control system was tasked with implementing the optimal velocity profile as opposed to the driver. Improved traffic information would also enable such a system to achieve its potential.

In the course of this work it was necessary to test the DP algorithm on a number of roads varying both in topography and artificial properties such as legal speed limit. A methodology is described in chapter 4 for processing data recorded in a vehicle using either Raspberry Pi hardware or commercial datalogging equipment to reconstruct road profiles driven in real journeys and by utilising online map data providers these can be supplemented with road data such as legal speed limit. By using such real road profiles and vehicle data recorded during the same journey the scenario can be replicated in simulation and used for assessing the realistic impact of the DP algorithm prior to testing in a real vehicle.

In summary, the fuel consumption of current and future vehicles can be improved by applying eco-driving techniques and the additional mental load of correctly implementing such techniques can be minimised by the use of a real-time in-vehicle optimisation algorithm. A system is presented in this work that incorporates such an algorithm and provides guidance with a view to ultimately integrate with an adaptive cruise control system to ensure that the optimal velocity and gear selection profile are followed.

Chapter 7

Further Work

The results of the work presented here provide a number of opportunities for further research that build on the projects contributions. A Dynamic Programming algorithm was developed to optimise velocity and gear selection to minimise fuel consumption and journey time using upcoming road data and a vehicle model that balanced computational load and accuracy. Within this framework it was possible to minimise the computational load of the algorithm so as to allow real-time in-vehicle deployment as well as test scenarios in simulation programs developed as part of the project. Identified during the course of the literature review and the algorithm development and testing are a number of areas of further work:

- Integration with commercial eHorizon systems based on the Advanced Driver Assistance System Interface Specification (ADASIS).
- Improvement of traffic and road environment data by incorporating vehicle sensors and vehicle communication systems.
- Investigation of environmental, economic and customer experience impact of eco-guidance assisted and eco-cruise systems compared to unaided driving.
- Development of hybrid and fully electric vehicle models in place of conventional vehicle model.
- Emissions such as NO_x included in cost function along with fuel con-

sumption by incorporating a control oriented model of emissions generation .

7.1 Commercial eHorizon System Integration

The work presented utilised a minimal set of Electronic Horizon data based on that which would be available from a Horizon Reconstructor using ADASIS. This approach was taken due to the cost and additional complexity involved in integrating a commercial system which would reduce the time available for the core activities of the project. Developing the existing system in conjunction with a commercial ADASIS eHorizon provider would open up a complete network of real road scenarios for testing and improving versions of the DP algorithm. The potential for testing would only be limited by the size of the map database available. Additionally, successfully combining the DP algorithm with a supplier's standard eHorizon system would increase the deployment opportunities for the DP algorithm in test vehicles, as the ADASIS standard provides the framework for such interoperability.

7.2 Vehicle Communication Systems Integration

The work presented here resulted in a system that guides vehicle velocity and gear selection based on road information such as legal speed limit and road gradient [57, 74]. In testing however, the ability to follow this optimal velocity profile is often compromised by dynamic factors such as congestion [72], traffic light timings [163] and variable speed limits which cannot be contained in a static database of road information. Increasingly connected vehicles are beginning to exploit the possibility to communicate a wide range of data both to and from vehicles primarily for safety and entertainment purposes [164]. In order to best utilise such communication channels for energy use optimization it is necessary to understand which data has the most impact on energy use.

As seen in section 5.6 there is a disturbance in the vehicle velocity due to traffic that is insufficiently reported through the traffic information service used. The cause of this disturbance is congested traffic conditions and

highlights the need for additional road information to be available to velocity optimization systems, information that can be provided by a network of connected vehicles. Publications investigating this issue have appeared in recent years [165]. Vehicle Ad-hoc Networks (VANETs) provide the facility that allows both vehicle to vehicle (V2V) and vehicle to infrastructure (V2I) communication [166]. While the mechanisms by which messages can be exchanged are being investigated extensively, an investigation of what messages would be of most benefit to vehicle energy optimisation could be simulated as an extension of the work presented here.

7.3 Eco-guidance and eco-cruise control impact

While the simulation work presented in this project utilised real driving data and a test drive with the eco-guidance system was undertaken, due to the limited availability of test vehicles, there is scope for further investigations into the impact of both eco-guidance and with the suitable risk assessment and safety systems in place, eco-cruise control. A range of drivers approximating the variation in driving styles in the general population would be necessary to test the impact of the eco-driving systems developed in this project. A benchmark could be produced for a set of journeys that are driven by a selection of drivers first without any guidance on eco-driving either in vehicle or otherwise, then each test repeated with eco-guidance and finally with eco-cruise. Assuming the variation of the test scenarios suitably approximates the driving of the general population a reasonable estimate of the impact of such eco-driving systems could be found. There is much scope for improving the user interface that provides the eco-guidance, for instance haptic feedback can be investigated to guide use of the accelerator pedal [167] as well as improvements to the audio and visual eco-guidance systems.

7.4 Alternative Powertrain Vehicles

As the project was focussed solely on a vehicle powered by a 4-cylinder diesel engine, one of the possible areas of further work is the development of

Hybrid Electric Vehicle (HEV) and fully electric vehicle models that can be integrated into the Dynamic Programming framework. These models can build on existing literature and the sensitivity analyses approaches detailed in this work to identify the most influential elements of such vehicle models. This will allow the development of a model that provides sufficient accuracy while minimising computation time for incorporation in the real time dynamic programming framework. Fully electric, plug-in and conventional HEVs can be investigated, with different aims depending on whether the destination is known to the system and if charging is available at the destination. In addition to modelling the energy demands for propulsion the vehicle models could also be extended to include auxiliary loads such as air conditioning to improve the quality of results in real world driving.

7.5 Emissions Model inclusion in Vehicle Control Optimisation

Minimising fuel consumption has been the main focus of research in hybrid vehicle optimization in the last decade, however with the increase in diesel hybrid electric vehicles and recent studies that have highlighted the importance of reducing NO_x emissions, [168] the requirements of vehicle design are more complex. Regulations on NO_x emissions are in place [169] and set to be enforced in real driving situations in addition to in laboratory tests. In order to minimise NO_x emissions using an optimization algorithm, a control oriented model representing NO_x production is required. The production of NO_x is a complex process that depends on temperature, fuel mixing and pressure [2]. Historically, modelling this process has involved crank angle based models which generate a high computational load such that they cannot be applied in real-time in a vehicle. The use of neural networks to model NO_x emissions [170] improve on this but still present implementation challenges for production vehicles. In a similar manner to that of eco-driving guidance where the complex systems involved in fuel consumption are simplified into straightforward driving instructions, if the use of NO_x models can identify operating behaviour that reduces such emissions and this can be in-

corporated in the optimisation process the benefits, particularly in polluted urban areas are clear.

The further work described in this chapter builds on the foundation of work produced in this project and extends its impact into alternative powertrain vehicles and emissions reduction as well as guiding the utilisation and development of new technologies such as ADAS and V2V.

