Applications of Optimal Control for Emissions Management in Road Vehicles



Gareth J. Pease

New College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity 2019

I'r rhai a gollwyd ar hyd y ffordd:
Bert, Molly, Peter and Cyril

Acknowledgements

I would like the acknowledge the help of a number of different people who aided in more ways than they can possibly imagine to make this thesis happen. Whether through distractions away from work, discussions of the finer details of control in general.

Firstly I would like to thank Professor David Limebeer who supervised this thesis. His knowledge and expertise was invaluable during all stages of the work completed here. Also for his belief in me to allow me to pursue a DPhil after spending year supervising me as an undergraduate.

Secondly, I would also like to acknowledge the help of Peter Fussey and Justin Seabrook for their help with the specifics of engine modelling and control. Especially Peter who has been a second supervisor during this four year process and granted me the opportunity to learn the ropes down at Ricardo.

The day to day study has been made easier with good colleagues in the Control Group, both permanent and visiting. Particularly to talk vehicle dynamics and optimal control with Mehdi, Sebastian and Luca.

Thanks to Jon for the countless discussions and all of the time spent collecting sets of vehicle measurement data. This was much appreciated.

Outside of engineering, thanks need to be given to all those in the Middle and Senior Common Rooms at New College. Special thanks to Katie, George, Will, Arup and Charlie for the inane chats and pub trips.

An acknowledgement has to be made to the engineering fellows at New College who allowed me the opportunity to inspire the next generation of young engineers.

I would like to thank my family for their support during, what must have been, the longest four years. Especially to my father, Mark, who has provided inspiration throughout my studies and my mother, Catherine, who has proofread this thesis.

A special mention has to made to Epstein and Barr who have made writing this thesis a memorable journey.

Abstract

Improvements in the availability of sensed and actuated vehicular systems have increased the opportunity for applying automatic control. Classic methods for designing control architectures are no longer sufficient due to the highly non-linear behaviour and heavy coupling between sub-systems. As a result, there is scope to employ computationally powerful optimisation techniques to investigate design trade-offs in these increasingly complex systems.

This thesis intends to take a multi-faceted look into the use of powerful optimal control techniques to solve some of the control problems arising in the automotive sector. In particular, a number of different control problems focused on emissions management are considered.

The first scenario relates to the fuel optimal control of a road vehicle, with a specific focus generating fuel-optimal drive cycle trajectories incorporating emissions constraints. A simplified model will be used initially to examine the solution structure and highlight any important features which occur in the solving of such eco-driving problems. Additionally, there are comparisons to be made between using a continuously variable transmission and a more traditional discrete gearbox.

The second scenario considers an optimal driving problem for a vehicle equipped hybrid electric powertrain with a thermal aftertreatment model. A large percentage of tailpipe emissions are produced when the three-way catalytic converter is cold. Increased flexibility of power delivery enables the engine to be powered greater than the tractive demand, which prompts quicker heating and thus producing less total emissions. The drawback to this method is an increase in the tailpipe flowrates during this heating period. The possibility and scope for the reduction of emissions is considered.

The final scenario considered in this thesis will cover the optimal control of gasoline engines with combustion mode switching events. In particular, a turbocharged gasoline engine equipped dual-independent variable valve timing is considered. For engines operating in a stoichiometric regime, three-way catalytic converters are effective at reducing emissions. If the combustion mode is switched to lean, there is an increase in the tailpipe nitrous oxides (NO_x). In order to avoid this, a lean-NO_x trap is installed in the tailpipe, which requires periodic regeneration to ensure its effectiveness. Dynamic compensation of all actuators vastly improves the performance of the switching control.

This thesis presents a number of suitable modelling techniques to investigate these scenarios, the results of which can be used to aid in the design of real-time control algorithms.

Contents

List of Figures **xiii**

Abbreviations **xix**

Notation **xxi**

1 Introduction **1**

1.1 Vehicle Design and Testing 2

1.1.1 Internal Combustion Engines 3

1.1.2 Hybridisation 6

1.1.3 Vehicle Testing 8

1.2 Motivation and Structure 10

1.3 Related Work 13

2 Numerical Optimal Control **15**

2.1 Optimisation 16

2.1.1 Unconstrained Optimality 17

2.1.2 Iterative Methods 18

2.1.3 Constrained Optima 20

2.1.4 Numerical Methods for Nonlinear Progams 23

2.2 Optimal Control 26

2.2.1 Optimality Conditions 27

2.2.2 Constrained Minimum Principle 29

2.3 Numerical Methods for Optimal Control 32

2.3.1 Shooting Methods 33

2.3.2 Numerical Integration 35

2.3.3 Indirect Methods 37

2.3.4 Direct Methods 38

2.3.5 Numerical Considerations 45

2.4 Conclusion 52

3 Minimum Energy Optimal Control 55
3.1 Minimum energy optimal control . . . 57
3.1.1 Arrival time constraint . . . 65
3.1.2 Slope Effects . . . 69
3.1.3 Computational Validation . . . 71
3.2 Variable Velocity Limit . . . 74
3.3 Variable Road Grade . . . 78
3.4 Conclusions . . . 84

4 Fuel-Optimal Driving 87
4.1 Car Model . . . 90
4.1.1 Vehicle Dynamics . . . 90
4.1.2 Powertrain Model . . . 91
4.2 Driving Optimal Control Problem . . . 97
4.2.1 Multi–phase Numerical Optimal Control . . . 101
4.2.2 Discontinuous Functions . . . 102
4.3 Results . . . 103
4.3.1 Power Transmission Comparison . . . 103
4.3.2 Emissions Sensitivity . . . 107
4.3.3 Enforced Stops . . . 110
4.4 Conclusions . . . 114

5 Optimal Hybrid Vehicle Driving and Aftertreatment Warmup 117
5.1 Model . . . 120
5.1.1 Vehicle Model . . . 120
5.1.2 Powertrain Model . . . 122
5.1.3 Engine Model . . . 125
5.1.4 Energy Store . . . 129
5.1.5 Three-way Catalytic Converter . . . 130
5.1.6 Route Data . . . 131
5.1.7 Optimal Control Problem Setup . . . 132
5.2 Results . . . 134
5.2.1 Baseline Study: Fuel-Optimal . . . 135
5.2.2 Case Study 1- Emissions Constraints . . . 141
5.2.3 Case Study 2 - Route B2 . . . 148
5.3 Conclusions . . . 153

6 Optimal Control of Gasoline Engine Combustion Mode Switching 157
6.1 Engine Model . . . 160
6.1.1 Cylinder Model . . . 166
6.1.2 Emissions formation . . . 171
6.1.3 Aftertreatment Devices . . . 172
6.2 Optimal Control Formulation . . . 173
6.2.1 Scaling . . . 175
6.3 Optimisation Results . . . 176
6.3.1 Constant power . . . 176
6.3.2 Drive Cycle . . . 181
6.4 Conclusions . . . 184

7 Conclusions 187
7.1 Future work . . . 189

Appendices

A Distance Domain Reformulation 197
A.1 Time domain . . . 197
A.1.1 Distance domain . . . 198

B Estimation of Route Data 203
B.1 Road Curvature . . . 204
B.1.1 Results . . . 206
B.2 Road Grade . . . 210

C Velocity Profile Validation 215

References 219

List of Figures

1.1 Predicted split of electric vehicles and internal combustion engined vehicles for passenger vehicles [5]. 2
1.2 Complete vehicle overview including powertrain and control and feedback signals. 3
1.3 Possible hybrid powertrain configurations. Electrical power connections are red and mechanical power links are blue. Arrow directions indicate the allowable power flow directions. 7
1.4 An example velocity profile generated in Chapter 4 (solid) compared to the NEDC (dashed) and the WLTC (dotted). 9

2.1 A comparison of two approximations of the $\max(x, 0)$ function. The hyperbolic approximation (2.152) is red, the square-root approximation (2.151) is blue and the original function is black. 52

3.1 A vehicle ascending a slope of constant incline. 57
3.2 Bang-singular-zero control (solid line), the switching function Φ (dashed line) and the velocity (dotted line). 62
3.3 Left: This is a figure to show the variation of switching time and arrival time for a sweep of θ. The contours are plots of equation (3.26). Right: figure shows the variation of the final velocity on the same slopes with the minimum traversal time final velocity indicated. The contours are plots of (3.30). These results are for the unconstrained case. 69
3.4 Left: Variation of the second switching time τ_2 in the velocity constrained case (solid coloured). The first switching times are dotted lines. Right: Variation of the final velocity v_f for a set of arrival times and road grades. 70
3.5 A comparison of the numerically computed solution to the problem (circles) compared to the analytically derived solutions of the previous section. The error between each state, control and costate is plotted below each. 72

3.6 Collocation point locations (crosses) and mesh points for each mesh iteration (circles). The legend states the order of each mesh interval (and thus number of collocation points). 73
3.7 The lemniscate for the lateral acceleration limited problem. The start point is at the bottom of the track at the point $(a, 0, 0)$ and ends directly above at $(a, 0, 10)$. The track is 100m long. 76
3.8 Left: Optimal velocity profile for the curvature constrained case in solid with the velocity limit plotted as the black dashed line. Right: Optimal force trajectory. 76
3.9 Left: Variation of the optimal energy input for the ascent based on the final arrival time. Minimum energy E^* is dotted. Right : Computed Hamiltonian for the ascent under variable velocity limits. 77
3.10 Example of an undulating road with elevation $y = 15$m at $x = 350$m. The peak at (A) has elevation $y = 10$m at $x = 100$m, while the dip at (B) has elevation $y = 5$m at $x = 250$m. The road length is $s_f = 352$m. 79
3.11 Left: The optimal velocity profile for an arrival time of 26.5s. Right: The optimal force control profile with the singular force (dashed). . 80
3.12 Left: The effect of maximum arrival time against the optimal energy input cost. The minimum cost, E^*, is plotted as a dotted line. Right: The terminal velocity variation of the vehicle against the arrival time. 81
3.13 Driving and braking force with optimal dissipative braking. The component of the weight parallel to the slope is dashed. 82
3.14 Left: The effect of maximum arrival time against the optimal energy input cost. The minimum cost, E^*, is plotted as a dotted line. The braking energy is plotted as a dashed line. Right: The terminal velocity variation of the vehicle against the arrival time. 83

4.1 The closed driving route (denoted Route A) around Oxford, UK, utilised in this chapter. 89
4.2 Forces acting on the vehicle on a general undulating road. 90
4.3 Brake specific FC (BSFC) map [g/kWh] including maximum torque line and lines of constant power. 92
4.4 Brake specific emissions maps [g/kWh] (Left - unburnt hydrocarbons (HC), Right - NO_x; the maximum torque line is overlaid. 94
4.5 Brake specific emissions maps [g/kWh] (Left - soot, Right - CO); the maximum torque line is overlaid. 95
4.6 Transmission layouts for both optimisation scenarios. 95
4.7 Engine speed as a function of forward velocity; the overall drive ratios are shown as dotted lines. The condition (4.15) is clear as the horizontal line up to the dashed line corresponding to first gear. . . 96

4.8 A diagram of how the distance based speed limit information is shown here. Each speed limit is separated by a boundary where either the speed is continuous or required to be zero. An admissible speed trajectory is shown as the dashed black curve. 101
4.9 Optimal velocity and power profiles for CVT (blue dotted) and AMT (red solid) transmission systems. Combined legislative and curvature-related speed limits are plotted as a dashed black line on the velocity figure. The route elevation is plotted as a dot-dash curve on the power figure. Individual phases are represented by alternating light grey regions. 104
4.10 Scatter of optimal engine operating points on the BSFC [g/kWh] map. The CVT results are plotted as circles with the AMT plotted as crosses. 105
4.11 Optimisation gear ratio results 106
4.12 Total percentage changes in each of: fuel, NO_x CO HC and PM for each of the constrained optimisations. 107
4.13 Segments of the emissions-constrained optimal driving solution for the 5km—8km section, and the highway driving section 14km—18km. The baseline fuel-optimal solution is plotted in red. The 90% constraints cases on NO_x CO, HC and PM are shown, respectively, as dot-dash, dotted, solid, and dashed. 108
4.14 Optimal velocity and power profiles for zero stops (black dotted), 7 random stops - Run 4 in Table 4.3 (blue solid) and all stops (red dot-dashed). Velocity limits are plotted as dashed black lines. . . . 112
4.15 Percentage changes in the total masses of fuel, NO_x, CO, HC and PM relative to driving without stopping. 113
4.16 Percentage changes in the shifting speeds relative to driving without stopping. 114

5.1 Driving route for this chapter, denoted Route B. 120
5.2 A model diagram of the available power flows in the hybrid power-train under consideration. Mechanical power links are blue arrows, electrical links are yellow arrows. 122
5.3 The layout of the electrical machines and internal combustion engine gearing. The ICE is coupled to the planetary carrier and planet gears (shown in red), the MU is attached to the ring gear (shown in green) and the GU is attached to the sun gear (shown in blue). 123

5.4 The relation between the rotational speeds of the three power devices for power generation. The motor speed is dictated by the vehicle speed, and the selection of either the generator speed fixes the engine speed or vice versa. The black line represents an example solution to (5.6). The region in which the powertrain can operate in all electric mode is shown in blue. The red region shows the available operating points for the hybrid operating mode. The engine is required to start to move from the blue to the red dashed line. The maximum all-electric speed is approximately 60km/h. 124
5.5 Polynomial efficiency maps for MU (left) and GU (right). Contours are levels of constant efficiency. Power limits for both machines are plotted as black dashed lines. 126
5.6 Emissions mass rate flow maps [g/s]. Left-to-right: CO, HC and NOx. 129
5.7 Catalytic converter temperature dependencies. 131
5.8 Curvature and elevation for Route B. 132
5.9 Top: Optimal speed (km/h) profile (solid). The route speed limits are dotted, taking the minimum of the legislative speed limit and the lateral acceleration limit. Middle: Tractive power (kW) is plotted along with the power delivery from each device. Bottom left: The SOC [%] of the energy store is plotted. Bottom right: TWC temperature. 137
5.10 Left: Motor unit operating points on contours of constant efficiency. Centre: Generator unit operating points on contours of constant efficiency. Right: Engine operating points overlaid on fuel map. The engine torque is net torque - whenever it is negative it corresponds to that required to start the engine. 138
5.11 Red shows the flow rates of each emissions species: solid is engine out (Equations (5.22),(5.23) and (5.24)), and dashed is tailpipe (Equations (5.37), (5.38) and (5.39)). The blue trace is the cumulative tailpipe emissions. 139
5.12 Total fuel consumption and total regularisation percentage against the regularisation weight on the engine speed. 140
5.13 Changes in the engine switching strategy (by comparison of the optimal engine speed trajectories) for different cases of r_4 in(5.48). 141
5.14 HC Constraint Results 143
5.15 CO Constraint Results 144
5.16 NO_x Constraint Results 145
5.17 Comparison of engine output powers for each of the limiting cases for each emissions type. 146

5.18 Changes to the driving part of the optimisation. The baseline trajectory is in black. Limiting constraints are imposed, NO_x is dashed, CO is dotted and HC is dot-dashed. 146
5.19 Red is the original result of the previous example (Route B1) and blue is the same route in reverse (Route B2). 148
5.20 Engine power trajectories for the limiting emissions reduction cases for Route B2. 149
5.21 Comparison of the percentage changes in fuel consumption for each of three emissions constraints for the traversal of the route in both directions. Red - NO_x, blue - CO, green - HC. 150
5.22 Part of the driving changes arising from the limiting constraint cases for Route B2. 152

6.1 Engine controller and signal flow layout. Engine speed N is a measured input to the system and is governed by the vehicle speed. Higher controllers determine the requested torque, $\tilde{\tau}_e$ and tailpipe λ (shown in red). 158
6.2 The engine model layout. Control input signals are in red, output signals are in blue. 161
6.3 Turbocharger maps. 163
6.4 A comparison of the true flow rate correction function in red (Equation (6.13)), against the approximation function in black (Equation (6.16)). 165
6.5 Intake side empirical functions taken from [99]. 167
6.6 Maps of engine mass flow rate corrections taken from [99]. 168
6.7 Torque lambda efficiency. 169
6.8 Engine maps for temperature and spark timing. 170
6.9 Engine torque maps 170
6.10 Engine emissions maps, contours are mass flow rates [g/s]. 171
6.11 Left: The dependence of engine out emissions on in-cylinder λ_e. Right: The dependence of TWC reduction efficiency on the local air quality λ_c. 172
6.12 Airpath state trajectory comparison for the optimised controller action (solid blue), compared to the baseline (P+I) (dashed black). 177
6.13 Control actions for baseline (black dashed) and optimised (solid blue). 178
6.14 Output tracking trajectories for the constant power output case. The left hand figure compares the optimised trajectories for λ_e (red), λ_c blue dotted compared to the baseline (black dashed). The right hand figure compares the torque tracking for the baseline controller (black) and optimised trajectory (blue). 179