References

- [1] T. Stocker, D. Qin, G. Plattner, M. Tignor, S. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex, and P. Midgley, “Summary for Policymakers. Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change,” *IPCC Localised*, vol. 9, p. 2013, 2013.
- [2] J. B. Heywood, *Internal combustion engine fundamentals*. Mcgraw-hill New York, 1988, vol. 930.
- [3] F. An and A. Sauer, “Comparison of passenger vehicle fuel economy and greenhouse gas emission standards around the world,” *Pew Center on Global Climate Change*, vol. 25, 2004.
- [4] *REGULATION (EC) No 715/2007 - on type approval of motor vehicles with respect to emissions from light passenger and commercial vehicles (Euro 5 and Euro 6) and on access to vehicle repair and maintenance information*, THE EUROPEAN PARLIAMENT Std., 2007.
- [5] European Environment Agency, “Air quality in europe 2016 report,” Tech. Rep., 2016.
- [6] T. Klier and J. Linn, “The price of gasoline and new vehicle fuel economy: evidence from monthly sales data,” *American Economic Journal: Economic Policy*, vol. 2, no. 3, pp. 134–153, 2010.
- [7] T. Levermore, M. N. Sahinkaya, Y. Zweiri, and B. Neaves, “Real-time velocity optimization to minimize energy use in passenger

- vehicles,” *Energies*, vol. 10, no. 1, p. 30, 2017. [Online]. Available: <http://www.mdpi.com/1996-1073/10/1/30>
- [8] “Passenger car world vehicles in use,” OICA, Tech. Rep., 2015. [Online]. Available: <http://www.oica.net/wp-content/uploads/pc-inuse-2015.xlsx>
- [9] “Economic and Market Report Q1 2016,” European Automobile Manufacturer’s Association, Tech. Rep., December 2016.
- [10] K. Mollenhauer and H. Tschöke, *Handbook of Diesel Engines*. Springer-Verlag Berlin Heidelberg, 2010.
- [11] DEFRA, “Emissions of air pollutants in the uk, 1970 to 2015,” Department for Environment, Food and Rural Affairs, Tech. Rep., 2016.
- [12] “CO₂ emissions from fuel combustion highlights 2015,” International Energy Agency, Tech. Rep., 2015. [Online]. Available: <http://www.iea.org/publications/freepublications/publication/CO2EmissionsFromFuelCombustionHighlights2015.pdf>
- [13] *Tier 3 Motor Vehicle Emission and Fuel Standards, Vol. 80, No. 33*, Environmental Protection Agency Std., February 2015.
- [14] *Regulation (EC) No 443/2009 of the European Parliament and of the Council of 23 April 2009*, European Commission Std., 2009.
- [15] *Uniform provisions concerning the approval of passenger cars powered by an internal combustion engine only, or powered by a hybrid electric power train with regard to the measurement of the emission of carbon dioxide and fuel consumption and/or the measurement of electric energy consumption and electric range, and of categories M1 and N1 vehicles powered by an electric power train only with regard to the measurement of electric energy consumption and electric range*, United Nations Economic Commission for Europe (UNECE) Std. E/ECE/324-E/ECE/TRANS/505/Rev.2/Add.100/Rev.2, 2011.
- [16] G. Pinto and M. T. Oliver-Hoyo, “Using the relationship between vehicle fuel consumption and co2 emissions to illustrate chemical

- principles,” *Journal of Chemical Education*, vol. 85, no. 2, p. 218, 2008. [Online]. Available: <http://dx.doi.org/10.1021/ed085p218>
- [17] *2017 and Later Model Year Light Duty Vehicle Greenhouse Gas Emissions and Corporate Average Fuel Economy Standards, Federal Register, Vol. 77, No. 199*, Environmental Protection Agency Std. 199, October 2012.
- [18] (2015) Global passenger vehicle standards comparison tool. International Council on Clean Transportation (ICCT). [Online]. Available: <http://www.theicct.org/info-tools/global-passenger-vehicle-standards>
- [19] S. Pagerit, P. Sharer, and A. Rousseau, “Fuel economy sensitivity to vehicle mass for advanced vehicle powertrains,” in *SAE Technical Paper*. SAE International, 04 2006. [Online]. Available: <http://dx.doi.org/10.4271/2006-01-0665>
- [20] D. Greene, J. Hooker, A. Rose, and G. Samuels, “Energy conservation in transportation,” *Advances in Energy Systems and Technology*, vol. 3, p. 187, 2013.
- [21] J. Bishop, A. Nedungadi, G. Ostrowski, B. Surampudi, P. Armiroli, and E. Taspinar, “An engine start/stop system for improved fuel economy,” SAE Technical Paper, Tech. Rep., 2007.
- [22] A. R. Carrico, P. Padgett, M. P. Vandenberg, J. Gilligan, and K. A. Wallston, “Costly myths: An analysis of idling beliefs and behavior in personal motor vehicles,” *Energy Policy*, vol. 37, no. 8, pp. 2881 – 2888, 2009. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0301421509001633>
- [23] C. C. Chan, “The state of the art of electric, hybrid, and fuel cell vehicles,” *Proceedings of the IEEE*, vol. 95, no. 4, pp. 704–718, April 2007.
- [24] J. N. Barkenbus, “Eco-driving: An overlooked climate change initiative,” *Energy Policy*, vol. 38, no. 2, pp. 762 – 769, 2010.

- [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0301421509007484>
- [25] D. Stacy, D. Susan, and R. Boundy, “Transportation energy data book: Edition 34,” Oak Ridge National Laboratory, Tech. Rep., 2015. [Online]. Available: <http://cta.ornl.gov/data/download34.shtml>
- [26] M. Sivak and B. Schoettle, “Eco-driving: Strategic, tactical, and operational decisions of the driver that influence vehicle fuel economy,” *Transport Policy*, vol. 22, July 2012. [Online]. Available: <http://dx.doi.org/10.1016/j.tranpol.2012.05.010>
- [27] E. Ericsson, H. Larsson, and K. Brundell-Freij, “Optimizing route choice for lowest fuel consumption potential effects of a new driver support tool,” *Transportation Research Part C: Emerging Technologies*, vol. 14, no. 6, pp. 369 – 383, 2006. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0968090X06000799>
- [28] “GNSS Market Report Issue 4,” European Global Navigation Satellite Systems Agency (GSA), Tech. Rep. 4, 2015.
- [29] W. Zeng, T. Miwa, and T. Morikawa, “Prediction of vehicle CO2 emission and its application to eco-routing navigation,” *Transportation Research Part C: Emerging Technologies*, vol. 68, pp. 194–214, 2016.
- [30] S. Park and H. Rakha, “Energy and environmental impacts of roadway grades,” *Transportation Research Record: Journal of the Transportation Research Board*, no. 1987, pp. 148–160, 2006.
- [31] M. André and U. Hammarström, “Driving speeds in europe for pollutant emissions estimation,” *Transportation Research Part D: Transport and Environment*, vol. 5, no. 5, pp. 321 – 335, 2000. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S136192090000002X>
- [32] T. Hiraoka, Y. Terakado, S. Matsumoto, and S. Yamabe, “Quantitative evaluation of eco-driving on fuel consumption based

- on driving simulator experiments,” in *Proceedings of the 16th World Congress on Intelligent Transport Systems*, 2009, pp. 21–25. [Online]. Available: [www.symbalab.sys.i.kyoto-u.ac.jp%2F~hiraoka%2Fpdf%2FITSWC2009.pdf](http://www.symbalab.sys.i.kyoto-u.ac.jp/~hiraoka%2Fpdf%2FITSWC2009.pdf)
- [33] B. Beusen, S. Broekx, T. Denys, C. Beckx, B. Degraeuwe, M. Gijssbers, K. Scheepers, L. Govaerts, R. Torfs, and L. I. Panis, “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, no. 7, pp. 514 – 520, 2009. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1361920909000698>
- [34] F. Mensing, E. Bideaux, R. Trigui, J. Ribet, and B. Jeanneret, “Eco-driving: An economic or ecologic driving style?” *Transportation Research Part C: Emerging Technologies*, vol. 38, pp. 110 – 121, 2014. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0968090X13002374>
- [35] P. Seewald, J. Josten, A. Zlocki, and L. Eckstein, “User acceptance evaluation approach of energy efficient driver assistance systems.”
- [36] M. P. Manser, M. Rakauskas, J. Graving, and J. W. Jenness, “Fuel economy driver interfaces: Develop interface recommendations,” National Highway Traffic Safety Administration, Tech. Rep., 2010.
- [37] 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, no. 2, pp. 141–146, 2009.
- [38] J. W. Jenness, J. Singer, J. Walrath, and E. Lubar, “Fuel economy driver interfaces: Design range and driver opinions (report on task 1 and task 2),” Tech. Rep., 2009.
- [39] I. M. Berry, “The effects of driving style and vehicle performance on the real-world fuel consumption of u.s. light-duty vehicles,” Master’s thesis, Massachusetts Institute of Technology, 2010. [Online]. Available: <http://hdl.handle.net/1721.1/58392>

- [40] (2013) Ecowill brochure (d9.5). ECOWILL Project. [Online]. Available: http://www.ecodrive.org/download/downloads/ecowill_brochure.pdf
- [41] *Regulation (EC) No 661/2009 of the European Parliament and of the Council of 13 July 2009 concerning type-approval requirements for the general safety of motor vehicles, their trailers and systems, components and separate technical units intended therefor*, European Commission Std., 2009.
- [42] 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*, vol. 226, no. 12, pp. 1630–1639, 2012. [Online]. Available: <http://pid.sagepub.com/content/226/12/1630.abstract>
- [43] C. Raubitschek, N. Schütze, E. Kozlov, and B. Bäker, “Predictive driving strategies under urban conditions for reducing fuel consumption based on vehicle environment information,” in *Integrated and Sustainable Transportation System (FISTS), 2011 IEEE Forum on*, June 2011, pp. 13–19.
- [44] P. Shakouri, A. Ordys, P. Darnell, and P. Kavanagh, “Fuel efficiency by coasting in the vehicle,” *International Journal of Vehicular Technology*, vol. 2013, p. 14, 2013. [Online]. Available: <http://dx.doi.org/10.1155/2013/391650>
- [45] R. C. Dorf and R. H. Bishop, *Modern Control Systems*, 9th ed. Upper Saddle River, NJ, USA: Prentice-Hall, Inc., 2000.
- [46] L. Guzzella and A. Sciarretta, *Vehicle Propulsion Systems: Introduction to Modeling and Optimization*. Springer, 2005. [Online]. Available: <http://books.google.co.uk/books?id=Qpq7jhbNtk8C>
- [47] R. Fellini, N. Michelena, P. Papalambros, and M. Sasena, “Optimal design of automotive hybrid powertrain systems,” in *Environmentally Conscious Design and Inverse Manufacturing, 1999. Proceedings*.

- EcoDesign '99: First International Symposium On*, Feb 1999, pp. 400–405.
- [48] A. Schwarzkopf and R. Leipnik, “Control of highway vehicles for minimum fuel consumption over varying terrain,” *Transportation Research*, vol. 11, no. 4, pp. 279 – 286, 1977. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/0041164777900934>
- [49] L. S. Pontryagin, *Mathematical theory of optimal processes*. CRC Press, 1987.
- [50] A. Stoicescu, “On fueloptimal velocity control of a motor vehicle,” *International Journal of Vehicle Design*, vol. 16, no. 2/3, pp. p. 229–56–, 1995. [Online]. Available: <https://trid.trb.org/view/426247>
- [51] J. N. Hooker, A. B. Rose, and G. F. Roberts, “Optimal control of automobiles for fuel economy.” *Transportation Science*, vol. 17, no. 2, p. 146, 1983.
- [52] R. Bellman, *Dynamic Programming*. Princeton, New Jersey: Princeton Univ Press, 1957.
- [53] J. Hooker, “Optimal driving for single-vehicle fuel economy,” *Transportation Research Part A: General*, vol. 22, no. 3, pp. 183–201, 1988.
- [54] P. Kock, A. Ordys, G. Collier, and R. Weller, “Intelligent predictive cruise control application analysis for commercial vehicles based on a commercial vehicles usage study,” *SAE Int. J. Commer. Veh.*, vol. 6, pp. 598–603, 10 2013. [Online]. Available: <http://dx.doi.org/10.4271/2013-01-9022>
- [55] A. Fröberg, E. Hellström, and L. Nielsen, “Explicit fuel optimal speed profiles for heavy trucks on a set of topographic road profiles,” SAE Technical Paper, Tech. Rep., 2006.
- [56] M. Ivarsson, J. Åslund, and L. Nielsen, “Look-ahead control consequences of a non-linear fuel map on truck fuel consumption,” *Proceedings of the Institution of Mechanical Engineers, Part D:*

- Journal of Automobile Engineering*, vol. 223, no. 10, pp. 1223–1238, 2009. [Online]. Available: <http://pid.sagepub.com/content/223/10/1223.abstract>
- [57] E. Hellström, M. Ivarsson, J. Aslund, and L. Nielsen, “Look-ahead control for heavy trucks to minimize trip time and fuel consumption,” *Control Engineering Practice*, vol. 17, no. 2, pp. 245 – 254, 2009. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0967066108001251>
- [58] S. Terwen, M. Back, and V. Krebs, “Predictive powertrain control for heavy duty trucks,” in *Proceedings of IFAC symposium in advances in automotive control*, 2004, pp. 451–457.
- [59] A. Sciarretta and L. Guzzella, “Control of hybrid electric vehicles,” *IEEE Control Systems*, vol. 27, no. 2, pp. 60–70, April 2007.
- [60] G. Paganelli, T. M. Guerra, S. Delprat, J.-J. Santin, M. Delhom, and E. Combes, “Simulation and assessment of power control strategies for a parallel hybrid car,” *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering*, vol. 214, no. 7, pp. 705–717, 2000. [Online]. Available: <http://pid.sagepub.com/content/214/7/705.abstract>
- [61] C. Musardo, G. Rizzoni, Y. Guezennec, and B. Staccia, “A-ECMS: An adaptive algorithm for hybrid electric vehicle energy management,” *European Journal of Control*, vol. 11, no. 45, pp. 509 – 524, 2005. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0947358005710487>
- [62] V. H. Johnson, K. B. Wipke, and D. J. Rausen, “Hev control strategy for real-time optimization of fuel economy and emissions,” in *SAE Technical Paper*. SAE International, 04 2000. [Online]. Available: <http://dx.doi.org/10.4271/2000-01-1543>
- [63] J.-S. Won, R. Langari, and M. Ehsani, “An energy management and charge sustaining strategy for a parallel hybrid vehicle with cvt,” *IEEE*

- Transactions on Control Systems Technology*, vol. 13, no. 2, pp. 313–320, March 2005.
- [64] F. R. Salmasi, “Control strategies for hybrid electric vehicles: Evolution, classification, comparison, and future trends,” *Vehicular Technology, IEEE Transactions on*, vol. 56, no. 5, pp. 2393–2404, 2007.
- [65] C.-C. Lin, H. Peng, J. W. Grizzle, and J.-M. Kang, “Power management strategy for a parallel hybrid electric truck,” *IEEE Transactions on Control Systems Technology*, vol. 11, no. 6, pp. 839–849, Nov 2003.
- [66] C.-C. Lin, H. Peng, and J. W. Grizzle, “A stochastic control strategy for hybrid electric vehicles,” in *American Control Conference, 2004. Proceedings of the 2004*, vol. 5, June 2004, pp. 4710–4715 vol.5.
- [67] S. Ross, *Introduction to Stochastic Dynamic Programming*. Academic Press, 1983. [Online]. Available: <https://books.google.co.uk/books?hl=en&lr=&id=bBLjBQAAQBAJ&oi=fnd&pg=PP1&dq=stochastic+dynamic+programming+markov&ots=vT4YCFLfMv&sig=20eWLLj84QCQ0sfzCmEYRRZZyLE#v=onepage&q=markov&f=false>
- [68] I. V. Kolmanovsky and D. P. Filev, “Stochastic optimal control of systems with soft constraints and opportunities for automotive applications,” in *2009 IEEE Control Applications, (CCA) Intelligent Control, (ISIC)*, July 2009, pp. 1265–1270.
- [69] G. E. Katsargyri, I. V. Kolmanovsky, J. Michelini, M. L. Kuang, A. M. Phillips, M. Rinehart, and M. A. Dahleh, “Optimally controlling hybrid electric vehicles using path forecasting,” in *2009 American Control Conference*, June 2009, pp. 4613–4617.
- [70] J. M. Rowell, “Applying map databases to advanced navigation and driver assistance systems,” *The Journal of Navigation*, vol. 54, no. 3, pp. 355–363, 09 2001, copyright - Copyright Cambridge University Press Sep 2001; Last updated - 2015-08-15. [Online]. Available: <https://search.proquest.com/docview/229599973?accountid=14557>