6.15 Mass flow rates of NO_x, HC and CO (descending); engine out on the left and out of the tailpipe on the right. Response under baseline control is black dashed. Dynamic optimised response is solid blue. . 180
6.16 The drive cycle engine torque (blue) and speed (red) trajectories for the transient optimisation case. 181
6.17 Optimal control trajectories (blue) for the drive cycle case compared to the nominal baseline controller (black dashed). 182
6.18 Output tracking trajectories - left is the tailpipe λ, blue is optimised, black dashed is baseline. Right is output torque, blue is optimised, black dashed is baseline control. The torque target is plotted as dotted green. 182
6.19 Emissions trajectories for the drive cycle case. The optimised trajectories are blue, the baseline controller are black dashed. 183

B.1 The geometry of the differential transformation required to estimate the road curvature. The road is drawn as the dashed line. The tangent to the road at point A is $\mathbf{t}$, and it makes an angle ψ with the global x-axis. 205
B.2 Route reconstruction with the original GPS data as black crosses. The insets demonstrate over-fitting of the data if w_c is chosen too large. 208
B.3 Curvature estimates plotted as a function of distance along the spine. 209
B.4 Route B elevation reconstruction based on the solution to the optimal filtering problem. Highlights on the right-hand side show the effects of over or under-reconstruction. 212
B.5 Comparison of the rapidly changing inclination angles for the elevation reconstruction. Insets are at the same points as Figure B.4. . . 213
B.6 Route A estimated curvature and inclination. 214

C.1 Equipment used for the validation in this appendix. 216
C.2 The original GPS (red) and OBD (blue) data is plotted in the top graph. Bottom: Scaled GPS (red) and OBD (blue) data, is compared to one simulation run for all stops (black). 217
C.3 Adjusted logged data for Route A (dotted). Original speed limit information is red and an optimisation required to stop at all stops ($T = 2000$s) is plotted in solid blue. 218

Abbreviations

AMT . . . . . Automated Manual Transmission
BEV . . . . . Battery Electric Vehicle
CAN . . . . . Controller Area Network
CO . . . . . Carbon Monoxide
CVT . . . . . Continuously Variable Transmission
ECU . . . . . Engine Control Unit
ECMS . . . . . Equivalent Consumption Minimisation Strategy
EM . . . . . Electrical Machine
ES . . . . . Energy Store
FC . . . . . Fuel Consumption
GPM . . . . . Gaussian Process Model
GU . . . . . Generator Unit
HC . . . . . Unburnt Hydrocarbons
HEV . . . . . Hybrid Electric Vehicle
ICE . . . . . Internal Combustion Engine
LNT . . . . . Lean NO_x Trap
MU . . . . . Motor Unit
NEDC . . . . . New European Drive Cycle
NO_x . . . . . Nitrous Oxides
PMP . . . . . Pontryagin's Maximum (Minimum) Principle
OBD . . . . . Onboard Diagnostics
OCP . . . . . Optimal control problem
PM . . . . . Particulate Matter
RDE . . . . . Real Driving Emissions
TPBVP . . . . Two-point Boundary Value Problem
TWC . . . . . Three-way Catalytic Converter
WLTC . . . . . World Light harmonised Test Cycle

Notation

In this thesis, the following notation is employed.

Lower case letters such as x, denoted scalar quantities and functions. The notation $\mathbb{R}$ denotes the set of real numbers, $\mathbb{R}^n$ is the set of real value vectors of dimension n. Bold face lower case represent column vectors $\mathbf{x} = [x_1, x_2, \ldots, x_n]^\top$. A scalar function mapping the vector x, has gradient denoted by the row vector

$$\nabla f(\mathbf{x} = \Big[\frac{\partial f}{\partial x_1}, \frac{\partial f}{\partial x_2}, \ldots, \frac{\partial f}{\partial x_n}\Big].$$

A vector valued function is written

$$\mathbf{f}(\mathbf{x}) = [f_1(\mathbf{x}), f_2(\mathbf{x}), \ldots, f_n(\mathbf{x})]^\top.$$

The derivative of this an $n \times n$ square matrix. An uppercase bold character is a matrix $\mathbf{X}$, whose entries of this matrix are denoted x_{ij}.

Dots refer to the derivative with respect to the independent variable $\dot{\mathbf{x}}$, which is either time, t or elapsed distance s. Superscript star refers to optimal values of parameters or functions, for example $\mathbf{x}^*$.

...the substitution of the internal combustion engine for the horse marked a very gloomy milestone in the progress of mankind.

— Winston Churchill, Speaking in the House of Commons, June 1952

1 Introduction

Contents

1.1 Vehicle Design and Testing 2
1.1.1 Internal Combustion Engines 3
1.1.2 Hybridisation 6
1.1.3 Vehicle Testing 8
1.2 Motivation and Structure 10
1.3 Related Work 13

Currently the road vehicle industry is under heavy scrutiny. Concern over climate change and air quality has led to a shift in public opinion away from internal combustion engines (ICEs). This has been driven in part by sensationalist news items [1] and a desire for lower emissions in urban areas [2]. Despite the push for new technologies, the market uptake of green technologies is slow: 90% of ground-transportation energy will continue to come from ICEs in the year 2040 [3].

Figure 1.1 shows a prediction market share of consumer vehicles by propulsion type. While the total penetration of battery electric vehicles (BEV) is expected to increase, there is still a large percentage of total vehicles requiring ICEs. Hybrid electric vehicles (HEV) will still require an ICE. It is clear that the use of ICEs cannot be completely avoided in the near future.

The quality of the atmosphere, and particularly the production of harmful

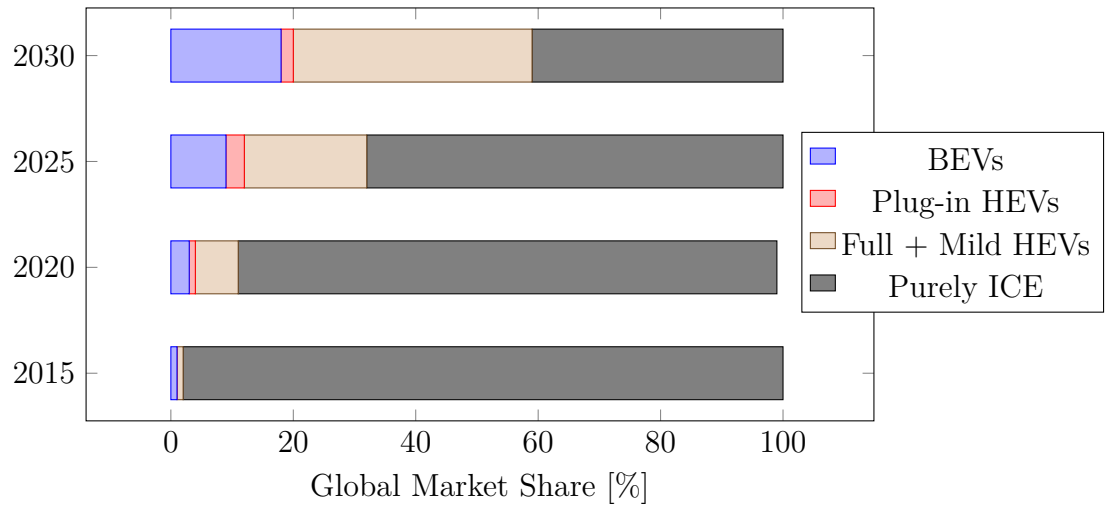


Figure 1.1: Predicted split of electric vehicles and internal combustion engined vehicles for passenger vehicles [5].

emissions by road vehicles has become an important topic in recent years. The gap between on the road emissions and those during testing has increased over the last 18 years [4]. Furthermore, fuel economy is poorer in real driving scenarios than quoted during tests. This has further tarnished the public opinion of ICEs.

As ICE propelled vehicles will be required for the foreseeable future, it is now the remit of engineers to find favourable ways to reduce emissions, improve efficiency and ensure the cost effectiveness of latest generation ICEs. Not only do improvements to hardware effect the performance of ICEs, better control schemes and greater flexibility of actuated and sensed systems can further abate harmful pollutants. Further changes to testing and design stages are also required to ensure better correlation between tests and real world driving. This brings with it a number of control challenges.

1.1 Vehicle Design and Testing

Modern powertrains now come in an increasing number of configurations and designs. A larger disconnect is seen between the driver and powertrain in modern road vehicles as an increased number of decisions are taken by the vehicle itself. Power is required not only for propulsion, but also to drive ancillary devices in the passenger cabin. A

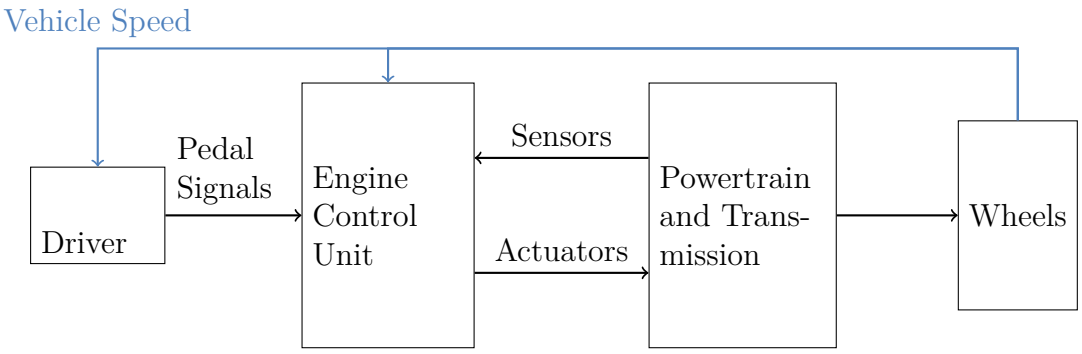


Figure 1.2: Complete vehicle overview including powertrain and control and feedback signals.

typical powertrain layout is shown in Figure 1.2. Pedal requests from the driver are acted upon by the engine control unit (ECU). This onboard computer decides how best to actuate the powertrain based on a number of feedback sensors, obtaining powertrain information and ambient conditions. The challenge for control engineers is ensuring the ECU can take full advantage of new hardware and system design by using state of the art control techniques. Hardware changes include novel/modern combustion systems and multiple energy sources (in the case of HEVs).

1.1.1 Internal Combustion Engines

At a high level, the internal combustion engine within an automotive powertrain converts the chemical potential energy stored in liquid fuel into work done to move the vehicle. This work overcomes friction in the powertrain, and that due to air and rolling resistances. The total stored energy in the fuel is not completely transformed into work; some is used for pumping and heat transfer within the ICE.

Inside the ICE, air is fed into a set of cylinders by the air-path which includes turbochargers, intercoolers, throttles and filters. This air is mixed with fuel inside the cylinder and burned to produce power. Two separate fluid systems arise - one for fuel, and one for air. During the combustion process gaseous products are formed and released by the tailpipe into the environment. In order to reduce this effect, the exhaust gas can be cleaned by aftertreatment devices. Three-way catalytic converters (TWC) have been standard in road vehicles for a number of years. These devices make use of a metal bed to create surfaces upon which conversion of gases

may take place [6]. This relies on the presence of oxygen in the exhaust flow, which is in turn related to the air to fuel ratio in the cylinders. Lambda (λ) is the metric used to determine the combustion regime and is defined as the ratio of the current air-fuel-ratio (AFR) to that at stoichiometry. Standard gasoline engines operate at stoichiometry ($\lambda = 1$). While TWC are effective at reducing tailpipe emissions, most of the required reactions require elevated temperatures to achieve good reduction. During the period while the catalytic beds are cold, a large amount of the total tailpipe emissions are produced.

One means to improve the fuel economy and reduce emissions is to make the engine smaller - this is referred to as *downsizing* [7]. A clear reduction in engine mass results, meaning lower power requirements from the engine. The overall efficiency of the engine is increased due to lower losses. For example, the smaller cylinders give rise to reduced friction losses. Due to the improvements in fuel economy engine emissions are lower. These improvements come at the price of power output. However this can be made up for by the incorporation of a pressure boosting system such as a turbocharger.

Another means to reduce the fuel consumption and improve emissions is to operate the engine in a lean combustion regime ($\lambda > 1$) [8]. This improves fuel economy by allowing the engine to operate at an increased compression ratio and decreased throttling losses. A downside to this regime is the increased instability of combustion due to the higher combustion temperatures favouring knock and the increased chance of misfire due to the smaller amount of fuel. Furthermore, the temperature of combustion increases leading to an increase in NO_x production. Performance of a TWC is degraded when the engine operates away from stoichiometry, resulting in poor NO_x reduction. Lean NO_x traps (LNT) are used to reduce NO_x tailpipe emission for lean combustion regimes appearing in both diesel and gasoline engines [9]. High NO_x reduction of the LNT is ensured by periodically absorbing NO_x in a lean mode, then regenerating the trap during rich running. The control challenge here is to ensure that switching combustion

modes is done effectively to ensure there is little increase in other emissions species during the transient phase, and that the power output is not degraded.

For all ICEs, there are a number of control challenges, especially during transient operation. The interaction of control paths on any number of outputs is difficult to determine independently of the other input/output relationships. Control of gasoline combustion, or fuel injection quantity, is achieved via feedback from the tailpipe [10]. A delayed measurement of the exhaust gas quality is taken at some point in the exhaust piping, and the fuelling adjusted to ensure combustion occurs at the correct AFR.

Torque control is achieved by controlling the quantity of air that is aspirated by the engine [10]. Output torque is estimated based on the intake manifold air, which is controlled by the throttle. The mass of fuel for combustion is then adjusted by the λ control loop. Any remaining actuators are controlled in an open-loop manner. Feed-forward look-up tables are populated based on steady state engine test data captured on an engine dynamometer. Emissions and/or fuel consumption is minimised for each operating point, and the required control inputs noted. This calibration process is an important step in the control of ICEs.

As the dynamics of the airpath side of the engine determine the overall system response it is more important to design and implement closed loop controllers for this system. Model based predictive controllers are increasingly common for airpath control; for example [11] and [12]. The ability to incorporate constraints and utilise multiple control inputs for multiple outputs is a key feature. MPC has also been utilised in an architecture for closed loop AFR control for a lean gasoline engine [13].

As the engine control interactions are complex, there is scope for dynamic optimal control to be applied this control problem. This enables complex tradeoffs to be examined for a variety of scenarios, saving engine testing time. The optimisation of control maps and the generation of optimal transient trajectories are two means by which this can help; for example [14] or [15].

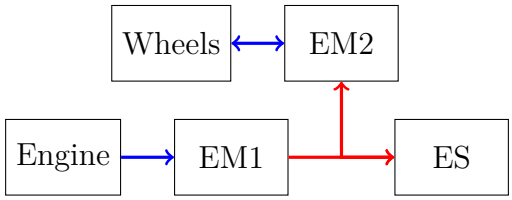
1.1.2 Hybridisation

Hybrid electric vehicles comprising an ICE and electric machine (EM) have existed commercially for over twenty years. The idea behind this technology is to recover some of the energy that would otherwise be lost as heat during braking, and then re-deploy this instead of burning extra fuel. Another facet to this architecture is the ability to operate the engine in a less transient manner and at more efficient points in its operating space. Typically a hybrid electric powertrain consists an ICE, one or more EMs and an energy store (ES) which is usually a battery.

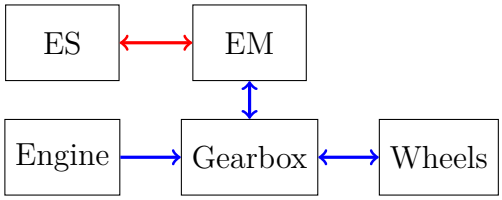
Many differing configurations can be realised with this framework [16]. The connection of each power generating device to the road wheels can offer differing strategies and constraints in terms of how devices operate. Examples of typical hybrid powertrain layouts are shown in Figure 1.3.

Series hybrids (Figure 1.3a) employ an internal combustion engine affixed to a generator unit. This is connected to the ES. This generator pair converts the chemical potential energy of the fuel into electrical potential energy which is then deployed for propulsion through the wheels using another EM. This second EM also provides regenerative braking, converting some of the kinetic energy of the vehicle back into electrical potential. One advantage of this configuration is the ability to ensure the engine operates at close to the optimal operating point. The downside is that power created by the ICE is reduced before it reaches the wheels by the operating efficiencies of each of the electrical machines. Chemical energy in the fuel is converted to mechanical energy in the ICE, which becomes electrical energy in the first EM. The energy is the converted back from electrical to mechanical energy in the second EM.

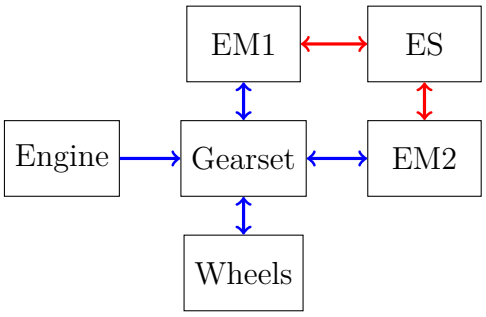
Parallel hybrids (Figure 1.3b) use a motor/generator unit that is coupled directly to the engine drive such that the tractive power output is the sum of the ICE and EM. Over the series configuration, the ability to power the wheels directly from the ICE reduces the losses. This layout can be achieved by placing the EM in a number of different places. In any case, the rotational speed of both of the power sources are fixed to each other during propulsion. Depending on the



(a) An example series hybrid configuration.



(b) An example parallel hybrid configuration.



(c) An example series-parallel hybrid configuration.