- [71] F. Mensing, R. Trigui, and E. Bideaux, "Vehicle trajectory optimization for application in eco-driving," in *2011 IEEE Vehicle Power and Propulsion Conference*, Sept 2011, pp. 1–6.
- [72] F. Mensing, E. Bideaux, R. Trigui, and H. Tattetrain, "Trajectory optimization for eco-driving taking into account traffic constraints," *Transportation Research Part D: Transport and Environment*, vol. 18, pp. 55 – 61, 2013. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1361920912001137>
- [73] J. Wollaeger, S. A. Kumar, S. Onori, D. Filev, Ü. Özgüner, G. Rizzoni, and S. D. Cairano, "Cloud-computing based velocity profile generation for minimum fuel consumption: A dynamic programming based solution," in *2012 American Control Conference (ACC)*, June 2012, pp. 2108–2113.
- [74] H. G. Wahl and F. Gauterin, "An iterative dynamic programming approach for the global optimal control of hybrid electric vehicles under real-time constraints," in *Intelligent Vehicles Symposium (IV), 2013 IEEE*, June 2013, pp. 592–597.
- [75] H. G. Wahl, K. L. Bauer, F. Gauterin, and M. Holzpfel, "A real-time capable enhanced dynamic programming approach for predictive optimal cruise control in hybrid electric vehicles," in *16th International IEEE Conference on Intelligent Transportation Systems (ITSC 2013)*, Oct 2013, pp. 1662–1667.
- [76] B. Passenberg, P. Kock, and O. Stursberg, "Combined time and fuel optimal driving of trucks based on a hybrid model," in *Control Conference (ECC), 2009 European*, Aug 2009, pp. 4955–4960.
- [77] M. A. S. Kamal, M. Mukai, J. Murata, and T. Kawabe, "Model predictive control of vehicles on urban roads for improved fuel economy," *IEEE Transactions on Control Systems Technology*, vol. 21, no. 3, pp. 831–841, May 2013.
- [78] B. Asadi and A. Vahidi, "Predictive cruise control: Utilizing upcoming traffic signal information for improving fuel economy and reducing trip

- time,” *IEEE Transactions on Control Systems Technology*, vol. 19, no. 3, pp. 707–714, May 2011.
- [79] H. T. Luu, L. Nouvelire, and S. Mammar, “Dynamic programming for fuel consumption optimization on light vehicle,” *{IFAC} Proceedings Volumes*, vol. 43, no. 7, pp. 372 – 377, 2010, 6th {IFAC} Symposium on Advances in Automotive Control. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1474667015368579>
- [80] T. van Keulen, B. de Jager, D. Foster, and M. Steinbuch, “Velocity trajectory optimization in hybrid electric trucks,” in *Proceedings of the 2010 American Control Conference*, June 2010, pp. 5074–5079.
- [81] E. Ozatay, S. Onori, J. Wollaeger, U. Ozguner, G. Rizzoni, D. Filev, J. Michelini, and S. Di Cairano, “Cloud-based velocity profile optimization for everyday driving: A dynamic-programming-based solution,” *Intelligent Transportation Systems, IEEE Transactions on*, vol. 15, no. 6, pp. 2491–2505, Dec 2014.
- [82] G. Heppeler, M. Sonntag, and O. Sawodny, “Fuel efficiency analysis for simultaneous optimization of the velocity trajectory and the energy management in hybrid electric vehicles,” *{IFAC} Proceedings Volumes*, vol. 47, no. 3, pp. 6612 – 6617, 2014, 19th {IFAC} World Congress. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1474667016426510>
- [83] E. Ozatay, U. Ozguner, D. Filev, and J. Michelini, “Analytical and numerical solutions for energy minimization of road vehicles with the existence of multiple traffic lights,” in *52nd IEEE Conference on Decision and Control*, Dec 2013, pp. 7137–7142.
- [84] D.-W.-I. S. Gausemeier, D.-I. K.-P. Jker, and P. D.-I. habil. Ansgar Trachtler, “Multi-objective optimization of a vehicle velocity profile by means of dynamic programming,” *{IFAC} Proceedings Volumes*, vol. 43, no. 7, pp. 366 – 371, 2010, 6th {IFAC} Symposium on Advances in Automotive Control. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1474667015368567>

- [85] S. Mandava, K. Boriboonsomsin, and M. Barth, "Arterial velocity planning based on traffic signal information under light traffic conditions," in *2009 12th International IEEE Conference on Intelligent Transportation Systems*, Oct 2009, pp. 1–6.
- [86] F. Mensing, "Optimal energy utilization in conventional, electric and hybrid vehicles and its application to eco-driving," Ph.D. dissertation, INSA de Lyon, 2013. [Online]. Available: <http://www.theses.fr/2013ISAL0106.pdf>
- [87] W. Dib, L. Serrao, and A. Sciarretta, "Optimal control to minimize trip time and energy consumption in electric vehicles," in *2011 IEEE Vehicle Power and Propulsion Conference*, Sept 2011, pp. 1–8.
- [88] J. Holland, *Adaptation in Natural and Artificial Systems*. University of Michigan Press, 1975.
- [89] D. Whitley, "A genetic algorithm tutorial," *Statistics and Computing*, vol. 4, no. 2, pp. 65–85, 1994. [Online]. Available: <http://dx.doi.org/10.1007/BF00175354>
- [90] Z. Michalewicz, C. Z. Janikow, and J. B. Krawczyk, "A modified genetic algorithm for optimal control problems," *Computers & Mathematics with Applications*, vol. 23, no. 12, pp. 83 – 94, 1992. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/089812219290094X>
- [91] A. Piccolo, L. Ippolito, V. zo Galdi, and A. Vaccaro, "Optimisation of energy flow management in hybrid electric vehicles via genetic algorithms," in *Advanced Intelligent Mechatronics, 2001. Proceedings. 2001 IEEE/ASME International Conference on*, vol. 1, 2001, pp. 434–439 vol.1.
- [92] M. Montazeri-Gh, A. Poursamad, and B. Ghalichi, "Application of genetic algorithm for optimization of control strategy in parallel hybrid electric vehicles," *Journal of the Franklin Institute*, vol. 343, no. 45, pp. 420 – 435, 2006, modeling, Simulation and Applied

- Optimization. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0016003206000329>
- [93] L. Fang, S. Qin, G. Xu, T. Li, and K. Zhu, “Simultaneous optimization for hybrid electric vehicle parameters based on multi-objective genetic algorithms,” *Energies*, vol. 4, no. 3, p. 532, 2011. [Online]. Available: <http://www.mdpi.com/1996-1073/4/3/532>
- [94] A. Gaier and A. Asteroth, “Evolution of optimal control for energy-efficient transport,” in *2014 IEEE Intelligent Vehicles Symposium Proceedings*, June 2014, pp. 1121–1126.
- [95] P. Othaganont, F. Assadian, and D. J. Auger, “Multi-objective optimisation for battery electric vehicle powertrain topologies,” *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering*, 2016. [Online]. Available: <http://pid.sagepub.com/content/early/2016/10/05/0954407016671275.abstract>
- [96] Y. V. Bocharnikov, A. M. Tobias, and C. Roberts, “Reduction of train and net energy consumption using genetic algorithms for trajectory optimisation,” in *IET Conference on Railway Traction Systems (RTS 2010)*, April 2010, pp. 1–5.
- [97] C. Sicre, A. Cucala, and A. Fernandez-Cardador, “Real time regulation of efficient driving of high speed trains based on a genetic algorithm and a fuzzy model of manual driving,” *Engineering Applications of Artificial Intelligence*, vol. 29, pp. 79 – 92, 2014. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0952197613001437>
- [98] A. Sahraeian, M. Shahbakhti, A. R. Aslani, S. Jazayeri, S. Azadi, and A. H. Shamekhi, “Longitudinal vehicle dynamics modeling on the basis of engine modeling,” in *SAE Technical Paper*. SAE International, 03 2004. [Online]. Available: <http://dx.doi.org/10.4271/2004-01-1620>
- [99] U. Kiencke and L. Nielsen, *Automotive control systems*. Springer, 2005, vol. 2.

- [100] R. Rajamani, *Vehicle dynamics and control*. Springer Science & Business Media, 2011.
- [101] A. McGordon, J. Poxon, C. Cheng, R. Jones, and P. Jennings, “Development of a driver model to study the effects of real-world driver behaviour on the fuel consumption,” *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering*, vol. 225, no. 11, pp. 1518–1530, 2011, cited By 14. [Online]. Available: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-84863011056&partnerID=40&md5=e963922e581378a267f9157c1bd013b6>
- [102] A. Fuhs, *Hybrid Vehicles: and the Future of Personal Transportation*. CRC Press, 2008. [Online]. Available: <https://www.crcpress.com/Hybrid-Vehicles-and-the-Future-of-Personal-Transportation/Fuhs/p/book/9781420075342>
- [103] K. Wipke, M. Cuddy, and S. Burch, “Advisor 2.1: a user-friendly advanced powertrain simulation using a combined backward/forward approach,” *Vehicular Technology, IEEE Transactions on*, vol. 48, no. 6, pp. 1751–1761, Nov 1999.
- [104] A. Fröberg and L. Nielsen, “Dynamic vehicle simulation -forward, inverse and new mixed possibilities for optimized design and control,” in *SAE Technical Paper*. SAE International, 03 2004. [Online]. Available: <http://dx.doi.org/10.4271/2004-01-1619>
- [105] L. Guzzella and A. Amstutz, “Cae tools for quasi-static modeling and optimization of hybrid powertrains,” *IEEE Transactions on Vehicular Technology*, vol. 48, no. 6, pp. 1762–1769, Nov 1999.
- [106] H. O. List and P. Schoeggl, “Objective evaluation of vehicle driveability,” in *SAE Technical Paper*. SAE International, 02 1998. [Online]. Available: <http://dx.doi.org/10.4271/980204>
- [107] R. E. Dorey and E. J. Martin, “Vehicle driveability - the development of an objective methodology,” in *SAE Technical*

- Paper.* SAE International, 03 2000. [Online]. Available: <http://dx.doi.org/10.4271/2000-01-1326>
- [108] V. Wicke, C. Brace, M. Deacon, and N. Vaughan, “Preliminary results from driveability investigations of vehicles with continuously variable transmissions,” *Proceedings: CVT*, pp. 9–14, 1999.
- [109] V. Wicke, C. J. Brace, N. D. Vaughan, and I. James, “The characterisation of driveability of cvt powertrains during acceleration transients,” *VDI Berichte*, pp. 311–325, 2002. [Online]. Available: <http://opus.bath.ac.uk/2579/>
- [110] D. F. Opila, D. Aswani, R. McGee, J. A. Cook, and J. W. Grizzle, “Incorporating drivability metrics into optimal energy management strategies for hybrid vehicles,” in *Decision and Control, 2008. CDC 2008. 47th IEEE Conference on*, Dec 2008, pp. 4382–4389.
- [111] D. F. Opila, X. Wang, R. McGee, R. B. Gillespie, J. A. Cook, and J. W. Grizzle, “An energy management controller to optimally trade off fuel economy and drivability for hybrid vehicles,” *IEEE Trans. Contr. Syst. Technol.*, vol. 20, no. 6, pp. 1490–1505, 2012. [Online]. Available: <http://dx.doi.org/10.1109/TCST.2011.2168820>
- [112] X. Wei and G. Rizzoni, “Objective metrics of fuel economy, performance and driveability - a review,” in *SAE Technical Paper.* SAE International, 03 2004. [Online]. Available: <http://dx.doi.org/10.4271/2004-01-1338>
- [113] O. Hayat, M. Lebrun, and E. Domingues, “Powertrain driveability evaluation: Analysis and simplification of dynamic models,” in *SAE Technical Paper.* SAE International, 03 2003. [Online]. Available: <http://dx.doi.org/10.4271/2003-01-1328>
- [114] J. Zehetner, P. Schggl, M. Dank, and K. Meitz, “Simulation of driveability in real-time,” in *SAE Technical Paper.* SAE International, 04 2009. [Online]. Available: <http://dx.doi.org/10.4271/2009-01-1372>

- [115] C. Röss, D. Balzer, A. Bracht, S. Durekovic, and J. Löwenau, “ADASIS protocol for advanced in-vehicle applications,” in *15th World Congress on Intelligent Transport Systems*, 2008.
- [116] J. P. Lauffenburger, B. Bradai, A. Herbin, and M. Basset, “Navigation as a virtual sensor for enhanced lighting preview control,” in *2007 IEEE Intelligent Vehicles Symposium*, June 2007, pp. 1–6.
- [117] ADASIS Forum. [Online]. Available: <http://erticoprojects.wpengine.com/projects/adasisforum/>
- [118] S. Durekovic, *ADASIS v2 Protocol*, ADASIS Forum Std.
- [119] P. Kock, R. Weller, A. W. Ordys, and G. Collier, “Validation methods for digital road maps in predictive control,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 16, no. 1, pp. 339–351, Feb 2015.
- [120] E. Wood, E. Burton, A. Duran, J. Gonder *et al.*, “Appending high resolution elevation data to gps speed traces for vehicle energy modeling and simulation,” *National Renew Energy Lab*, 2014.
- [121] D. E. Kirk, *Optimal Control Theory*. Englewood Cliffs, New Jersey: Prentice-Hall, 1970.
- [122] D. Bertsekas, *Dynamic programming and optimal control*. Belmont, Mass: Athena Scientific, 2005.
- [123] V. Monastyrsky and I. Golownykh, “Rapid computation of optimal control for vehicles,” *Transportation Research Part B: Methodological*, vol. 27, no. 3, pp. 219 – 227, 1993. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/0191261593900315>
- [124] E. Hellström, “Look-ahead control of heavy vehicles,” Ph.D. dissertation, Linköping University, Vehicular Systems, The Institute of Technology, 2010.
- [125] R. T. Wragge-Morley, G. Herrmann, S. Burgess, and P. Barber, “Modelling and simulation of rapidly changing road gradients,” *SAE*