Figure 1.3: Possible hybrid powertrain configurations. Electrical power connections are red and mechanical power links are blue. Arrow directions indicate the allowable power flow directions.

location and size of EM, the ICE can be shut-off and clutched, allowing the vehicle to operate in all-electric mode. Additionally, a gearbox is required in order to efficiently propel the vehicle. A drawback of this method compared to a series layout is that inability to operate the ICE at its best operating point. Incorporating a continuously variable transmission (CVT) between the power sources and the wheels can overcome this at the cost of mechanical efficiency.

Parallel-series hybrids (Figure 1.3c) utilise the ICE and two EMs to form an electronic CVT by combining them in a planetary gearset [17]. Utilising this scheme, the ICE can be rotated independently of the wheels in a similar manner to a series hybrid vehicle. When power requirements are high, or the ES is low, the ICE can still be used directly for propulsion in combination with one of the electrical machines. The ability to rotate the engine independently of the wheels means the engine can operate along its fuel-optimal operation line.

Regardless of the mechanical layout all hybrid vehicles require the solution of

a power-splitting problem. To fulfil the tractive power demands, the power from each machine must be decided. This depends on the current ES state of charge and future power requirements. One of the main difficulties of solving this problem is the logic for switching the engine off and on again resulting in a problem with discrete and continuous control inputs. A number of different approaches are taken to solve the powersplit problem which are briefly listed here.

The equivalent minimization consumption strategy (ECMS) utilises optimal control theory to derive analytical solutions which can be implemented in real time [18]. The model formulation minimises the total fuel required from the ICE, and incorporates a single dynamic state relating to the state of charge of the ES. This method for power-splitting control is popular due to its simplicity. The ECMS method has seen a number of variations depending on differing scenarios; for example with emissions considerations [19].

Model predictive control is a suitable means to solve the power-split problem and has been extensively developed for this application [20, 21]. The model used still remains simple, as in the ECMS case. Emissions are difficult to include, so any studies focussed on this area are best computed offline. These studies can be useful in finding what is possible and solving problems that are too complex for online algorithms. Another useful facet of offline studies is the ability to investigate design trade-offs in hardware selection. For example [22] used dynamic programming to optimise not only the powersplit but also the maximum motor power which relates to the motor size. There are also parallels to be drawn between the ECMS strategy for minimising fuel consumption of a hybrid powertrain and an equivalent emissions minimisation strategy for a diesel powered engine [23].

1.1.3 Vehicle Testing

In order to homologate road vehicles, standardised test cycles have been enforced to ensure conformation to legislative emissions limits. Two typical examples are the New European Drive Cycle (NEDC) and the World Harmonised Light Test cycle Procedure (WLTP). These are pre-defined velocity profiles which a vehicle is

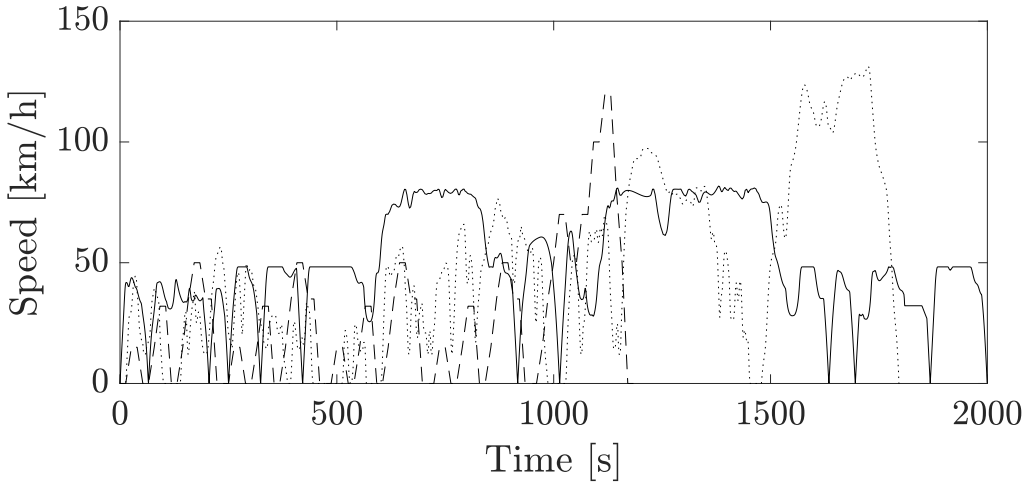


Figure 1.4: An example velocity profile generated in Chapter 4 (solid) compared to the NEDC (dashed) and the WLTC (dotted).

subjected to on a rolling test bed. Total fuel economy and tailpipe emissions are measured and the vehicle certified. One of the major issues with drive cycle testing is the poor correlation with real world data. This came to a head in September 2015 when it was found new road cars were producing emissions levels on average, five times higher than allowed by legislation [24]. In a bid to overcome these testing shortcomings, new legislation has been introduced which measures the real driving emissions (RDE) production during a road test. This is still coupled with standardised test cycles completed on the vehicle test bed [25].

This drive cycle information can be used for validation of control algorithms and parameter sizing studies in the simulation space. Since drive cycles are different, any designs created on one will not be optimal on another. The ability to generate additional driving information for simulation and design is key to unlocking further potential in powertrain design [26]. A comparison of the NEDC and WLTP drive cycles with one of the results from Chapter 4 are plotted in Figure 1.4. This shows the difference between the unrealistic fixed drive cycles and a more real driving scenario which can be created without the need to drive the route.

1.2 Motivation and Structure

A number of control challenges are present in automotive powertrains which relate to fuel economy and emissions management. Additionally, changes to testing procedures have afforded a change in the means to generate engine calibrations and design controllers for use in ECUs. Hardware advances have brought a number of key improvements to road vehicle powertrain structures. Recent developments in computational power and methods for solving complex optimal control problems have enabled the solution of large and complex optimal control problems.

In the standard engine calibration-testing cycle, simulation and optimisation procedures provide an effective method to reduce expensive test bed time and offer a benchmark for implementation comparison. One key area of interest surrounds the creation and evaluation of engine and vehicle performance on real driving emissions (RDE) tests.

This thesis will examine how dynamic optimal control can be used to **evaluate and determine** effective methods for emissions management. Optimal control is applied to joint control of driving and powertrain control on a prescribed route, and transient engine trajectory optimisation for emissions management during combustion regime switching. A variety of smooth model approaches are utilised to model engine fuel efficiency, emissions and in-cylinder processes.

This chapter offered a brief outline of the basis for all of the work pursued in this thesis. The remainder of the thesis is structured as follows. Initially, it is required that the background of optimal control and optimisation is examined. Chapter 2 will discuss some of the relevant background theory for this thesis. The basic necessary conditions for optimality in the presence of constraints are discussed along with algorithms for solving nonlinear programs. Continuous time optimal control is also presented which, when transcribed using a numerical integration scheme, can be reformulated as a nonlinear program. Other schemes for solving optimal control problems numerically are also discussed. The method of choice for this thesis, Legendre-Gauss-Radau collocation, is presented in detail. As well as a detailed consideration of the choice of numerical algorithm, some general

aspects relating to the calculation of derivative, problem scaling and handling model formulations with non-smooth functions are discussed.

Building on the theory laid down in Chapter 2, the application of optimal control to driving problems is initialised in Chapter 3. A very simple vehicle model, without resistive forces, is presented and the solution of a straightforward scenario is computed. Interesting analytical results are presented which indicate the significance of speed limits and arrival times on the solution structure. Extensions to this model are discussed which expose fundamental pieces of the solution to the later optimal driving problem. In subsequent chapters, the optimal control problem is reformulated into the elapsed distance domain. A comparison of the time and distance domain optimal control results for the simple example are examined in Appendix A.

A main result from this thesis is presented in Chapter 4. The model used in Chapter 3 is extended to the scenario of modelling real driving. A multiple phase optimal control problem is posed. Within each phase, a different legislative speed limit is enforced. Route data is obtained from GPS and transformed into curvature and elevation profiles in Appendix B. This method for evaluating real world driving represents a novel contribution to the literature.

High-fidelity emissions models are incorporated into the problem and predict the total emissions produced along the route. Vehicle speed and tractive power are optimised along with engine speed and torque. This effectively optimises both the driver and engine control simultaneously. The sensitivity of the arising solution profiles to constraints on the total emissions is considered. This represents an important step in the engine calibration cycle, especially when legislative emissions limits are considered. A further comparison is made between a continuously variable transmission and a discrete gearbox. Finally, the effect of stopping on the arising drive cycles is considered. Validation work for some of the arising drive cycle profiles is discussed in Appendix C. The results of the optimisation correlate well to the logged data. A series of improvements to subsequent optimisations are proposed based on the information gathered. This chapter demonstrates a method

to generate drive cycles, and engine calibration curves simultaneously for any given route. An example of which is plotted in Figure 1.4.

An extension to the problem of the previous chapter is proposed in Chapter 5. The thermal dynamic effects of heating the catalytic converter have a major impact on the total emissions produced. A parallel-series gasoline hybrid is utilised in order to improve the flexibility of power delivery to the road. A novel powertrain model is formulated to allow for solution of the engine switching problem within a continuous time framework. This contribution is novel and represents an important step in enabling the study of hybrid powertrains in the absence of discrete decision variables. Polynomial maps for electrical machine efficiency and emissions flow rates are included into the problem together with temperature emissions reduction characteristics for the TWC.

Promise is shown for the powertrain model used in this study; the resulting ICE operation is expected. The ICE is utilised when power demands are high and to ensure state-of-charge constraints on the ES can be met.

The sensitivity of the resulting speed profiles and powersplit is considered for a series of emissions constraints. All of the results of this chapter show that significant reductions in emissions can be made for a minimal increase in the total fuel consumption.

In the final emissions management scenario, optimal control is applied to the combustion mode switching of a lean gasoline engine. Chapter 6 poses an optimal control problem which seeks to minimise tracking error in output torque and tailpipe gas quality. This completes the engine calibration cycle offline, as the drive cycle data can be created using the methods of Chapter 4. Typical dynamic engine control only makes use of a few actuators with the majority determined by static feedforward maps. The optimiser here can utilise all of the state and output information to determine the best dynamic control to apply in combination with a feedforward controller. Additional path constraints related to combustion stability are enforced. Results of the optimisation show that vast improvements in output tracking can be achieved by combined control of input actuators. Pre-emptive

control of slow dynamics is a key aspect to effective switch over between lean and rich combustion and vice versa.

Finally, conclusions are drawn and future avenues of work are discussed in Chapter 7.

1.3 Related Work

Work from this thesis has contributed to a number of publications in the field as listed below.

1. David JN Limebeer, and Matteo Massaro. *Dynamics and Optimal Control of Road Vehicles*. Oxford University Press, 2018. Material from Chapters 3 and 4 of this thesis appear in Chapter 9 of the book.

2. Gareth Pease, David Limebeer, and Peter Fussey. "Diesel Engine Optimisation, with Emissions Constraints, on a Prescribed Driving Route". In: *2017 IEEE Africon* . Cape Town, South Africa, 2017, pp. 410-415.

3. Gareth Pease, David Limebeer, and Peter Fussey. "Fuel Consumption Minimisation for Diesel Powered Vehicles on a Prescribed Driving Route". In: *IEEE Transactions on Control Systems Technology* 19.3 (2019), pp. 545-555.

4. Gareth Pease, David Limebeer, and Peter Fussey. "Optimal Hybrid Thermal Warmup Control" *In preparation.*

The following presentations were also given

1. Gareth Pease, David Limebeer, and Peter Fussey. "Application of Optimal Control to Lambda Mode Switching". Symposium on Combustion Control, Aachen Germany, June, 2018.

. . . nothing whatsoever takes place in the universe in which some relation of maximum and minimum does not appear.

— Leonhard Euler, *De Curvis Elasticis*, 1744

2 Numerical Optimal Control

Contents

2.1 Optimisation 16
2.1.1 Unconstrained Optimality 17
2.1.2 Iterative Methods 18
2.1.3 Constrained Optima 20
2.1.4 Numerical Methods for Nonlinear Progams 23
2.2 Optimal Control 26
2.2.1 Optimality Conditions 27
2.2.2 Constrained Minimum Principle 29
2.3 Numerical Methods for Optimal Control 32
2.3.1 Shooting Methods 33
2.3.2 Numerical Integration 35
2.3.3 Indirect Methods 37
2.3.4 Direct Methods 38
2.3.5 Numerical Considerations 45
2.4 Conclusion 52

Control is the field of engineering primarily based on the treatment, design and analysis of controlled systems. External inputs provide a means to influence and actuate complex dynamic systems in a number of different ways. Optimal control is the part of control theory related to deciding the best means to drive a system subject to a number of differing criteria. Examples of optimal control problems in the automotive field might be: the fastest means to complete a

racing lap, the trajectory one must follow in order to minimise fuel consumption over a pre-determined route or choosing the correct trade-off between parameter values such as vehicle power-to-weight. This thesis is concerned with computing solutions to minimum-fuel, emission-constrained manoeuvres and scenarios. In order to compute viable solutions to such problems, it is necessary to review some of the fundamental theory of optimal control. This chapter is by no means a comprehensive examination of the theory, but highlights important facets which are relevant to the problems studied here.

The literature on the subject of optimisation and optimal control is vast: there are a number of good texts on optimisation [27, 28]. Texts on optimal control are also well known [29–31].

First, this will require a brief review of the basics of numerical optimisation. This theory is well known and is repeated here for reference. An exposition of the necessary conditions for optimality of a nonlinear program, along with prevalent solution methods. Following this overview, a similar discussion will take place for dynamic optimal control. Again the necessary conditions for optimality are discussed for a nonlinear performance index constrained by a nonlinear dynamics and constraints. Solution of an optimal control problem numerically requires either a direct or indirect method. In order to compute an optimal trajectory, a suitable numerical integration scheme is required. Quadrature integration rules are used extensively and are discussed. The method used later in this thesis is Legendre-Gauss-Radau global collocation which is discussed in more detail. Finally, there are a number of practical considerations regarding the problem structure, differentiation of the system equations and problem scaling. These are addressed at the end of this chapter.

2.1 Optimisation

Mathematical optimisation deals with the minimisation or maximisation of a function over a finite set of numerical parameters. The function to be minimised is referred to as the *cost (or objective) function* and the input parameters over which the cost

function is minimised are known as *decision variables*. An extremum of a function is a point which is higher or lower than those immediately surrounding it. In the remainder of this chapter, only the minimisation of functions will be considered. Maximisation can be thought of as negative minimisation.

2.1.1 Unconstrained Optimality

A group of n decision variables can be compiled into a vector of length n such that

$$\mathbf{x} = [x_1, x_2, x_3, \ldots, x_n]^\top. \tag{2.1}$$

The set of vectors $\mathbf{x}$ reside in the vector space $\mathbb{R}^n$. Consider the cost function

$$J(\mathbf{x}) : \mathbb{R}^n \to \mathbb{R} \tag{2.2}$$

which maps the decision variable onto the domain of real scalar values. It is intuitive to say that for a minimum $\mathbf{x}^*$ to have been found, all function evaluations in the immediate vicinity must be larger than at the minimum. If this was not the case, the minimum could be found elsewhere and the condition checked at this new point. Formally this can be written as

$$J(\mathbf{x}^*) \leq J(\mathbf{x}). \tag{2.3}$$

Another way to think of this is consideration of the slope of the function at the point $\mathbf{x} = \mathbf{x}^*$. If it were the case that the slope was negative in some direction, then one could move in this direction and decrease the cost. Therefore it is a *first order necessary* condition that $\mathbf{x}^*$ is a minimum if the gradient of the cost function is zero in all directions. The second derivative of $J(\mathbf{x}^*)$ will determine whether the point found is a minimum, maximum or saddle point. If the second derivative is strictly positive definite, then it is sufficient to say that the point $\mathbf{x}^*$ is a minimum. The gradient of the cost function at the point $\mathbf{x}$ is

$$\mathbf{g}(\mathbf{x}) = \nabla J(\mathbf{x}) = \left[\frac{\partial J}{\partial x_1}, \ldots, \frac{\partial J}{\partial x_n}\right]^\top, \tag{2.4}$$

and $\mathbf{H}(\mathbf{x})$ is the Hessian matrix of the cost function at the point $\mathbf{x}$. It is defined as

$$\mathbf{H}(\mathbf{x}) = \nabla^2 J(\mathbf{x}) = \begin{bmatrix} \frac{\partial^2 J}{\partial x_1^2} & \frac{\partial^2 J}{\partial x_1 \partial x_2} & \cdots & \frac{\partial^2 J}{\partial x_1 \partial x_n} \\ \frac{\partial^2 J}{\partial x_2 \partial x_1} & \frac{\partial^2 J}{\partial x_2^2} & \cdots & \vdots \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial^2 J}{\partial x_n \partial x_1} & \frac{\partial^2 J}{\partial x_n \partial x_2} & \cdots & \frac{\partial^2 J}{\partial x_n^2} \end{bmatrix}. \tag{2.5}$$

For optimality the following conditions must hold [32]. *Necessary* conditions for optimality are:

$$\mathbf{g}(\mathbf{x}^*) = 0 \quad \text{and} \quad \mathbf{H}(\mathbf{x}^*) \succeq 0, \tag{2.6}$$

while *sufficient* conditions are that

$$\mathbf{g}(\mathbf{x}^*) = 0 \quad \text{and} \quad \mathbf{H}(\mathbf{x}^*) \succ 0. \tag{2.7}$$

These conditions define a local minimum. If the condition (2.3) extends over the whole domain, then $\mathbf{x}^*$ is a global minimiser. It is important in this analysis that the minimised function is at least twice differentiable, and is smooth over the whole domain in question.

2.1.2 Iterative Methods

So far the necessary and sufficient conditions for optimality have been discussed. As the minimum is not known beforehand, it is necessary to take a guess at where it is, and iteratively update this until the minimum is reached. An intuitive means to do this would consider the gradient and curvature of the function in question at the current point $\mathbf{x}$. The next step would be a certain distance (called *step length*) in a computed direction and chosen such that the cost function decreases. This exercise could be repeated until the changes in the cost function are sufficiently small. *Newton's iterative method* is a well known second order method to achieve this [32, Chapter 1.6]. Consider the quadratic Taylor expansion of the function about a point $\mathbf{x}$

$$J(\mathbf{x}') = J(\mathbf{x}) + \mathbf{g}^\top(\mathbf{x})(\mathbf{x}' - \mathbf{x}) + \frac{1}{2}(\mathbf{x}' - \mathbf{x})^\top \mathbf{H}(\mathbf{x})(\mathbf{x}' - \mathbf{x}) \tag{2.8}$$

The new point $\mathbf{x}'$ can be chosen such that the gradient is zero (i.e. $J(\mathbf{x}') = J(\mathbf{x})$). This then simplifies to

$$\mathbf{g}(\mathbf{x}) + \mathbf{H}(\mathbf{x})(\mathbf{x}' - \mathbf{x}) = 0 \tag{2.9}$$

which is also a first order expansion of the gradient at the new point and setting to zero. This can be rearranged to give the iterative update:

$$\mathbf{x}' = \mathbf{x} - \mathbf{H}^{-1}(\mathbf{x})\mathbf{g}(\mathbf{x}) \tag{2.10}$$

As there is no certainty that the direction chosen is toward a minimum, a descent condition is required:

$$\mathbf{g}(\mathbf{x})^\top(\mathbf{x}' - \mathbf{x}) < 0 \tag{2.11}$$

Utilising (2.9), this results in

$$-\mathbf{g}(\mathbf{x})^\top\mathbf{H}^{-1}(\mathbf{x})\mathbf{g}(\mathbf{x}) < 0. \tag{2.12}$$

This condition only holds in the case that $\mathbf{H} \succ 0$. As the iterate uses second order information, it converges quadratically. However, the update rule (2.10) requires knowledge of both the gradient and curvature of the cost function in question. Construction of the Hessian matrix is, in general, a costly operation and should ideally be avoided to decrease the computational burden.

One means to achieve lower computing time would be to select the descent direction based solely on the gradient of the cost function at the current point. This first order method is referred to as *steepest descent* which has the following update rule:

$$\mathbf{x}' = \mathbf{x} - \alpha_s\mathbf{g}(\mathbf{x}). \tag{2.13}$$

The step-length α_s is either fixed for the whole routine, or itself updated with a line search method, which chooses α_s such that the cost function is minimised along the search direction $\mathbf{g}(\mathbf{x})$. While steepest descent has a lower computational effort compared to that of the Newton update, it converges more slowly. *Quasi-Newton*

methods [33], are a means to compute and update an approximation to the Hessian and/or Jacobian matrices. This incorporates the convergence properties of the full Newton update with the computational speed of the gradient descent. An example update rule is known as the *BFGS update*. A discussion of this, along with other considerations for implementation can be found in [32, Chapter 1.7].

2.1.3 Constrained Optima

So far the optimisation of functions in the absence of constraints has been considered. Many applicable real life situations where optimisation is applied requires the implementation of constraints. Within the engineering field, physical quantities are limited by actuator travels or available input powers. These are inequality constraints, where the decision variable has to remain within a set of bounds. Equality constraints are required to be fulfilled at every feasible point in the optimisation. The first order necessary conditions for a general non-linear program are given by the Karush-Kuhn-Tucker conditions [32, 33].

Equality Constraints

An equality constrained nonlinear program (NLP) is formulated by computing the minimum of $J(\mathbf{x})$ subject to n_f-vector valued equality constraints

$$\mathbf{f}(\mathbf{x}) = \mathbf{0}, \tag{2.14}$$

$\mathbf{f} : \mathbb{R}^n \to \mathbb{R}^{n_f}$. It is required that $J(\mathbf{x})$ and $\mathbf{f}(\mathbf{x})$ are smooth, and twice differentiable in the domain $\mathbb{R}^n$ and that $n_f \leq n$. The *Lagrangian* function is formed by appending the constraints to the cost function in the following manner

$$L(\mathbf{x}, \boldsymbol{\lambda}) = J(\mathbf{x}) + \boldsymbol{\lambda}^\top \mathbf{f}(\mathbf{x}) \tag{2.15}$$

$\boldsymbol{\lambda}$ is an n_f-vector of Lagrange multipliers [32, Chapter 1.8]. In a manner analogous to the unconstrained case, a minimiser of the constrained problem is a stationary

point of the Lagrangian. A multivariable version of the necessary gradient condition (2.6) can be applied which leads to the following,

$$\nabla_x L(\mathbf{x}^*, \boldsymbol{\lambda}^*) = \mathbf{0}, \tag{2.16}$$

$$\nabla_{\boldsymbol{\lambda}} L(\mathbf{x}^*, \boldsymbol{\lambda}^*) = \mathbf{0}, \tag{2.17}$$

for the optimum $\mathbf{x}^*, \boldsymbol{\lambda}^*$ pair. Explicitly, this leads to the necessary conditions for optimality

$$\mathbf{g}(\mathbf{x}^*) + \mathbf{F}(\mathbf{x}^*)^\top \boldsymbol{\lambda}^* = \mathbf{0} \tag{2.18}$$

$$\mathbf{f}(\mathbf{x}^*) = \mathbf{0}. \tag{2.19}$$

where $\mathbf{F}(\mathbf{x})$ is the gradient of the constraint function (2.14). Similarly to the unconstrained case, further information is required for determination of the nature of the stationary point. The second derivative of the Lagrangian leads to its Hessian matrix,

$$\mathbf{H}_L(\mathbf{x}, \boldsymbol{\lambda}) = \nabla^2_{\mathbf{xx}} L(\mathbf{x}, \boldsymbol{\lambda}) = \nabla^2_{\mathbf{xx}} J(\mathbf{x}) + \boldsymbol{\lambda}^\top \nabla^2_{\mathbf{xx}} \mathbf{f}(\mathbf{x}) \tag{2.20}$$

For the point $\mathbf{x}^*$ to be minimum, this is required to be positive definite

$$\mathbf{v}^\top \mathbf{H}_L(\mathbf{x}^*, \boldsymbol{\lambda}^*) \mathbf{v} > 0 \tag{2.21}$$

for any vector $\mathbf{v}$ in tangent space of the constraint set [32, Chapter 1.8].

Inequality Constraints

The final case to be considered is when the function is to be minimised in the presence of both equality and inequality constraints. The new optimisation problem requires the minimisation of the cost function $J(\mathbf{x})$, subject to n_f vector valued equality constraints

$$\mathbf{f}(\mathbf{x}) = \mathbf{0} \tag{2.22}$$

and n_h vector valued inequality constraints

$$\mathbf{h}(\mathbf{x}) \leq \mathbf{0}. \tag{2.23}$$

Where $\mathbf{h} : \mathbb{R}^n \to \mathbb{R}^{n_h}$. Contrarily to the case of equality constraints, the number of inequality constraints is not necessarily required to be less than or equal to the number of variables. This is due to the fact that an infinite number of valid inequality constraints could overlap at a feasible point which lies in the interior of the set.

In the case that the point $\mathbf{x}^*$ is the solution optimisation problem, then the set of constraints can be divided based on whether it is *active* or *inactive*. Active inequality constraints behave as equality constraints at optimum,

$$h_i(\mathbf{x}^*) = 0 \quad \text{for} \quad i \in \mathcal{A} \tag{2.24}$$

where $\mathcal{A}$ is the *active set* of constraints. The remaining $n_h - i$ constraints must be inactive. In the case of optimisation, if one can deduce which constraints are active at the optimum, then the inactive constraints can be removed, and the problem is reduced to an equality constrained one. If one appends the inequality constraints onto the cost function with the Lagrange multipliers $\boldsymbol{\mu}$, then there exists the condition that

$$\mu_i^* > 0 \quad \text{for} \quad i \in \mathcal{A} \tag{2.25}$$

else the multiplier is 0. This can be written concisely as

$$\mathbf{h}(\mathbf{x}^*) \leq 0 \tag{2.26}$$

$$\boldsymbol{\mu}^* \geq 0 \tag{2.27}$$

This is referred to as the *complementary slackness condition*. It is important to note that the point $\mathbf{x}^*$ is a regular point for the constraints, i.e one that causes all active constraints to have linearly independent gradients. For a general NLP of the form (2.2), (2.22) and (2.23), the complete KKT conditions read as follows [32]

$$\nabla_{\mathbf{x}} J(\mathbf{x}^*, \mathbf{u}^*, \boldsymbol{\lambda}^*, \boldsymbol{\mu}^*) = \mathbf{0} \tag{2.28}$$

$$\mathbf{f}(\mathbf{x}^*) = \mathbf{0} \tag{2.29}$$

$$\mathbf{h}(\mathbf{x}^*) \leq \mathbf{0} \tag{2.30}$$

$$\boldsymbol{\mu}^* \geq \mathbf{0} \tag{2.31}$$

$$(\boldsymbol{\mu}^*)^\top \mathbf{h}(\mathbf{x}^*) = \mathbf{0} \tag{2.32}$$

2.1.4 Numerical Methods for Nonlinear Progams

Optimality conditions for a general nonlinear program (NLP) have been presented in the previous section. In a similar manner to the unconstrained case in Section 2.1.2, iterative methods to compute the solution for constrained optimisation problems are required. Two leading algorithms will be considered, both of which can be used for sparse NLPs as the later work in this thesis requires it.

Sequential Quadratic Programming

Sequential quadratic programming (SQP) iteratively solves the KKT conditions the equality constrained problem (2.18) and (2.19). In order to deal with inequality constraints, the active set is determined at each iteration. SQP algorithms rely on the solution of an auxiliary problem or rule set (based on the Lagrange multipliers) to do so. A popular implementation of the active set SQP method is available in the SNOPT package [34]. This is particularly efficient at solving large scale optimisation problems.

A brief description of this method is discussed below. Beginning with the KKT conditions for the equality constrained NLP[1] a first order approximation to these conditions, about the point $\mathbf{x}, \boldsymbol{\lambda}$ gives

$$\mathbf{g}(\mathbf{x}) + \mathbf{F}(\mathbf{x})^\top \boldsymbol{\lambda} + \mathbf{H}_L(\mathbf{x})(\mathbf{x}' - \mathbf{x}) + \mathbf{F}(\mathbf{x})^\top(\boldsymbol{\lambda}' - \boldsymbol{\lambda}) = \mathbf{0} \tag{2.33}$$

$$\mathbf{f}(\mathbf{x}) + \mathbf{F}(\mathbf{x})^\top(\mathbf{x}' - \mathbf{x}) = \mathbf{0} \tag{2.34}$$

Simplifying, one can obtain the KKT system [27]

$$\begin{bmatrix} \mathbf{H}_L(\mathbf{x}) & \mathbf{F}(\mathbf{x}) \\ \mathbf{F}(\mathbf{x})^\top & \mathbf{0} \end{bmatrix} \begin{bmatrix} -\mathbf{p} \\ \boldsymbol{\lambda}' \end{bmatrix} = - \begin{bmatrix} \mathbf{g}(\mathbf{x}) \\ \mathbf{f}(\mathbf{x}) \end{bmatrix} \tag{2.35}$$

In this notation, $\mathbf{p} = \mathbf{x}' - \mathbf{x}$ is the search direction. This linear matrix problem can be solved iteratively to give the solution to the KKT conditions which would be the same as the solution to the original problem. This Newton iterate is based on a linear approximation these optimality conditions, this approximation is equivalent to computing a second order approximation to the original problem functions. The

[1] This now includes those active inequality constraints.

same linear set of equations could be derived by solving the quadratic program minimising over the vector $\mathbf{p}$, the quadratic cost

$$\frac{1}{2}\mathbf{p}^\top \mathbf{H}\mathbf{p} + \mathbf{g}(\mathbf{x})^\top \mathbf{p} \tag{2.36}$$

subject to the equality constraint

$$\mathbf{F}^\top(\mathbf{x})\mathbf{p} + \mathbf{f}(\mathbf{x}) = \mathbf{0}. \tag{2.37}$$

The resulting KKT conditions from this new quadratic program are identical to those in (2.35). As this method is iterative, the inequality constraints can be checked at each step as the program proceeds towards the optimum.

Interior Point Methods

Interior point methods avoid the explicit calculation of the active set of constraints that are required in the SQP case. A set of the KKT conditions are solved for with an approximation made for the set of inequality constraints using an *indicator function* [27]. The function is required to have the properties which match that for the slackness condition (2.26) and (2.27). These can be written

$$B(\mathbf{x}) = \infty,\ \mathbf{x} > 0 \tag{2.38}$$

$$B(\mathbf{x}) = 0,\ \mathbf{x} \leq 0 \tag{2.39}$$

An example of this could be

$$B(\mathbf{x}) = \rho \ln(-\mathbf{x}), \tag{2.40}$$

which in the limit as the positive barrier parameter $\rho \to 0$ meets the criteria. Using this function, the inequality constraints of the NLP can be appended to the cost to give the following *barrier problem*.

$$\min \quad J(\mathbf{x}) - \rho \sum_{i=1}^{n_h} \ln(-h_i(\mathbf{x})) \tag{2.41}$$

$$\text{subject to} \quad \mathbf{f}(\mathbf{x}) = 0 \tag{2.42}$$

The Lagrangian is defined as

$$L(\mathbf{x}, \boldsymbol{\lambda}, \rho) = J(\mathbf{x}) - \rho \sum_{i=1}^{n_h} \ln(-h_i(\mathbf{x})) + \boldsymbol{\lambda}^\top \mathbf{f}(\mathbf{x}) \tag{2.43}$$

Examination of the new KKT conditions for this new optimisation problem lead to the following,

$$\nabla_{\mathbf{x}} L(\mathbf{x}^*, \boldsymbol{\lambda}^*, \rho) = 0 \tag{2.44}$$

$$\mathbf{f}(\mathbf{x}^*) = 0 \tag{2.45}$$

where the gradient of the Lagrangian is

$$\nabla_{\mathbf{x}} J(\mathbf{x}^*) - \rho \sum_{i=1}^{n_h} \nabla_{\mathbf{x}} h_i(\mathbf{x}^*) h_i(\mathbf{x}^*)^{-1} + \mathbf{F}(\mathbf{x}^*)^\top \boldsymbol{\lambda}^* = 0. \tag{2.46}$$

Comparing this with the original inequality constrained KKT conditions, the multipliers can be approximated by

$$\mu_i^* = \rho h_i(\mathbf{x}^*)^{-1}, \quad \text{for } i = [1, \ldots, n_h] \tag{2.47}$$

As $h_i(x)$ tends to zero, the objective will be very large, regardless of the value of ρ. The inequality constraint will therefore sit on the inside of the feasible region. It is possible that solving the new barrier problem KKT conditions using a Newton type algorithm results in ill-conditioning of the problem for small ρ. It is better practice to solve the full set of perturbed conditions of the original problem and solve for the multipliers μ_i in tandem. This means solving the set of equations:

$$\nabla_{\mathbf{x}} J(\mathbf{x}^*) + \nabla_{\mathbf{x}} \mathbf{h}(\mathbf{x}^*)^\top \boldsymbol{\mu}^* + \mathbf{F}(\mathbf{x}^*)^\top \boldsymbol{\lambda}^* = 0, \tag{2.48}$$

$$\mathbf{f}(\mathbf{x}^*) = \mathbf{0}, \tag{2.49}$$

$$(\boldsymbol{\mu}^*)^\top \mathbf{h}(\mathbf{x}^*) + \rho = \mathbf{0} \tag{2.50}$$

IPOPT is one of the most widely used optimisation solvers using this algorithm [35]. IPOPT uses a filter method in combination with a line search technique to ensure that new steps only improve the objective function or constraint violation and ensure that appropriate step lengths are utilised. This thesis will make use of interior point methods to solve the sparse nonlinear programs which arise in the discretisation of continuous time optimal control problems.

2.2 Optimal Control

So far optimisation has been concerned with parameter problems. These consider the optimisation of a function over a finite set of decision variables. Optimal control, by contrast, is the field devoted to the determination of optimal trajectories in time. A large number of similarities exist between optimal control and optimisation. In the optimal control setting, the decision variable is now a curve in time that takes values at every instance in time. This effectively leads to an infinite dimension parameter optimisation problem.