- Int. J. Passeng. Cars - Mech. Syst.*, vol. 9, pp. 392–401, 04 2016.
[Online]. Available: <http://dx.doi.org/10.4271/2016-01-1663>
- [126] T. D. Gillespie, *Fundamentals of vehicle dynamics*. SAE International, 1992.
- [127] “Passenger car fuel economy, EPA and road: a report to the congress in response to the national energy conservation policy act of 1978, public law 95-619, title iv, part 1, section 404,” United States Environmental Protection Agency, Tech. Rep., 1980.
- [128] *Design Manual for Roads and Bridges: Volume 6 Road Geometry Section 1 Links TD 9/93 Highway Link Design*, The Highways Agency Std., 2002.
- [129] A. Pretschner, M. Broy, I. H. Kruger, and T. Stauner, “Software engineering for automotive systems: A roadmap,” in *Future of Software Engineering, 2007. FOSE '07*, May 2007, pp. 55–71.
- [130] (2016) Beagleboard black. BeagleBoard. [Online]. Available: <https://beagleboard.org/black>
- [131] (2016) Genuino mega 2560. Arduino. [Online]. Available: <https://www.arduino.cc/en/Main/ArduinoBoardMega2560>
- [132] Raspberry Pi Foundation. [Online]. Available: <https://www.raspberrypi.org/products/raspberry-pi-2-model-b/>
- [133] M. Broy, I. H. Kruger, A. Pretschner, and C. Salzmann, “Engineering automotive software,” *Proceedings of the IEEE*, vol. 95, no. 2, pp. 356–373, Feb 2007.
- [134] Software in the Public Interest, Inc. [Online]. Available: <https://www.debian.org/>
- [135] Raspberry Pi Foundation. [Online]. Available: <https://www.raspberrypi.org/downloads/raspbian/>
- [136] GPSD service daemon. [Online]. Available: <http://www.catb.org/gpsd/>

- [137] Python-can socketcan user documentation. Dynamic Controls. [Online]. Available: <https://python-can.readthedocs.io/en/latest/socketcan.html>
- [138] GNU Project. GCC, the GNU Compiler Collection. [Online]. Available: <https://gcc.gnu.org/>
- [139] Python Software Foundation. (2016) The python standard library - release 2.7.12. [Online]. Available: <https://docs.python.org/2/library/multiprocessing.html>
- [140] BU-353S4 USB GPS Receiver. US GlobalSat Inc. [Online]. Available: <http://usglobalsat.com/p-688-bu-353-s4.aspx>
- [141] A. Savitzky and M. J. E. Golay, “Smoothing and differentiation of data by simplified least squares procedures,” *Analytical Chemistry*, vol. 36, no. 8, pp. 1627–1639, 1964. [Online]. Available: <http://dx.doi.org/10.1021/ac60214a047>
- [142] W. Gander and J. Hrebicek, *Solving Problems in Scientific Computing Using Maple and MATLAB*. Springer-Verlag Berlin Heidelberg, 2004.
- [143] C. C. Robusto, “The cosine-haversine formula,” *The American Mathematical Monthly*, vol. 64, no. 1, pp. 38–40, 1957. [Online]. Available: <http://www.jstor.org/stable/2309088>
- [144] Google Maps API. [Online]. Available: <https://developers.google.com/maps/documentation/roads/speed-limits>
- [145] M. Haklay and P. Weber, “Openstreetmap: User-generated street maps,” *IEEE Pervasive Computing*, vol. 7, no. 4, pp. 12–18, Oct 2008.
- [146] HERE routing API developer’s guide. HERE. [Online]. Available: <https://developer.here.com/rest-apis/documentation/routing>
- [147] *Guidelines for the Use of the C Language in Critical Systems*, MISRA Std., 2012.

- [148] *PiCAN 2 User Guide*, 1st ed., SK Pang Electronics, 2016. [Online]. Available: http://skpang.co.uk/catalog/images/raspberrypi/pi_2/PICAN2UGB.pdf
- [149] C. Ebert and C. Jones, “Embedded software: Facts, figures, and future,” *Computer*, vol. 42, no. 4, pp. 42–52, April 2009.
- [150] *CAN Specification v2.0*, Robert Bosch Std., 1991. [Online]. Available: http://www.bosch-semiconductors.de/media/ubk_semiconductors/pdf_1/canliteratur/can2spec.pdf
- [151] BUSMASTERph. ETAS. [Online]. Available: http://www.etas.com/en/products/applications_busmaster.php
- [152] (2016) USB CAN Interface - Hardware Manual. IXXAT. [Online]. Available: <https://www.ixxat.com/docs/librariesprovider8/default-document-library/products/can-pc-interfaces/usb-to-can-v2-manual-english.pdf?sfvrsn=2>
- [153] M. Virbel, T. Hansen, and O. Lobunets, “Kivy—a framework for rapid creation of innovative user interfaces,” *Mensch & Computer 2011*, p. 69.
- [154] (2016) Raspberry pi touch display. [Online]. Available: <https://www.raspberrypi.org/products/raspberry-pi-touch-display/>
- [155] eSpeak text to speech. eSpeak. [Online]. Available: <http://espeak.sourceforge.net/>
- [156] “Average wind speed and deviations from the long term mean,” Department for Business, Energy & Industrial Strategy, Tech. Rep., 2016.
- [157] F.-A. Fortin, F.-M. De Rainville, M.-A. Gardner, M. Parizeau, and C. Gagné, “DEAP: Evolutionary algorithms made easy,” *Journal of Machine Learning Research*, vol. 13, pp. 2171–2175, Jul 2012.
- [158] Met office data archive. [Online]. Available: <https://data.gov.uk/metoffice-data-archive>

- [159] Google maps elevation API. Google. [Online]. Available: <https://developers.google.com/maps/documentation/elevation/intro>
- [160] M. K. Liubakka, D. S. Rhode, J. R. Winkelman, and P. V. Kokotovic, "Adaptive automotive speed control," *IEEE Transactions on Automatic Control*, vol. 38, no. 7, pp. 1011–1020, Jul 1993.
- [161] G. Buttazzo, *Hard real-time computing systems: predictable scheduling algorithms and applications*. Springer Science & Business Media, 2011, vol. 24.
- [162] Google Streetview API. Google. [Online]. Available: <https://developers.google.com/maps/documentation/streetview/>
- [163] T. Tielert, M. Killat, H. Hartenstein, R. Luz, S. Hausberger, and T. Benz, "The impact of traffic-light-to-vehicle communication on fuel consumption and emissions," in *Internet of Things (IOT), 2010*, Nov 2010, pp. 1–8.
- [164] R. Stanica, E. Chaput, and A.-L. Beylot, "Simulation of vehicular ad-hoc networks: Challenges, review of tools and recommendations," *Computer Networks*, vol. 55, no. 14, pp. 3179 – 3188, 2011, deploying vehicle-2-x communication. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1389128611001629>
- [165] F. Jiménez, W. Cabrera-Montiel, and S. Tapia-Fernandez, "System for road vehicle energy optimization using real time road and traffic information," *Energies*, vol. 7, no. 6, p. 3576, 2014. [Online]. Available: <http://www.mdpi.com/1996-1073/7/6/3576>
- [166] S. Zeadally, R. Hunt, Y.-S. Chen, A. Irwin, and A. Hassan, "Vehicular ad hoc networks (vanets): status, results, and challenges," *Telecommunication Systems*, vol. 50, no. 4, pp. 217–241, 2012. [Online]. Available: <http://dx.doi.org/10.1007/s11235-010-9400-5>
- [167] E. Adell and A. Vrheiyi, "Driver comprehension and acceptance of the active accelerator pedal after long-term use," *Transportation Research Part F: Traffic Psychology and Behaviour*, vol. 11, no. 1, pp.