Important results for optimal control rely on the theory of the *calculus of variations* which, like parameter optimisation, is a mature and extensive subject. Good coverage can be found in [29–31, 36]. The analysis here follows that in [36]. A controlled non-linear system is described by an ordinary differential equation as such

$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t), \mathbf{u}(t), t), \quad \mathbf{x}(0) = \mathbf{x}_0, \quad t > 0. \tag{2.51}$$

The dot notation in this case refers to the derivative with respect to time. Within this formulation the system *state*, $\mathbf{x}(t)$ is a function mapping $\mathbf{x} : [0, \infty) \to \mathbb{R}^{n_x}$, while the system control $\mathbf{u}(t)$ is a curve mapping $\mathbf{u} : [0, \infty) \to \mathbb{R}^{n_u}$. The state has initial condition $\mathbf{x}_0 \in \mathbb{R}^{n_x}$. The state dynamics are described by the function $\mathbf{f} : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \times \mathbb{R} \to \mathbb{R}^{n_x}$.

To determine what is 'best' in some sense, a cost function is required. In a general form, one could be defined as such

$$J(\mathbf{x}, \mathbf{u}, t) = m(\mathbf{x}(T), T) + \int_0^T l(\mathbf{x}(t), \mathbf{u}(t), t)dt \tag{2.52}$$

This cost[2] comprises two quantities: the terminal cost, $m : \mathbb{R}^{n_x} \times \mathbb{R} \to \mathbb{R}$, (referred to as the Mayer term) and the running cost, $l : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \times \mathbb{R} \to \mathbb{R}$ (the Lagrange term). The cost function could be chosen to reflect a physical quantity such as input energy. The final time, T must be greater than 0 but may be prescribed or free.

[2]This is referred to as the Bolza form.

The goal of the optimal control is to find a $\mathbf{u}^*(t)$, which has corresponding state $\mathbf{x}^*(t)$ satisfying (2.51), such that the following is true:

$$J(\mathbf{x}^*, \mathbf{u}^*, t) \leq J(\mathbf{x}, \mathbf{u}, t). \tag{2.53}$$

Initially, the control is assumed unbounded.

2.2.1 Optimality Conditions

In the case that the system is to be controlled subject to only the dynamic equations (2.51) without constraints on the control, conditions for optimality can be derived using a variational calculus framework. This has strong parallels with Hamiltonian mechanics [36].

The dynamics (2.51) can be appended to the cost by the use of a *co-state* vector $\boldsymbol{\lambda}(t) : [0, \infty) \to \mathbb{R}^{n_x}$. The initial state $\mathbf{x}_0$ is fixed along with the final time T. The new, appended cost is written as

$$\begin{aligned} J'(\mathbf{x}, \mathbf{u}, \boldsymbol{\lambda}, t) = m(\mathbf{x}(T), T) \\ + \int_0^T l(\mathbf{x}(t), \mathbf{u}(t), t) + \boldsymbol{\lambda}(t)^\top (\mathbf{f}(\mathbf{x}(t), \mathbf{u}(t), t) - \dot{\mathbf{x}}(t)) dt. \end{aligned} \tag{2.54}$$

Time dependency is now omitted for clarity. The integral can be reduced by integrating by parts in the following way

$$J'(\mathbf{x}, \mathbf{u}, \boldsymbol{\lambda}) = m(\mathbf{x}(T)) + \int_0^T l(\mathbf{x}, \mathbf{u}) + \boldsymbol{\lambda}^\top (\mathbf{f}(\mathbf{x}, \mathbf{u}) - \dot{\mathbf{x}})\, dt \tag{2.55}$$

$$= m(\mathbf{x}(T)) + \int_0^T l(\mathbf{x}, \mathbf{u}) + \boldsymbol{\lambda}^\top \mathbf{f}(\mathbf{x}, \mathbf{u})\, dt + \int_0^T \dot{\boldsymbol{\lambda}}^\top \mathbf{x}\, dt - \left[\boldsymbol{\lambda}^\top \mathbf{x}\right]_0^T \tag{2.56}$$

$$= m(\mathbf{x}(T)) + \int_0^T \mathcal{H}\, dt + \int_0^T \dot{\boldsymbol{\lambda}}^\top \mathbf{x}\, dt - \left[\boldsymbol{\lambda}^\top \mathbf{x}\right]_0^T \tag{2.57}$$

The *control Hamiltonian* $\mathcal{H}$ is introduced and is defined as follows

$$\mathcal{H}(\mathbf{x}, \mathbf{u}, \boldsymbol{\lambda}) = l(\mathbf{x}, \mathbf{u}) + \boldsymbol{\lambda}^\top \mathbf{f}(\mathbf{x}, \mathbf{u}). \tag{2.58}$$

Consider now taking a variation in the control $\mathbf{u}$ which is denoted $\delta\mathbf{u}$. This in turn causes a variation in the state $\delta\mathbf{x}$ and a variation in J given by

$$\delta J' = \left[\left(\frac{\partial m}{\partial \mathbf{x}} - \boldsymbol{\lambda} \right)^\top \delta\mathbf{x} \right]_{t=T} + \int_0^T \left(\frac{\partial \mathcal{H}}{\partial \mathbf{x}}^\top \delta\mathbf{x} + \frac{\partial \mathcal{H}}{\partial \mathbf{u}}^\top \delta\mathbf{u} + \dot{\boldsymbol{\lambda}}^\top \delta\mathbf{x} \right) dt. \tag{2.59}$$

If the following is defined,

$$\dot{\boldsymbol{\lambda}} = -\frac{\partial \mathcal{H}}{\partial \mathbf{x}}, \tag{2.60}$$

along with

$$\frac{\partial m}{\partial \mathbf{x}}(T) = \boldsymbol{\lambda}(T), \tag{2.61}$$

then (2.59) is reduced to

$$\delta J' = \int_0^T \frac{\partial \mathcal{H}}{\partial \mathbf{u}}^\top \delta \mathbf{u}\, dt \tag{2.62}$$

This results in the first-order necessary condition for a local optimal control $\mathbf{u} = \mathbf{u}^*$:

$$\frac{\partial \mathcal{H}}{\partial \mathbf{u}} = 0. \tag{2.63}$$

for $t \in [0, T]$. The optimal state trajectory $\mathbf{x}^*(t)$ is computed by integrating the system dynamics from the initial condition. The costate equation is integrated from the terminal point using (2.61). The optimal solution is therefore the triple that satisfies these complete first order necessary conditions;

$$\dot{\mathbf{x}}^* = \mathbf{f}(\mathbf{x}^*, \mathbf{u}^*, t), \tag{2.64}$$

$$\dot{\boldsymbol{\lambda}}^* = -\frac{\partial \mathcal{H}}{\partial \mathbf{x}}\Big|_{\mathbf{u}=\mathbf{u}^*, \mathbf{x}=\mathbf{x}^*, \boldsymbol{\lambda}=\boldsymbol{\lambda}^*}, \tag{2.65}$$

$$\mathbf{0} = \frac{\partial \mathcal{H}}{\partial \mathbf{u}}\Big|_{\mathbf{u}=\mathbf{u}^*, \mathbf{x}=\mathbf{x}^*, \boldsymbol{\lambda}=\boldsymbol{\lambda}^*}. \tag{2.66}$$

The nature of the extremal, by analogy with the parameter optimisation case, can be determined by the curvature considerations of the Hamiltonian in the vicinity of the extrema. In the case of minimisation, the Legendre-Clebsch condition must hold:

$$\frac{\partial^2 \mathcal{H}}{\partial \mathbf{u}^2}\Big|_{\mathbf{u}=\mathbf{u}^*, \mathbf{x}=\mathbf{x}^*, \boldsymbol{\lambda}=\boldsymbol{\lambda}^*} \geq 0, \tag{2.67}$$

at optimality. To find the optimal control, computation of (2.66) and (2.67) is required. The dynamic equations (2.64) and (2.65) form a $2n$ dimension boundary value problem. This has mixed boundary conditions which lead to the name *two-point-boundary-value-problem* (TPBVP) as the state is initialised at $t = 0$, and the costates at the end of the horizon at $t = T$.

It is important to note here that if the Hamiltonian is not an explicit function of time, it is constant along the optimal control trajectory. Starting with (2.58) one can differentiate with the respect to time

$$\frac{d\mathcal{H}}{dt} = \frac{\partial \mathcal{H}}{\partial t} + \frac{\partial l}{\partial \mathbf{x}}^{\top} \dot{\mathbf{x}} + \frac{\partial l}{\partial \mathbf{u}}^{\top} \dot{\mathbf{u}} + \boldsymbol{\lambda}^{\top} \frac{\partial \mathbf{f}}{\partial \mathbf{x}} \dot{\mathbf{x}} + \boldsymbol{\lambda}^{\top} \frac{\partial \mathbf{f}}{\partial \mathbf{u}} \dot{\mathbf{u}} + \dot{\boldsymbol{\lambda}}^{\top} \mathbf{f} \tag{2.68}$$

Grouping the Hamiltonian back together and noting the conditions for optimality (2.65) and (2.66)

$$\frac{d\mathcal{H}}{dt} = \frac{\partial \mathcal{H}}{\partial t} + \left(\frac{\partial l}{\partial \mathbf{x}}^{\top} + \boldsymbol{\lambda}^{\top} \frac{\partial \mathbf{f}}{\partial \mathbf{x}} \right) \dot{\mathbf{x}} + \left(\frac{\partial l}{\partial \mathbf{u}}^{\top} + \boldsymbol{\lambda}^{\top} \frac{\partial \mathbf{f}}{\partial \mathbf{u}} \right) \dot{\mathbf{u}} + \dot{\boldsymbol{\lambda}}^{\top} \dot{\mathbf{x}} \tag{2.69}$$

$$= \frac{\partial \mathcal{H}}{\partial t} + \left(\frac{\partial \mathcal{H}}{\partial \mathbf{x}}^{\top} + \dot{\boldsymbol{\lambda}}^{\top} \right) \dot{\mathbf{x}} + \left(\frac{\partial \mathcal{H}}{\partial \mathbf{u}}^{\top} \right) \dot{\mathbf{u}} \tag{2.70}$$

$$= \frac{\partial \mathcal{H}}{\partial t}. \tag{2.71}$$

As was shown in [30], a number of different other free and fixed time conditions can arise, resulting in further conditions for optimality. For example, if the end time was free then the condition

$$\mathcal{H}(\mathbf{x}^*, \mathbf{u}^*, \boldsymbol{\lambda}^*) = -\frac{\partial m}{\partial T}(T) \tag{2.72}$$

would be fulfilled.

2.2.2 Constrained Minimum Principle

The derivation of the optimality conditions in the previous section had no guarantee that the variation $\delta\mathbf{u}$ would remain in the new control set $\mathcal{U}$. If this were the case, then the condition (2.66) can no longer provide the necessary condition for optimality. In practice, a large number of problems require constraints be placed on actuators. Saturation of current, power limits on mechanical devices and limited available thrust are some prime examples. The seminal work of Pontryagin is required to overcome difficulties in this case [37]. This principle is known as

Pontryagin's Minimum Principle (PMP). The Hamiltonian must be minimised by the optimal control over the whole control set

$$\mathcal{H}(\mathbf{x}^*, \mathbf{u}, \boldsymbol{\lambda}^*, t) \geq \mathcal{H}(\mathbf{x}^*, \mathbf{u}^*, \boldsymbol{\lambda}^*, t), \quad \text{for} \quad t \in [0, T] \quad \text{and} \quad \mathbf{u} \in \mathcal{U} \tag{2.73}$$

which can equivalently be written as

$$\mathbf{u}^*(t) = \arg\min_{\mathbf{u} \in \mathcal{U}} \mathcal{H}(\mathbf{x}^*, \mathbf{u}^*, \boldsymbol{\lambda}^*). \tag{2.74}$$

In the case that the minimum to the problem lies within the control set, then the optimal control can be recovered with (2.66). Else if the control is on the boundary of the set $\mathcal{U}$, then PMP is required.

Now a more general formulation of the optimal control problem can be presented along with its conditions for optimality. Functions describing boundary conditions on the state can be combined with general nonlinear inequality constraints and a performance index which requires initial and final time and state.

In sum, the full optimal control problem is to minimise the cost

$$J(\mathbf{x}, \mathbf{u}, t) = m(\mathbf{x}(t_0), \mathbf{x}(t_f), t_0, t_f) + \int_{t_0}^{t_f} l(\mathbf{x}(t), \mathbf{u}(t), t) dt \tag{2.75}$$

subject to the system dynamics

$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t), \mathbf{u}(t), t). \tag{2.76}$$

Where $m : \mathbb{R}^{n_x} \times \mathbb{R}^{n_x} \times \mathbb{R} \times \mathbb{R} \to \mathbb{R}$ is the new Mayer function. The system dynamics and Lagrange term in the cost remains as before in (2.51) and (2.52). The system is now subject to boundary equality constraints

$$\mathbf{b}_e(\mathbf{x}(t_0), \mathbf{x}(t_f), t_0, t_f) = \mathbf{0}, \tag{2.77}$$

where $\mathbf{b}_e : \mathbb{R}^{n_x} \times \mathbb{R}^{n_x} \times \mathbb{R} \times \mathbb{R} \to \mathbb{R}^{n_{b_e}}$, in combination with boundary inequality constraints

$$\mathbf{b}_i(\mathbf{x}(t_0), \mathbf{x}(t_f), t_0, t_f) \leq \mathbf{0}, \tag{2.78}$$

where $\mathbf{b}_i : \mathbb{R}^{n_x} \times \mathbb{R}^{n_x} \times \mathbb{R} \times \mathbb{R} \to \mathbb{R}^{n_{b_i}}$. Finally a set of nonlinear path constraints are imposed on the system

$$\mathbf{c}(\mathbf{x}(t), \mathbf{u}(t), t) \leq \mathbf{0}, \tag{2.79}$$

where $\mathbf{c} : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \times \mathbb{R} \to \mathbb{R}^{n_c}$.

The control Hamiltonian can be formed in the same fashion as previously, with new multipliers introduced to append the additional constraints. Time dependence is again omitted for clarity.

$$\mathcal{H}(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}, \mathbf{u}) = l(\mathbf{x}, \mathbf{u}) + \boldsymbol{\lambda}^\top \mathbf{f}(\mathbf{x}, \mathbf{u}) + \boldsymbol{\mu}^\top \mathbf{c}(\mathbf{x}, \mathbf{u}) \tag{2.80}$$

The first order necessary conditions read as

$$\dot{\mathbf{x}} = \frac{\partial \mathcal{H}}{\partial \boldsymbol{\lambda}} \tag{2.81}$$

$$\dot{\boldsymbol{\lambda}} = -\frac{\partial \mathcal{H}}{\partial \mathbf{x}} \tag{2.82}$$

$$\mathbf{u}^* = \arg\min_{\mathbf{u} \in \mathcal{U}} \mathcal{H}(\mathbf{x}^*, \mathbf{u}^*, \boldsymbol{\lambda}^*, \boldsymbol{\mu}^*) \tag{2.83}$$

Either or both of the following boundary conditions are enforced in the case that the state at the boundaries is fixed.

$$\boldsymbol{\lambda}(t_0) = -\frac{\partial m}{\partial \mathbf{x}}(t_0) + \left(\frac{\partial \mathbf{b}_e}{\partial \mathbf{x}}(t_0)\right)^\top \boldsymbol{\nu}_e \tag{2.84}$$

$$\boldsymbol{\lambda}(t_f) = +\frac{\partial m}{\partial \mathbf{x}}(t_f) - \left(\frac{\partial \mathbf{b}_e}{\partial \mathbf{x}}(t_f)\right)^\top \boldsymbol{\nu}_e \tag{2.85}$$

The multiplier for the equality boundary constraints is $\boldsymbol{\nu}_e \in \mathbb{R}^{n_{b_e}}$, In the case that either the initial and final times are fixed then the following also have to hold:

$$\mathcal{H}(t_0) = +\frac{\partial m}{\partial t_0} - \boldsymbol{\nu}_e^\top \frac{\partial \mathbf{b}_e}{\partial t_0} \tag{2.86}$$

$$\mathcal{H}(t_f) = -\frac{\partial m}{\partial t_f} + \boldsymbol{\nu}_e^\top \frac{\partial \mathbf{b}_e}{\partial t_f} \tag{2.87}$$

These results are directly analogous to the constrained optimisation of Section 2.1.3. The path constraint attracts the time-varying multiplier $\boldsymbol{\mu}(t)$. The path constraint

set may be active or inactive along the optimal trajectory. It can be generalised in the same manner that a complementary slackness condition holds for these constraints.

$$\boldsymbol{\mu}(t)^\top \mathbf{c}(\mathbf{x}, \mathbf{u}, t) = \mathbf{0} \tag{2.88}$$

and

$$\boldsymbol{\mu}(t) \leq \mathbf{0} \quad \text{and} \quad \mathbf{c}(\mathbf{x}, \mathbf{u}, t) \leq 0 \tag{2.89}$$

Inequality constraints on the boundary state and time can be treated in the same manner, again a new multiplier is required, $\boldsymbol{\nu}_i \in \mathbb{R}^{n_{b_i}}$:

$$\boldsymbol{\nu}_i^\top \mathbf{b}_i(\mathbf{x}(t_0), \mathbf{x}(t_f), t_0, t_f) = \mathbf{0} \tag{2.90}$$

and

$$\boldsymbol{\nu}_i \leq 0 \quad \text{and} \quad \mathbf{b}_i(\mathbf{x}(t_0), \mathbf{x}(t_f), t_0, t_f) \leq \mathbf{0}. \tag{2.91}$$

The necessary conditions for an optimum for a general nonlinear system subject to a number of initial and terminal constraints have been presented. Clearly for a highly complex system, the optimal solution is difficult to immediately jump to. Therefore numerical, iterative methods are required to compute the optimal trajectories.

2.3 Numerical Methods for Optimal Control

For simple problems of relatively low dimension and complexity, PMP can be directly applied to the optimal control problem, and a solution determined which can be written in closed form. The corresponding system dynamics can be integrated by hand to determine the optimal state and costate trajectories. However, in the case of more accurate dynamic system models and realistic scenarios, the integration of complex dynamics may not lead to neat algebraic expressions. Increasing the number of dynamic states and control inputs also does not lend itself to solution this way. To overcome this issue, one of two competing methods for the computation of numerical optimal control solutions must be employed [32].

Indirect methods indirectly compute the solution by solving the two point boundary value problem from the first order necessary conditions (2.81) to (2.91).

This requires explicit computation of the costate equations and their boundary conditions which may be problematic for complex problems.

Direct methods directly solve the optimal control problem by discretising it and solving the arising nonlinear program. The dynamics are discretised to form a set of nonlinear constraints, which, together with the path and boundary constraints, form the set of constraints for a large nonlinear program. This avoids the requirement for explicit derivation of the co-state equations. It is often the case that co-states are computed during the optimisation routine.

A sensible means to integrate the system dynamics is required for either of these strategies, which are touched upon here. Time-marching simulation requires the state at the next time step to be approximated based on the state at one or more previous time steps. The dynamics are integrated along the whole system horizon and the boundary conditions at the end checked. Iteratively better approximations can be made until the constraints are satisfied. The solution to the optimal control problem can be achieved using either a direct or indirect method.

Numerical integration schemes allow for the entire problem horizon to be transcribed on a defined mesh and the entire set of dynamics is solved for simultaneously, respecting any constraints. Again, either an indirect or direct method can be used for the computation of the optimal control.

One could also solve the two point boundary value problem with dynamic programming and the use of the Hamilton-Jacobi-Bellman equation. While this computes the solution for a problem with rigor, the major drawback is the curse of dimensionality [29]. Instead this thesis focuses on using numerical methods to approximate the problem on a finite mesh and solve the arising non-linear program.

2.3.1 Shooting Methods

The simplest case of explicit solution of the system dynamics, is one where the state is propagated forward in time (referred to as time-marching). The system state at the next timestep is computed based on the current one. The TPBVP requires boundary conditions at either end of the time interval must be met. Perhaps

the most simple means to integrate the dynamics is *single-shooting* [32]. The whole optimisation horizon is solved based on a guess of either the initial or final boundary conditions. In the context of optimal control it is likely that the final state boundary conditions are unknown, but the initial condition on the state is. This directional shooting is referred to as backward shooting. The method first guesses the final state conditions and propagates these backward through time (or shoots them). The resulting initial condition is then compared to the known initial condition. A root finding method is used to adjust the variable conditions to minimise the error at the known boundary.

This method of integrating along the whole problem horizon can present numerical difficulties. It may be the case that small changes in initial condition result in a large change in the one at the other end of the horizon. This results in a highly nonlinear relationship between boundary conditions even for linear systems.

A simple means to improve this method, is to split the horizon into segments, then apply single shooting within each. This method is referred to as multiple shooting. A guess at the boundary condition is required at the beginning of each integration interval.

Furthermore, at the boundary of each interval, defect constraints are applied to ensure continuity in the state between propagated dynamics and boundary conditions. If the time interval $[0, T]$ is split into K intervals, $\mathcal{T}_k = [t_{k-1}, t_k]$. The shooting method can be applied within each sub-interval $\mathcal{T}$. Defect constraints of the form

$$x(t_k^-) - x(t_k^+) = 0, \quad k = 1, \ldots, K-1. \tag{2.92}$$

are necessary. Despite the increase in program size due to multiple integration intervals, the shortness of each interval means errors in boundary guesses have a decreased effect on the solution. This increases the size of the root finding problem, depending on the number of states required and the number of shooting intervals. For a Newton method, for example, the Jacobian matrix is required for search direction computation. Fortunately, as the dynamics are effectively decoupled between mesh intervals - as only neighbouring intervals interact - this matrix is sparse.

2.3.2 Numerical Integration

For both shooting and multiple shooting methods, it was required that the dynamics were propagated from one end of the interval to the other. In the case that the system can be solved algebraically, then this is readily achieved. Often, however, it is required that numerical integration schemes are used to approximate the problem as a closed form solution may not exist for even simple looking dynamics. Rather than integrate the dynamics within each segment algebraically, the problem is discretised and solved as a large number of algebraic defect constraints.

Splitting the time horizon into intervals, not necessarily of equal size:

$$t \in [t_0, t_f] \rightarrow t_0 < t_1 < t_2 < \cdots < t_I \tag{2.93}$$

$t_I = t_f$ and the union of all of the intervals is equal to the time horizon: $\mathcal{T}_i = [t_{i-1}, t_i]$, $i \in [1, I]$. The system state can be solved in each mesh interval:

$$\mathbf{x}_{i+1} = \mathbf{x}_i + \int_{t_i}^{t_{i+1}} \dot{\mathbf{x}}(t)dt = \mathbf{x}_i + \int_{t_i}^{t_{i+1}} \mathbf{f}(\mathbf{x}, t)dt. \tag{2.94}$$

This interval can by further split into K subintervals,

$$\tau_j = t_i + h_i \rho_j \tag{2.95}$$

for $1 \leq j \leq K$. The integration interval width is $h_i = t_{i+1} - t_i$ and

$$0 \leq \rho_1 \leq \rho_2 \leq \ldots \rho_K \leq 1 \tag{2.96}$$

The dynamics in each interval, and subinterval can be approximated with quadrature formulae

$$\int_{t_i}^{t_{i+1}} \mathbf{f}(\mathbf{x}, t) \approx h_i \sum_{j=1}^{K} \beta_j \mathbf{f}(\mathbf{x}(\tau_j), \tau_j) \tag{2.97}$$

and

$$\int_{t_i}^{\tau_j} \mathbf{f}(\mathbf{x}, t) \approx h_i \sum_{l=1}^{K} \alpha_{jl} \mathbf{f}(\mathbf{x}(\tau_l), \tau_l) \tag{2.98}$$

This results in the K-stage Runge Kutta integration scheme

$$\mathbf{x}_{i+1} = \mathbf{x}_i + h_i \sum_{j=1}^{K} \beta_j \mathbf{f}_{ij} \tag{2.99}$$

where the intermediary values of the dynamics are

$$\mathbf{f}_{ij} = \mathbf{f}\left(\left(\mathbf{x}_i + h_i \sum_{l=1}^{K} \alpha_{jl}\mathbf{f}_{il}\right), (t_i + h_i\rho_j)\right) \tag{2.100}$$

Each of the set $\rho_j, \beta_j, \alpha_{jl}$ are known and tabulated for well known integration schemes, Butcher arrays provide a concise way to represent these parameters [32, Chap. 3.8]. In the case that $\alpha_{jl} = 0$ for $l \geq j$ then the scheme is explicit, otherwise it is implicit. Example, explicit schemes are the Euler method ($K = 1$) and Runge-Kutta-4 (RK4) ($K = 4$). Implicit schemes include trapezoidal ($K = 2$) and Hermite-Simpson ($K = 3$).

Another way to think of Runge-Kutta schemes is as a $(K+1)^{\text{th}}$ order polynomial over each interval $\mathcal{T}_i$ approximating the state, which can then easily be integrated

$$\tilde{\mathbf{x}}(t) = \sum_{k=0}^{K} a_k (t - t_i)^k. \tag{2.101}$$

The parameters a_k are chosen such that the initial value of the polynomial is same as the state at the initial time. Additionally, the derivative of the polynomial is the same as the state dynamics evaluated at each point, t_i. This method is referred to as a *collocation* method. The RK4 method is a collocation scheme and the solution is represented by a piecewise polynomial of 4^{th} order in each interval. This bears some resemblance to the multiple shooting method, but has the added constraint that the gradient of the state is equivalent to the function at that point. A Radau integration scheme places a collocation point at one of the mesh boundaries. Consider the example of the Euler forward scheme. There is a single collocation point (one at which the gradient is approximated). The state dynamics are integrated thus

$$\mathbf{x}_{i+1} = \mathbf{x}_i + h\mathbf{f}(\mathbf{x}_i, t_i) \quad \text{for} \quad i = 0, 1, \ldots, (I-1), \tag{2.102}$$

which can be transformed into a set of defect constraints

$$\zeta_i = \mathbf{x}_{i+1} - \mathbf{x}_i - h\mathbf{f}(\mathbf{x}_i, t_i) = \mathbf{0}, \quad \text{for} \quad i = 0, 1, \ldots, (I-1). \tag{2.103}$$

The transcription scheme could be either implicit or explicit, and non-linear defect constraints would still arise. Further constraints between mesh intervals must also

be respected (i.e $\mathbf{x}(t_i^+) = \mathbf{x}(t_i^-)$) in the same manner as multiple-shooting. This, in combination with the defect constraints is solved for using a root solving method as the problem takes the form $\mathbf{f}(\mathbf{z}) = \mathbf{0}$.

2.3.3 Indirect Methods

Now that numerical methods for integrating system dynamics have been discussed, the first of two competing methods for solving optimal control problems can be discussed.

Indirect methods are used to solve the first order necessary conditions for the optimal control problem. This then indirectly solves for the optimal control for the original problem. The two boundary value problem consisting (2.81) and (2.82) can be solved using any of the methods discussed above. Depending on the constraints applied to the problem different boundary conditions will occur.

Indirect shooting starts by guessing an initial boundary condition for the costates and integrates these along with the state dynamics across the time horizon. The terminal constraints are checked and the initial conditions can be adjusted and the exercise repeated until convergence. The resulting solution is that for the original problem but it has been found indirectly. Ill-conditioning of the dynamics can occur for small errors in the initial conditions [32]. This can be overcome by applying multiple shooting intervals, to reduce the sensitivity of the solution. As each interval introduces new boundary conditions, this increases the size of the problem. Each interval can be integrated in parallel to also reduce the computational effort.

A slightly different method dubbed indirect collocation requires the use of a discretised state and costate, for example by collocating with polynomials. The problem is to find the state, control and costate to solve the strong form of the PMP (2.66). This large-scale algebraic system is of the from $F(\mathbf{z}) = 0$ and can be solved for using root finding methods. The decision variable $\mathbf{z}$ in this case is a vector of the coefficients of the piecewise polynomials for both the state and control.

The problem can be simplified by altering the path constraints, and appending them using a barrier function in the same manner as interior point methods. This

has the benefit of enforcing the solution to lie within the constraint set and ensures that the strong form of the necessary conditions always holds.

An example software code that employs indirect methods is discussed in [38]. This is an optimal control suite that employs indirect methods with both state and control collocations.

2.3.4 Direct Methods

Direct methods utilising explicit solution methods are known as direct shooting or direct multiple shooting. The control is parametrised with piecewise polynomials, the parameters of which are to be determined via optimisation. The dynamics can then be solved for using a time-marching method or by a numerical integration scheme. Whichever integration scheme is used to solve the dynamics of the problem, an approximation of the cost is given in the same way. In the case shooting over a single interval, this method may struggle when the solution changes very quickly or the time interval over which it is applied is long. This drawback can be somewhat alleviated by splitting the interval into multiple shooting intervals.

Over the past two decades, direct collocation methods have become the *de facto* standard for discretising and solving optimal control problems [36]. This is an implicit method with an integration scheme of the form (2.99) and (2.100). Both the state and control are collocated on an integration mesh and the differential equations are discretised and become algebraic defect constraints. This results in an NLP which minimises the discrete cost function subject to these defect constraints. Path, state and control constraints are also enforced at each discretisation point.

Revisiting the Euler transcription earlier (2.102), the problem was split into N step points with width h_k. The time horizon is split into a number of intervals, $t_0 < t_1 < \cdots < t_{N-1}$ and at each of these points collocation is performed. These are the points where the state is approximated along with its gradient. There is an uncollocated point at t_N, but is still included in the decision variable. A general cost function can be integrated on the same mesh described

$$J = \sum_{k=0}^{N-1} h_k \mathcal{L}(\mathbf{x}_k, \mathbf{u}_k, t_k) \tag{2.104}$$

The decision variable consists both the state and control

$$\mathbf{z} = [\mathbf{x}_0, \mathbf{u}_0, \mathbf{x}_1, \mathbf{u}_1, \ldots, \mathbf{u}_{N-1}, \mathbf{x}_N]^\top. \tag{2.105}$$

The discrete dynamics along with any other problem constraints can be transformed into a constraint vector $\mathbf{g}(\mathbf{z})$. The resulting NLP reads

$$\min_{\mathbf{z}} F(\mathbf{z}) \tag{2.106}$$

$$\mathbf{g}(\mathbf{z}) \leq \mathbf{0} \tag{2.107}$$

Tractability of this NLP is ensured by the fact that it is of sparse structure. This arises due to the fact that most of the elements of the Jacobian of the NLP are the gradients of the defect constraints. There is a low amount of coupling between constraints, a perturbation of a single decision variable effects neighbouring constraints diminishingly. In the case of the Euler method in the example, only adjacent constraints are directly effected by perturbing the state at a given point. This decoupling implies that most of the Jacobian is populated by zero terms. This sparsity is key to ensuring the tractability of the direct methods. Both SQP and interior point algorithms can be used for sparse computing [32].

Orthogonal Collocation

Fixed order methods rely on alterations to the number and width of mesh intervals to improve the accuracy of the approximation, and converge to the solution. The order of the approximating polynomial within each interval is fixed beforehand. Pseudospectral or orthogonal collocation methods use polynomials of varying degree within each mesh interval across the optimisation horizon. The order of polynomial can then be increased to reduce error in the estimation which enables the exploitation of the convergence properties of Gaussian quadrature [36]. Collocation is performed at the Gauss-quadrature collocation points, which themselves are roots of orthogonal polynomials (for example Legendre polynomials). Orthogonal basis functions are chosen to approximate the state. Typically Lagrange polynomials are chosen, but others such as Chebyshev are valid.

This thesis will utilise the Legendre-Gauss-Radau method to solve complex optimal control problems. This presentation follows that in [39, 40]. Within the LGR method, each of the intervals has a collocation point at one of the mesh interval boundaries. Other methods can be employed, such as Legendre-Gauss include neither endpoint and the Legendre-Gauss-Lobatto method which includes both endpoints. [41]. The LGR method is useful when the problem is split into multiple meshing intervals, the first point is collocated on each mesh while the last is not. This means that all points are considered once. The only non-collocated point is at the very end of the optimisation horizon.

For a polynomial function $f(\tau)$, an LGR quadrature rule can be written,

$$\int_{-1}^{1} f(\tau)d\tau \approx \sum_{i=1}^{N} w_i f(\tau_i). \tag{2.108}$$

The continuous integral is approximated as the weighted sum of evaluations of the function up to order N. The location of evaluation nodes τ_i are given by the zeros of the Legendre polynomial $P_{N+1}(-\tau) + P_N(-\tau)$ [3]. The required quadrature weights w_i are

$$w_1 = \frac{2}{N^2}, \quad w_i = \frac{1}{(1-\tau_i)[P'_{N-1}(\tau_i)]^2} \quad \text{for} \quad i = 2, \ldots, N. \tag{2.109}$$

Where $P'(\tau)$ is the derivative of $P(\tau)$ with respect to τ. Using LGR collocation, the highest order f that can be exactly integrated is $2N-2$ [39].

The transcription and assembly of the NLP can now be undertaken. A general Lagrange cost function can be written

$$J = \frac{t_f - t_0}{2} \int_{-1}^{1} \mathcal{L}(\mathbf{x}(\tau), \mathbf{u}(\tau), \tau, t_0, t_f)\, d\tau. \tag{2.110}$$

The timescale has been transformed using the affine transformation $t = \frac{(t_f - t_0)}{2}\tau + \frac{(t_0+t_f)}{2}$ for $\tau \in [-1, 1]$. The problem is to minimise the cost functional J by

[3] Legendre polynomials are given by

$$P_n(x) = 2^n \sum_{k=0}^{n} x^k \binom{n}{k} \binom{\frac{n+k-1}{2}}{n}$$

determining the optimal state $\mathbf{x} \in \mathbb{R}^{n_x}$ and optimal control $\mathbf{u} \in \mathbb{R}^{n_u}$, subject to the system dynamics

$$\frac{d\mathbf{x}(\tau)}{d\tau} = \frac{t_f - t_0}{2}\mathbf{f}(\mathbf{x}(\tau), \mathbf{u}(\tau), \tau, t_0, t_f), \tag{2.111}$$

path constraints

$$\underline{\mathbf{c}} \leq \mathbf{c}(\mathbf{x}(\tau), \mathbf{u}(\tau), \tau, t_0, t_f) \leq \overline{\mathbf{c}}, \tag{2.112}$$

and boundary conditions

$$\underline{\mathbf{b}} \leq \mathbf{b}(\mathbf{x}(-1), \mathbf{x}(+1), t_0, t_f) \leq \overline{\mathbf{b}}. \tag{2.113}$$

The problem horizon can be split into K mesh intervals, which results in the partitioning of the time horizon $\mathcal{T}_k = [T_{k-1}, T_k]$ for $k = 1, \ldots, K$. It should be noted that $T_0 = -1$ and $T_K = +1$. The state and control in $\mathcal{T}_k$ can be written $\mathbf{x}^{(k)}(\tau)$ and $\mathbf{u}^{(k)}(\tau)$ respectively. The LGR method [40] will now be detailed for a single integration interval, this can be repeated across to the entire optimisation horizon.