- 37 – 51, 2008. [Online]. Available: [//www.sciencedirect.com/science/article/pii/S1369847807000319](http://www.sciencedirect.com/science/article/pii/S1369847807000319)
- [168] A. Chaloulakou, I. Mavroidis, and I. Gavriil, “Compliance with the annual {NO₂} air quality standard in athens. required {NO_x} levels and expected health implications,” *Atmospheric Environment*, vol. 42, no. 3, pp. 454 – 465, 2008. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1352231007008655>
- [169] *Council Directive 1999/30/EC of 22 April 1999 relating to limit values for sulphur dioxide, nitrogen dioxide and oxides of nitrogen, particulate matter and lead in ambient air*, European Union Std. Official Journal L 163, 29/06/1999, 4160, 1999.
- [170] S. dAmbrosio, R. Finesso, L. Fu, A. Mittica, and E. Spessa, “A control-oriented real-time semi-empirical model for the prediction of {NO_x} emissions in diesel engines,” *Applied Energy*, vol. 130, pp. 265 – 279, 2014. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0306261914005418>

Appendices

Appendix A

Vehicle Data

The parameters used in the vehicle model are listed in Table A.1 and the engine model is based on the performance data shown in Figure A.1 and Figure A.2.

Item	Value	Units	
Vehicle Mass	1929	kg	
Tyre Radius	0.369	m	
Gear Ratios	1 st	4.713	-
	2 nd	2.842	
	3 rd	1.909	
	4 th	1.382	
	5 th	1.000	
	6 th	0.808	
	7 th	0.699	
	8 th	0.580	
	9 th	0.480	
Final Drive Ratio	3.944	-	
Aerodynamic Coefficient	0.354	-	
Frontal Area	2.63	m ²	
Engine displacement	1991	cm ³	

Table A.1: Vehicle Model Data

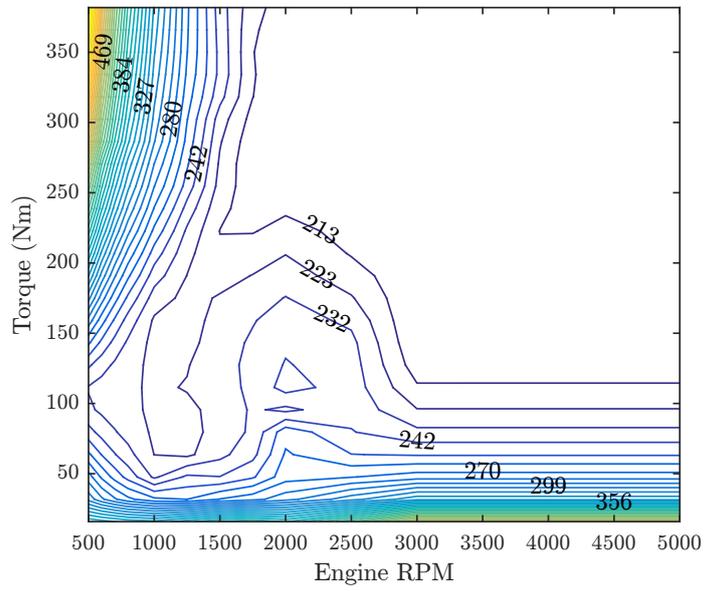


Figure A.1: Brake Specific Fuel Consumption with engine coolant 90 °C

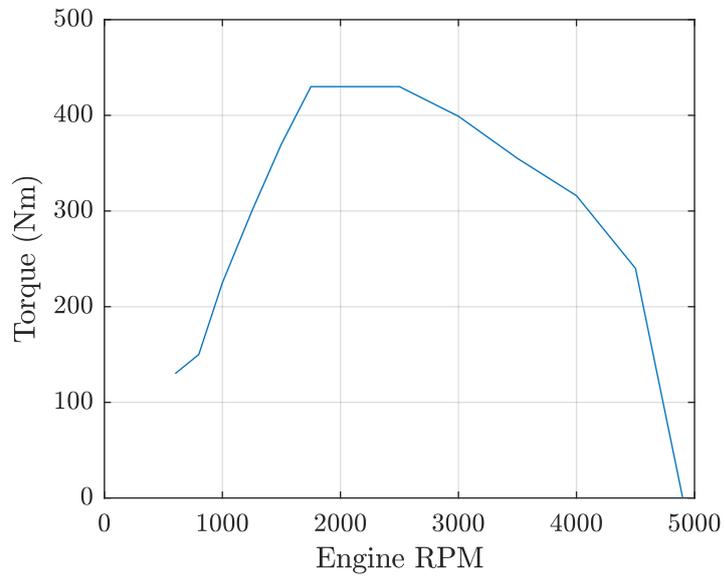


Figure A.2: Engine Maximum Torque Curve

Appendix B

Project Plan

ID	Task Name	Duration	Start	Finish	Predecessors	013																																						
						Half 1, 2014			Half 2, 2014			Half 1, 2015			Half 2, 2015			Half 1, 2016			Half 2, 2016																							
						O	N	D	J	F	M	A	M	J	J	A	S	O	N	D	J	F	M	A	M	J	J	A	S	O	N	D	J	F	M	A	M	J	J	A	S	O	N	D
1	PhD Start Date	0 days	Thu 03/10/13	Thu 03/10/13		03/10																																						
2	First Year	244 days?	Mon 14/10/13	Fri 03/10/14		[Task bar from 14/10/13 to 03/10/14]																																						
48	First Year Quarter Months 9-12	3.35 mons	Mon 30/06/14	Tue 30/09/14		[Task bar from 30/06/14 to 30/09/14]																																						
53	Second Year	290 days	Mon 06/10/14	Fri 13/11/15		[Task bar from 06/10/14 to 13/11/15]																																						
54	Second Year Start Date	0 days	Mon 06/10/14	Mon 06/10/14	21FS+10 wks	06/10																																						
55	Driver Related Research	255 days	Mon 06/10/14	Fri 25/09/15	54	[Task bar from 06/10/14 to 25/09/15]																																						
56	Benchmarking review	1 mon	Mon 06/10/14	Fri 31/10/14		[Task bar from 06/10/14 to 31/10/14]																																						
57	Driver Data Analysis (offline)	70 days	Mon 06/10/14	Fri 09/01/15		[Task bar from 06/10/14 to 09/01/15]																																						
58	Statistical Methods	6 wks	Mon 06/10/14	Fri 14/11/14		[Task bar from 06/10/14 to 14/11/14]																																						
59	Machine Learning Methods	6 wks	Mon 06/10/14	Fri 14/11/14		[Task bar from 06/10/14 to 14/11/14]																																						
60	Identification of patterns to be utilised by online module	2 mons	Mon 17/11/14	Fri 09/01/15	58,59	[Task bar from 17/11/14 to 09/01/15]																																						
61	Driver Analysis Module (online)	185 days	Mon 12/01/15	Fri 25/09/15		[Task bar from 12/01/15 to 25/09/15]																																						
62	Research applicable methods of online analysis	3 mons	Mon 12/01/15	Fri 03/04/15	57	[Task bar from 12/01/15 to 03/04/15]																																						
63	Plan module structure, interfaces and development process	1 wk	Mon 06/04/15	Fri 10/04/15	57,62	[Task bar from 06/04/15 to 10/04/15]																																						
64	Develop module	6 mons	Mon 13/04/15	Fri 25/09/15	62,63	[Task bar from 13/04/15 to 25/09/15]																																						
65	Preliminary testing	2 wks	Mon 07/09/15	Fri 18/09/15	62FS+5.5 mons	[Task bar from 07/09/15 to 18/09/15]																																						
66	Vehicle Model Development	260 days	Mon 06/10/14	Fri 02/10/15	54	[Task bar from 06/10/14 to 02/10/15]																																						
67	Integration of Simulink models with chosen commercial software package	1 mon	Mon 06/10/14	Fri 31/10/14		[Task bar from 06/10/14 to 31/10/14]																																						
68	Model Validation	3 wks	Mon 14/09/15	Fri 02/10/15		[Task bar from 14/09/15 to 02/10/15]																																						
69	Control System Research	180 days	Mon 06/10/14	Fri 12/06/15	54	[Task bar from 06/10/14 to 12/06/15]																																						
70	Electronic Horizon adaption	2 mons	Mon 06/10/14	Fri 28/11/14		[Task bar from 06/10/14 to 28/11/14]																																						
71	Comparison of Dynamic Programming with other optimal control methods	3 mons	Mon 06/10/14	Fri 26/12/14		[Task bar from 06/10/14 to 26/12/14]																																						
72	Formalise problem - constraints, cost function	0 days	Fri 26/12/14	Fri 26/12/14	71	26/12																																						
73	Development of optimisation algorithm in Simulink and C-Code	6 mons	Mon 29/12/14	Fri 12/06/15	71	[Task bar from 29/12/14 to 12/06/15]																																						
74	Publications	140 days	Fri 06/02/15	Thu 20/08/15	54	[Task bar from 06/02/15 to 20/08/15]																																						
75	Offline Driver Data Analysis Research Paper	1 mon	Fri 06/02/15	Thu 05/03/15	57	[Task bar from 06/02/15 to 05/03/15]																																						
76	Driveability Metrics for Control System Applications	0 days	Thu 20/08/15	Thu 20/08/15	75FS+6 mons	20/08																																						
77	Quarterly Reports	180 days	Fri 19/12/14	Thu 27/08/15		[Task bar from 19/12/14 to 27/08/15]																																						
82	Annual Monitoring Report	1 mon	Mon 15/06/15	Fri 10/07/15	54FS+9 mons	[Task bar from 15/06/15 to 10/07/15]																																						
83	End of Second Year	0 days	Wed 30/09/15	Wed 30/09/15		30/09																																						
84	Third Year	280 days	Mon 05/10/15	Fri 28/10/16		[Task bar from 05/10/15 to 28/10/16]																																						
85	Third Year Start Date	0 days	Mon 05/10/15	Mon 05/10/15	82FS+12 wks	05/10																																						
86	Driver Analysis System	140 days	Mon 24/08/15	Fri 04/03/16		[Task bar from 24/08/15 to 04/03/16]																																						
87	Implementation on hardware	3 mons	Mon 24/08/15	Fri 13/11/15	62FS+5 mons	[Task bar from 24/08/15 to 13/11/15]																																						