Within a mesh interval, the j^{th} component of the state is approximated by Lagrange polynomials as basis functions in the following way

$$x_j(\tau) \approx \sum_{i=1}^{N+1} x_{ij} L_i(\tau) \tag{2.114}$$

where

$$L_i(\tau) = \prod_{\substack{l \neq j \\ l=1}}^{N+1} \frac{\tau - \tau_l}{\tau_i - \tau_l} \tag{2.115}$$

The time horizon again is $\tau \in [-1, +1]$ for the k^{th} integration interval. The state is approximated based on the LGR collocation points which are the Legendre evaluation nodes.

The whole discrete state $\mathbf{X}(\tau)$ can be differentiated with respect to τ, giving

$$\frac{d\mathbf{X}(\tau)}{d\tau} = \sum_{i=1}^{N+1} \mathbf{X}_i \frac{dL_i(\tau)}{d\tau}. \tag{2.116}$$

The dynamics can be collocated at the LGR points to obtain

$$\sum_{i=1}^{N+1} D_{ij}\mathbf{X}_j - \frac{t_f - t_0}{2}\mathbf{f}(\mathbf{X}_i, \mathbf{U}_i, \tau_i, t_0, t_f) = 0 \quad j = 1, \ldots, N, \tag{2.117}$$

where D_{ij} are the entries of the $N \times (N+1)$ LGR differentiation matrix $\mathbf{D}$. A slightly different, but related formulation can lead more explicitly to the nonlinear constraints

$$\mathbf{X}_{j+1} = \mathbf{X}_1 + \frac{(t_f - t_0)}{2} \sum_{i=1}^{N} I_{ij} \mathbf{f}(\mathbf{X}_i, \mathbf{U}_i, \tau_i) \quad \text{for} \quad j = 1, \ldots, N. \tag{2.118}$$

This alternative *implicit integral form* [40] has an $N \times N$ LGR *integration matrix* $\mathbf{I}$. This matrix is formed by inversion of submatrices from the differentiation matrix

$$\mathbf{I} = \Big[\mathbf{D}_2, \ldots, \mathbf{D}_{N+1}\Big]^{-1}. \tag{2.119}$$

It is also the case that $\mathbf{I}\mathbf{D}_1 = -\mathbf{1}$. Finally, the cost function is integrated with a quadrature approximation. The following NLP has arisen from the discretisation of a single mesh. Minimise the functional

$$J \approx \frac{t_f - t_0}{2} \sum_{i=1}^{N} \omega_i \mathcal{L}(\mathbf{X}_i, \mathbf{U}_i, \tau_i, t_0, t_f) \tag{2.120}$$

subject to the discrete dynamic equations posed as defect constraints

$$\mathbf{X}_{j+1} - \mathbf{X}_1 - \frac{(t_f - t_0)}{2} \sum_{i=1}^{N} I_{ij} \mathbf{f}(\mathbf{X}_i, \mathbf{U}_i, \tau_i) = \mathbf{0} \quad \text{for} \quad j = 1, \ldots, N, \tag{2.121}$$

and the path constraints applied at each meshing point

$$\underline{\mathbf{c}} \leq \mathbf{c}(\mathbf{X}_i, \mathbf{U}_i, \tau_i, t_0, t_f) \leq \overline{\mathbf{c}} \quad j = 1, \ldots, N. \tag{2.122}$$

This formulation is repeated in each mesh interval, with the total cost summed over all integration intervals. The boundary constraints are applied to the boundaries of the first and last mesh intervals:

$$\underline{\mathbf{b}} \leq \mathbf{b}(\mathbf{X}_1^{(1)}, \mathbf{X}_N^{(K)}, t_0, t_f) \leq \overline{\mathbf{b}}. \tag{2.123}$$

Finally, there is an equality constraint that enforces state continuity at each interior boundary interval

$$\mathbf{X}_{N_k+1}^{(k)} = \mathbf{X}_1^{(k+1)} \quad \text{for} \quad k = 1, \ldots, K-1 \tag{2.124}$$

This constraint can be removed when implementing the problem in a computational environment by simply using the same decision variable for both quantities. It

may arise that both inequality and equality integral constraints are imposed on the optimal control problem. This could be the total elapsed time along a journey or the cumulative mass of some quantity. These take the form

$$\underline{\mathbf{q}} \leq \mathbf{q} = \int_{t_0}^{t_f} \mathbf{Q}(\mathbf{x}(\tau), \mathbf{u}(\tau), t_0, t_f) dt \leq \overline{\mathbf{q}}. \tag{2.125}$$

This can be added to the state vector and the constraints applied as state constraints:

$$\begin{bmatrix} \dot{\mathbf{x}} \\ \dot{\mathbf{q}} \end{bmatrix} = \begin{bmatrix} \mathbf{f}(\mathbf{x}, \mathbf{u}, \tau) \\ \mathbf{Q}(\mathbf{x}(\tau), \mathbf{u}(\tau), t_0, t_f) \end{bmatrix} \tag{2.126}$$

Parameters can also be included as decision variables in the optimal control formulation. These are effectively static control inputs with the same value at each point in time.

Several software packages implement pseudospectral methods for solving nonlinear optimal control problems (for example PROPT [42] or DIRCOL [43]). One such, GPOPS-II [39], will be used for the remainder of this thesis.

First Order Necessary Conditions

Collocation methods have proved themselves effective in both direct and indirect methodologies, for example in the case of solving minimum laptime problems [44].

It is not immediately obvious that the first order necessary conditions arising from the discretisation of the program is equivalent to a discretised version of PMP. This can be demonstrated for a simple problem of the form:

$$\min_{\mathbf{u}} J = \int_{t_0}^{t_f} \mathcal{L}(\mathbf{x}, \mathbf{u})\, dt\,. \tag{2.127}$$

subject to

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{u}) \tag{2.128}$$

The initial condition is $\mathbf{x}(t_0) = \mathbf{x}_0$. The problem horizon is split into K mesh intervals between the fixed initial and final times, t_0 and t_f. Euler transcription with interval width $h = (t_f - t_0)/M$ leads to the following discrete dynamics

$$\mathbf{X}_{k+1} = \mathbf{X}_k + h\mathbf{f}(\mathbf{X}_k, \mathbf{U}_k). \tag{2.129}$$

and discrete cost function

$$\tilde{J} \approx \sum_{k=1}^{K-1} h\mathcal{L}(\mathbf{X}_k, \mathbf{U}_k) \tag{2.130}$$

The initial condition $\mathbf{X}_1 = \mathbf{x}_0$ is also required. The Lagrangian for the NLP reads:

$$\tilde{L} = \sum_{k=1}^{K-1} \left(h\mathcal{L}(\mathbf{X}_k, \mathbf{U}_k) + \mathbf{\Lambda}_k^\top (\mathbf{X}_{k+1} - \mathbf{X}_k - h\mathbf{f}(\mathbf{X}_k, \mathbf{U}_k)) \right) \tag{2.131}$$

The following KKT conditions hold at optimality:

$$\frac{\partial \tilde{L}}{\partial \mathbf{U}_k} = h\left(\frac{\partial \mathcal{L}_k}{\partial \mathbf{U}_k} - \mathbf{\Lambda}_k^\top \frac{\partial \mathbf{f}_k}{\partial \mathbf{U}_k} \right) = 0 \quad k = 0, \ldots, K-1 \tag{2.132}$$

$$\frac{\partial \tilde{L}}{\partial \mathbf{X}_k} = \mathbf{\Lambda}_{k-1} - \mathbf{\Lambda}_k + h\left(\frac{\partial \mathcal{L}_k}{\partial \mathbf{X}_k} - \mathbf{\Lambda}_k^\top \frac{\partial \mathbf{f}_k}{\partial \mathbf{X}_k} \right) = 0 \quad k = 0, \ldots, K-1 \tag{2.133}$$

$$\frac{\partial \tilde{L}}{\partial \mathbf{\Lambda}_k} = \mathbf{X}_{k+1} - \mathbf{X}_k - h\mathbf{f}_k = 0 \quad k = 0, \ldots, K-1 \tag{2.134}$$

$$\frac{\partial \tilde{L}}{\partial \mathbf{X}_K} = \mathbf{\Lambda}_{K-1}^\top = 0 \tag{2.135}$$

Where the following simplifications have been made: $\mathbf{f}_k = \mathbf{f}(\mathbf{X}_k, \mathbf{U}_k)$ and $\mathcal{L}_k = \mathcal{L}(\mathbf{X}_k, \mathbf{U}_k)$.

In the limit as $h \to 0$, the discrete conditions tend to their continuous time counterparts. That is (2.132) becomes the control equation (2.66), (2.134) recovers the state dynamics (2.64), and (2.133) and (2.135) become (2.65) and (2.61) respectively. In the case that a path constraint is added this must be respected at each of the collocation points.

Note that the discrete state equation (2.134) is an explicit Euler scheme while the discrete adjoint equation (2.133) results in an implicit Euler scheme where both schemes together represent a symplectic Euler scheme (see [46]) applied to the state-adjoint system. Similar symplectic schemes are obtained for general Runge-Kutta discretisation as shown in [45, 47]. Furthermore, equivalence of direct and direct methods could be shown if symplectic discretisation schemes are applied directly to the state dynamics [46].

In indirect methods, the program requires the creation of analytical expressions for the costates, along with the optimality and boundary conditions. This can then be discretized. This can prove cumbersome for the user if the state equations

are complicated. This offers one means by which the user can select between the two methods. Within direct methods the method internally produced the costate equations which can later be extracted from the solution.

2.3.5 Numerical Considerations

When formulating and solving numerical optimisation methods there are series of considerations regarding mesh refinement techniques, scaling, singular arcs and the smoothing of functions with discontinuous gradients. In order to ensure tractability, these issues must be overcome, often at the expense of accuracy. This trade-off is fundamental to enabling complex optimal control problems to be solved numerically.

Derivative Information

For effective solution to any optimisation problem, it is necessary to provide first (and second) derivative information. The simplest method to produce the Jacobian and Hessian information is analytically. Analytical derivatives can be cumbersome to compute by hand but have the advantage of machine precision. Difficulty in derivation for a large number of systems states or complex model structure can be overcome by use of a symbolic computation software such as the MATLAB Symbolic Math toolbox. However, even in this state, evaluation of the derivatives can be computationally expensive and therefore hinder solution speed significantly.

A well known means to compute numerical derivatives is by finite differencing. A Taylor expansion of a function can be written:

$$f(x+h) \approx f(x) + h\frac{df}{dx} + h^2\frac{d^2f}{dx^2} + \dots \tag{2.136}$$

As h is small, higher terms than the first order are neglected. Rearranging gives rise to the forward difference relation:

$$\frac{df}{dx} \approx \frac{f(x+h) - f(x)}{h} + \mathcal{O}(h) \tag{2.137}$$

Similarly, the backward difference formula can be derived by considering a step $x - h$. Improvements to both of these relations can be achieved by the central

difference method which provides quadratic error convergence. This arises from the difference in forward and backward Taylor expansions.

$$\frac{df}{dx} \approx \frac{f(x+h) - f(x-h)}{2h} + \mathcal{O}(h^2) \tag{2.138}$$

It is known that numerical instability or impreciseness occurs unless the parameter h is sufficiently small. In this case, numerical issues can arise (the derivative becomes very large). Additionally, the error due to *subtractive cancellation* can start to dominate if the parameter is chosen too small.

Complex step differentiation overcomes this by computing a Taylor expansion about the real axis, into the complex plane [48]. This requires only one function evaluation compared to the two for the finite differencing methods. There is also machine precision as the h parameter becomes smaller.

$$\frac{df}{dx} \approx \mathrm{Im}\left(\frac{f(x+ih)}{h}\right) \tag{2.139}$$

This method is accurate to $\mathcal{O}(h^2)$.

Another method for computing derivatives is automatic or algorithmic differentiation [49]. The majority of suitable model functions for optimal control can be formed using the basic operations: addition, subtraction, multiplication and division, in combination with trigonometric functions, hyperbolic functions, exponentials or logarithms. Utilising the chain rule of differentiation, the derivative information for complex model structures can be determined. Indeed, it is easy to compute the total derivatives for the model easily and quickly, up to machine precision. This method retains the accuracy that analytical derivatives give. In this thesis, the *source transformation* method is used in the ADiGator software package [50]. This routine generates a set of derivative files based on a set of source files provided by the user.

Mesh refinement

Considering the non-linear nature of such optimal control problems, it would be wrong to assume that one could heuristically choose the correct integration mesh for the problem. An important step in the solution of optimal control problems is an

estimate of the accuracy of the current solution. After which it may be required that the mesh over which the problem was solved is adjusted to improve the estimate. In the presence of slow dynamic evolution, or rapid switching of control, it is beneficial to adaptively refine the integration mesh in order to converge efficiently to the solution. For a fixed order method, the approximating polynomial order is decided beforehand. Error reduction can only be achieved by altering the placement of mesh boundaries. By contrast, a pseudospectral method alters the order of the approximating polynomial across the optimisation horizon. Non-smoothness of the solution can pose problems for this method. A polynomial of extremely high order may result and, in the case of a single integration mesh covering the whole horizon, this would result in a loss of sparsity in the Jacobian matrix. In the case of fixed order methods, at points of non-smoothness, an increased number of mesh intervals would be required. Furthermore, increasing the density of the mesh results in a larger NLP which hinders solution speed.

In order to overcome these limitations, a combined *hp*-method can be utilised. These algorithms determine the location, width and order of each mesh interval using local error estimation. An example methodology for this is presented in [51]. An accurate approximation to the solution of the continuous time problem can be achieved efficiently and in a timely manner compared to either solely p or h methods. A variety of mesh refinement algorithms are present in GPOPS-II [39].

Singular Arcs

There is an important class of optimal control problems to consider where the performance index and system equations are linear functions of the one or more of the control inputs. The resulting Hamiltonian can be written

$$\mathcal{H}(\mathbf{x}, \mathbf{u}, \boldsymbol{\lambda}) = g(\mathbf{x}, \boldsymbol{\lambda}) + \Phi(\mathbf{x}, \boldsymbol{\lambda})\mathbf{u}. \tag{2.140}$$

As the control will be bounded, PMP states that the control will switch between boundaries provided the switching function $\Phi()$ remains non-zero. This can be

written compactly as

$$\mathbf{u}^*(t) = \begin{cases} \mathbf{u}_{min}, & \text{if} \quad \Phi(t) > 0 \\ ???, & \text{if} \quad \Phi(t) \equiv 0 \\ \mathbf{u}_{max}, & \text{if} \quad \Phi(t) < 0 \end{cases} \tag{2.141}$$

Cases where the control switches between the upper and lower limits are referred to as *bang-bang* control arcs.

When $\Phi(t) = 0$ for a finite interval, this is referred to as a *singular arc*. PMP provides no information about the solution of the problem. Numerical methods cannot recover the optimal control either, the optimisation routine will often lead to oscillatory control trajectories which prevents efficient computation. Mesh refinement algorithms can continue to increase the number of mesh intervals and increase the order or polynomials in order to try to recover a solution. The Hessian matrix of the Hamiltonian can be close to, rather than actually singular.

In order to properly deal with the case of a singular arc, a higher order minimum principle is needed [52]. By using higher derivatives of control variations, new necessary conditions can be obtained. A feedback law can be derived using this higher order condition and applied to the problem during the singular interval. Successive derivatives of Φ aid in the derivation of the control structure. An example of this is the Goddard rocket problem [53].

One of the main reasons for using a direct method is the avoidance of the explicit calculation of the optimality conditions and algebraic expressions. Additionally, the conditions for singular arcs arising would need to be determined. In some cases there may be several singular arcs of differing lengths which can themselves be optimised. Several numerical type solutions have been proposed in the literature. To avoid singularity in the Hessian of the Hamiltonian, a small positive definite matrix can be appended [54],

$$\frac{\partial^2 \mathcal{H}}{\partial \mathbf{u}^2} + 2\mathbf{R} \geq 0, \tag{2.142}$$

to make it positive definite, provided $\mathbf{R} > 0$. The corresponding modified cost function is

$$\tilde{J} = \int_{t_0}^{t_f} \left(\mathcal{L}(\mathbf{x}) + \mathbf{u}^\top \mathbf{R} \mathbf{u} \right) dt. \tag{2.143}$$

In order to preserve the original problems structure, it is required that $\mathbf{R}$ remain small. Practically, one would have to start with a large value of $\mathbf{R}$, and iteratively decrease this in size until a reasonable solution trajectory was obtained. A similar method to this, referred to as *piecewise derivative variation* [55], can be used. It was shown that the problem would reach the original problem in the limit as the number of mesh points increased.