Project: PhD Project Plan
Date: Mon 14/09/15

Task		Project Summary		Inactive Milestone		Manual Summary Rollup		Deadline	
Split		External Tasks		Inactive Summary		Manual Summary		Progress	
Milestone		External Milestone		Manual Task		Start-only			
Summary		Inactive Task		Duration-only		Finish-only			

Appendix C

Expanded Results Data

		R1						R2									
		Fuel		Time		Cost		%		Fuel		Time		Cost		%	
Velocity		Fixed	DP	Fixed	DP	Fixed	DP	Fixed	DP	Fixed	DP	Fixed	DP	Fixed	DP	Fixed	DP
20		62.7	86.9	72.5	50.7	44.1	36.2	17.9		57.1	83.5	72.5	50.7	43.4	35.9	17.3	
22		70.0	81.4	65.9	50.7	41.7	35.4	15.0		64.7	78.1	65.9	50.6	41.0	35.1	14.6	
24		75.8	77.2	60.4	50.3	39.7	34.6	12.7		72.1	73.1	60.4	50.5	39.2	34.3	12.5	
26		82.6	71.3	55.8	50.2	38.2	33.8	11.5		75.0	68.2	55.8	50.2	37.3	33.5	10.2	
28		90.7	64.8	51.8	50.2	37.2	32.9	11.6		80.6	62.7	51.8	50.1	36.0	32.6	9.3	
30		97.4	58.2	48.3	50.1	36.3	32.0	12.0		88.3	55.2	48.3	50.1	35.2	31.7	10.0	
		R3						R4									
		Fuel		Time		Cost		%		Fuel		Time		Cost		%	
Velocity		Fixed	DP	Fixed	DP	Fixed	DP	Fixed	DP	Fixed	DP	Fixed	DP	Fixed	DP	Fixed	DP
20		56.9	71.5	72.5	51.6	43.4	34.8	19.7		51.6	79.0	72.5	49.9	42.7	34.9	18.3	
22		61.6	67.1	65.9	51.4	40.7	34.1	16.2		58.1	73.3	65.9	50.1	40.2	34.2	15.0	
24		67.1	61.4	60.4	51.3	38.6	33.2	13.9		65.0	68.5	60.4	49.8	38.3	33.4	12.9	
26		71.5	55.9	55.8	51.2	36.8	32.4	12.0		68.9	63.4	55.8	49.6	36.5	32.6	10.7	
28		74.9	50.4	51.8	50.9	35.3	31.5	10.6		73.3	56.2	51.8	49.7	35.1	31.7	9.6	
30		80.3	43.3	48.3	50.9	34.2	30.6	10.6		79.2	50.4	48.3	49.5	34.1	30.8	9.7	

Table C.1: Total cost results from fixed velocity and unconstrained DP policies on real road profiles R1 to R4 with initial velocities from 20 m s^{-1} to 30 m s^{-1} . The cost difference between the two policies in percent is shown in bold.

		R1						R2									
		Fuel		Time		Cost		%		Fuel		Time		Cost		%	
Velocity		Fixed	DP	Fixed	DP	Fixed	DP	Fixed	DP	Fixed	DP	Fixed	DP	Fixed	DP	Fixed	DP
20		62.7	80.1	72.5	54.2	44.1	37.8	14.3		57.1	84.9	72.5	52.3	43.4	37.2	14.2	
22		70.0	74.9	65.9	54.0	41.7	36.9	11.5		64.7	74.4	65.9	53.5	41.0	36.1	12.0	
24		75.8	77.8	60.4	52.0	39.7	36.0	9.2		72.1	75.1	60.4	51.1	39.2	35.0	10.7	
26		82.6	66.7	55.8	53.1	38.2	34.8	8.8		75.0	70.6	55.8	50.5	37.3	34.1	8.6	
28		90.7	68.2	51.8	50.2	37.2	33.6	9.8		80.6	49.6	51.8	54.3	36.0	33.4	7.2	
30		97.4	54.5	48.3	51.9	36.3	32.5	10.4		88.3	51.8	48.3	51.1	35.2	32.1	8.9	
		R3						R4									
		Fuel		Time		Cost		%		Fuel		Time		Cost		%	
Velocity		Fixed	DP	Fixed	DP	Fixed	DP	Fixed	DP	Fixed	DP	Fixed	DP	Fixed	DP	Fixed	DP
20		56.9	70.0	72.5	53.7	43.4	36.6	15.6		51.6	81.6	72.5	51.5	42.7	36.2	15.2	
22		61.6	64.8	65.9	53.5	40.7	35.7	12.2		58.1	77.7	65.9	50.9	40.2	35.3	12.1	
24		67.1	58.2	60.4	53.5	38.6	34.7	10.1		65.0	69.9	60.4	50.4	38.3	34.4	10.4	
26		71.5	53.5	55.8	52.7	36.8	33.6	8.8		68.9	59.2	55.8	50.7	36.5	33.4	8.6	
28		74.9	47.1	51.8	52.6	35.3	32.7	7.1		73.3	56.1	51.8	50.0	35.1	32.4	7.6	
30		80.3	38.2	48.3	52.9	34.2	31.7	7.3		79.2	46.6	48.3	50.5	34.1	31.2	8.4	

Table C.2: Total cost results from fixed velocity and constrained DP policies on real road profiles R1 to R4 with initial velocities from 20 m s^{-1} to 30 m s^{-1} . The cost difference between the two policies in percent is shown in bold.