A variation of the method (2.143) is employed in this thesis. The system is controlled by the rate of change of the inputs and this is combined with a cost function term penalising the derivative of the control. An augmented state vector is now integrated along the system trajectory,

$$\mathbf{x}_a = [\mathbf{x}, \mathbf{u}]^\top \tag{2.144}$$

and the cost function is modified to read

$$\tilde{J} = \int_{t_0}^{t_f} \mathcal{L}(\mathbf{x}_a) + \dot{\mathbf{u}}^\top \mathbf{R} \dot{\mathbf{u}} dt. \tag{2.145}$$

This results in smooth controls chosen as the solution due to the penalty on gradient of control. One downside to this method is the increased dimensionality of the problem, this can be reduced by appending only those controls in which the problem is or is close to singular. It is important therefore to select the correct number of controls to drive with their derivatives and ensure the total fraction of the cost that they take up is insignificant. This operation has to be completed iteratively with a reasonable threshold for the regularisation component. A set of control derivatives and weights are chosen and the problem solved. This is then compared to other weightings and control derivative inputs.

Problem Scaling

In a similar manner to that for the singular arcs, an ill-conditioned Hessian matrix can lead to difficulty when solving numerically. In the case that the state vector has notionally large and small state quantities, the Hessian matrix can end-up being ill-conditioned. In this case, convergence will be slow. To alleviate this problem,

the system and cost function can be scaled such that the numerics of the problem are more suitable. A well used method makes all dimensioned quantities in the system dynamics and cost dimensionless using a basis of SI units [56]. This ensures that the whole problem is consistently scaled, this is especially useful when the dynamics model is large. Included in this scaling is an alteration of the time horizon itself, which may be formed by choosing a suitable rotational speed scale. If the rotational speed takes values in the range $[0, \omega_m]$ then it is prudent to scale the state by the maximum speed. In the case of a turbocharger, this is likely to be in the tens of thousands of Hz. The variable scaling would be as follows

$$\omega'(t') = \omega(t)/\omega_m \tag{2.146}$$

With the independent variable scaled as follows

$$t' = t\omega_m. \tag{2.147}$$

One means to overcome this, is by altering the state equations and such that the new state is reasonably sized, but does not involve fully non-dimensionalising the problem. A more detailed explanation of this appears in Chapter 6.

Function Smoothing

As was alluded to in the mesh refinement section, non-smoothness can pose problems for the solution of the NLP. A number of useful functions that have undefined gradients, for example the max() function and the absolute function, are required later in this thesis. Both of these functions have discontinuous first derivatives which result in non-smooth derivatives. As functions of this type can be formed using the sign() function a number of *sigmoid* functions can be used to provide continuous derivative information. The hyperbolic tangent is a useful approximation

$$\text{sign}(x) \approx \tanh(kx) \tag{2.148}$$

As $k \to \infty$ then the hyperbolic tangent approaches the sign() function. The absolute function can be written as

$$|x| = \text{sign}(x)x \tag{2.149}$$

Therefore an appropriate approximation might be

$$|x| = x \tanh(kx) \tag{2.150}$$

In [57], an approximation of the maximum function is presented

$$\max^1(x, 0) \approx \frac{x + \sqrt{x^2 + \epsilon}}{2}. \tag{2.151}$$

While another useful approximation takes the form

$$\max^2(x, 0) \approx \frac{x(1 + \tanh(kx))}{2} \tag{2.152}$$

Each of these has its advantages and disadvantages. In the case of the first, the function is always positive. This is beneficial if the optimisation routine requires strictly positive outputs from the function to ensure a quantity remains real for example. However, this approximation intersects the y-axis at $\frac{\sqrt{\epsilon}}{2}$. If the control enters the optimal control problem through this condition and is required to be 0. Then the approximation will result in a small positive quantity driving the system. By contrast, the hyperbolic approximation (2.152) fulfils the criterion that $\max(0, 0) = 0$. However, the function is not always positive. A minimum is found at $kx = -0.639$. Provided k is chosen large enough, the approximation is closer to the max function, and the minimum becomes smaller for small negative x.

Depending on the requirements, either of these functions can be used. The two approximations to the max function are compared in Figure 2.1. In this thesis, (2.152) is used extensively. Finally, a small transformation is required for the comparison of two arguments

$$\max(x, y) = \frac{x}{2}(1 - \tanh(k(x - y)) + y. \tag{2.153}$$

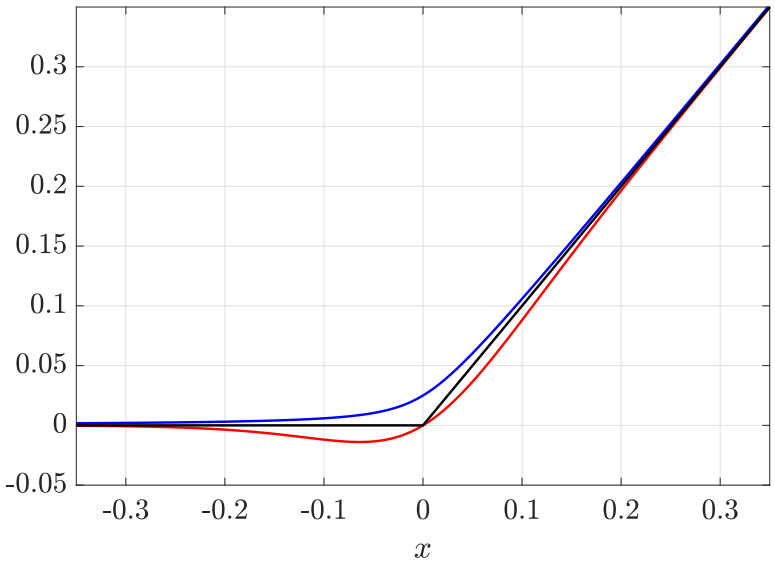


Figure 2.1: A comparison of two approximations of the $\max(x, 0)$ function. The hyperbolic approximation (2.152) is red, the square-root approximation (2.151) is blue and the original function is black.

2.4 Conclusion

This chapter has laid the foundations of the numerical solution of optimal control problems. The conditions for optimality for unconstrained optimisation and a means to solve these problems was discussed initially. This was followed by a brief outline of constrained optimality, including the KKT conditions. The solution to general nonlinear programs was briefly discussed and two leading algorithms for solving general NLPs were presented.

Continuous time optimal control theory was also laid out in this chapter, along with Pontryagin's Minimum Principle. This forms the first order necessary conditions for optimality for a nonlinear dynamic system. Numerical methods for discretisation of optimal control problems were discussed at some length. Direct and indirect methods are summarised, while some detailed attention is given to LGR pseudospectral methods. This discretisation method, along with the interior point algorithm form the basis of the GPOPS-II optimisation suite used extensively in the remaining chapters of this thesis. The equivalence of the first order necessary conditions arising from the discrete problem are compared to the continuous time

counterparts derived from PMP. Algorithmic differentiation, mesh refinement, non-smooth functions and singular arcs are a few important necessary considerations when implementing any of the aforementioned methods.

The next chapter in this thesis applies the knowledge gained here to optimal control problems pertaining to the minimisation of input energy. This is analogous to minimising the fuel consumption - something which is integral to the fuel-optimal driving of later chapters. The initial problem is simple, and can be solved analytically using PMP. Extensions to this problem are introduced which reflect the additional problem information required in optimal driving. These extensions required the computation of the optimal trajectories using the pseudospectral method detailed in this chapter. The validity of this method is outline by comparing a numerically computed case with an analytical solution.

Is it true that the particle doesn't just "take the right path" but that it looks at all the other possible trajectories?

Richard Feynman, The Feynman Lectures on Physics Vol II

3 Minimum Energy Optimal Control

Contents

3.1 Minimum energy optimal control 57
3.1.1 Arrival time constraint 65
3.1.2 Slope Effects 69
3.1.3 Computational Validation 71
3.2 Variable Velocity Limit 74
3.3 Variable Road Grade 78
3.4 Conclusions 84

THIS chapter focuses on the application of optimal control theory to vehicular energy minimisation problems. As the energy input in a road vehicle context refers to fuel mass, there are direct analogies to be drawn between minimum fuel and minimum energy optimal control problems. Typically in minimum fuel strategies three control arcs occur: maximum permitted acceleration to the speed limit is followed by adherence to the limit for the majority of the journey. Finally the vehicle is brought to rest along a braking arc [58]. The situation may arise where there are zero-tractive-force periods of coasting - for example downhill. This avoids increases in the total fuel consumption [36]. Fuel energy is required to overcome losses due to friction and changes in elevation along the route. In order to keep losses due to drag low, the journey should be completed with low average speeds.

It may be required that the route is traversed within an allocated time. This would cause an increase in fuel consumption. There is a fuel economy trade-off therefore between the journey time and the fuel consumption.

The purpose of this chapter is to analyse minimum energy optimal control problems of a similar nature to those that can occur in the optimisation of automotive scenarios for optimal driving and powertrain use. Simple examples similar to those given in this chapter are included in a number of well known texts, for example [29, Chap. 3.9] and [31, Chap. 5.5]. Four related vehicular control problems will be considered, which highlight some of the main features of time-constrained minimum-energy optimal control problems. In the interests of simplicity, aerodynamic forces and power loss effects are neglected. The first problem considers the minimum-fuel ascent of a hill with constant gradient. This is solved analytically in order to uncover the main structural features of the problem, particularly when speed restrictions are in place. Initially, analysis is completed with a fixed arrival time. Further analysis will consider the case when the terminal time is upper bounded rather than fixed. This results in a condition whereby the maximum arrival time is always reached. In the second example, the problem is then increased in complexity to provide insight into the solution when variable speed limits are enforced. This could be present in minimum fuel optimal control problems as a curvature-based acceleration limit. Finally, a more general case for traversal of a route with variable road grade is considered. This example enables insight into regenerative braking and more conventional dissipative braking. The former is an integral part of hybrid powertrain operation.

The goal here is to evaluate the important parts of such minimum fuel optimal control problems. It will be shown that the prescribed or bounded arrival time has a profound effect on the solution structure and optimal cost to complete the courses considered. The more simple examples of this chapter also enable the validation of the numerical methods which are required in later chapters of this thesis.

All of the analytical work in this chapter considers the solution of the presented problems in the time domain. However, it is sometimes advantageous to re-formulate

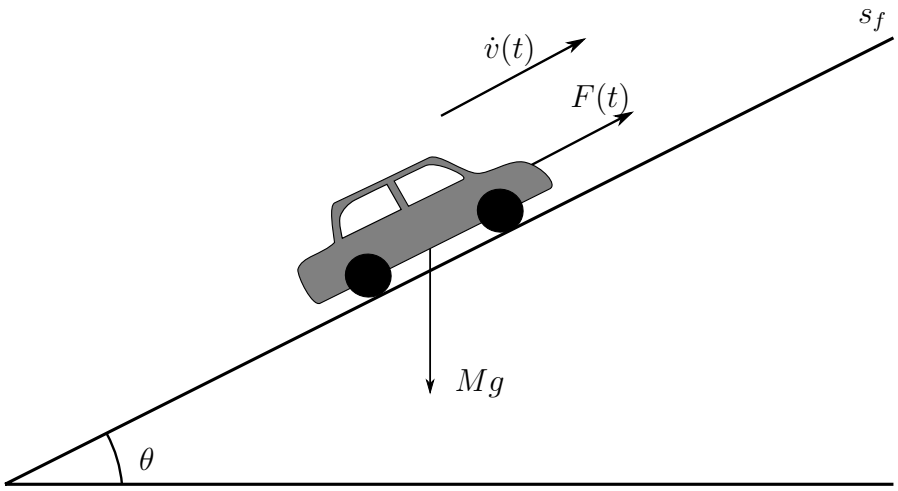


Figure 3.1: A vehicle ascending a slope of constant incline.

the problem into the elapsed distance domain. This property is exploited later in the thesis. Equivalent results are presented in Appendix A for the elapsed distance domain.

Some material from this chapter appears in [36, Chap. 9.2] and [59].

3.1 Minimum energy optimal control

Consider the problem of minimum energy control of the vehicle in Figure 3.1, in the absence of frictional forces. The car is required to cover the distance (s_f) within a prescribed time (t_f), while minimising the total energy input E. This example serves to highlight important features that are present in vehicular minimum-energy optimal control problems. There is a direct analogy to be drawn between minimising input energy and fuel. Clearly, energy is required to increase the cars potential energy up the length of the slope. This expenditure of fuel energy also increases the kinetic energy of the vehicle, thereby reducing the journey time.

A controllable force input $F(t)$ and the component of the car's weight parallel to the slope are the only forces which act on the car. The car's position and velocity are $s(t)$ and $v(t)$ respectively. The car has mass M and the acceleration due to gravity is g. For simplicity in the following analysis, let $K = g \sin\theta$ be the component of acceleration due to gravity parallel to the slope. The optimal

control problem is therefore to minimise

$$\min_{F(t)} E = \int_0^{t_f} F(t)v(t)\mathrm{d}t. \tag{3.1}$$

The equations of motion for the problem are (3.2) and (3.3). The velocity constraint is (3.4) and the driving force constraint is (3.5).

$$\dot{v}(t) = F(t)/M - K, \tag{3.2}$$

$$\dot{s}(t) = v(t), \tag{3.3}$$

$$v(t) \leq v_m(t), \tag{3.4}$$

$$0 \leq F(t) \leq F_m. \tag{3.5}$$

The boundary conditions on the states are:

$$s(0) = 0,\ s(t_f) = s_f, \tag{3.6}$$

$$v(0) = 0,\ v(t_f) = \text{free}. \tag{3.7}$$

This problem will be solved in its full generality. The control Hamiltonian (see Section 2.2) for the problem can be formed by appending the dynamic state equations and velocity inequality constraint to the cost function

$$\mathcal{H}(t) = F(t)v(t) + \lambda_v(t)\Big(F(t)/M - K\Big) + \lambda_s(t)v(t) + \mu(t)\Big(v(t) - v_m(t)\Big), \tag{3.8}$$

$\lambda_v(t)$ and $\lambda_s(t)$ are the costate variables associated with each of the system state variables. The multiplier $\mu(t)$ appends the speed-related inequality constraint to the control Hamiltonian, which must satisfy the following conditions:

$$\mu(t) = 0 \quad \text{for} \quad v(t) < v_m(t), \tag{3.9}$$

$$\mu(t) < 0 \quad \text{for} \quad v(t) = v_m(t), \tag{3.10}$$

$$\mu(t) > 0 \quad \text{for} \quad v(t) > v_m(t). \tag{3.11}$$

The sign of $\mu(t)$ ensures that violating the constraint (3.4) cannot decrease (3.8). Pontryagin's Maximum Principle (PMP) states that an optimal control is one that minimises the Hamiltonian over the problem horizon, i.e.

$$F^*(t) = \arg\min_{F(t)} \quad \mathcal{H}(v^*, s^*, F, \lambda_v^*, \lambda_s^*, \mu^*) \tag{3.12}$$

for $t \in [0, t_f]$. The superscript $(.^*)$ denotes the optimal trajectories for the control, state and co-state variables. The first-order necessary conditions for optimality are

$$\begin{aligned} \frac{d\mathcal{H}(t)}{dv} &= -\dot{\lambda}_v(t), \\ \frac{d\mathcal{H}(t)}{ds} &= -\dot{\lambda}_s(t), \end{aligned} \tag{3.13}$$

that is

$$\dot{\lambda}_v(t) = -(F(t) + \lambda_s(t) + \mu(t)) \tag{3.14}$$

$$\dot{\lambda}_s(t) = 0 \tag{3.15}$$

which implies that $\lambda_s(t)$ is constant. The final velocity is free, and so there is a transversality constraint, $\lambda_v(t_f) = 0$. It follows from PMP that the optimal control is given by

$$F^*(t) = \begin{cases} 0, & \text{if} \quad \Phi(t) > 0 \\ F_s(t), & \text{if} \quad \Phi(t) \equiv 0 \\ F_m, & \text{if} \quad \Phi(t) < 0 \end{cases} \tag{3.16}$$

where $\Phi(t) = v(t) + \lambda_v(t)/M$ is the switching function - consisting the terms preceding the control[1]. In the case of a singular arc, $\Phi(t) \equiv 0$ for some finite time interval, and so $\dot{\Phi} \equiv \ddot{\Phi} \equiv ... \equiv 0$, must also hold on this time interval. In the present case, the first derivative of the switching function is

$$\dot{\Phi}(t) = -(K + (\lambda_s(t) + \mu(t))/M). \tag{3.17}$$

Since θ and λ_s are constant, and $\mu = 0$ in the case of an inactive speed constraint (as is assumed initially), $\dot{\Phi}(t)$ is constant. It can be concluded that the control never appears following differentiation of the switching function (with respect to time) and thus that singular arcs do not exist for this problem. The resulting control is *bang-bang*.

At the start of the optimal control interval $F^*(t) = F_m$, which implies that $\Phi(0) < 0$, because in the case that $F^*(0) = 0$ the car would roll backwards down the incline. Since $\dot{\Phi}$ is constant, there is at most one control switch. In sum,

$$F^*(t) = \begin{cases} F_m, & t \in [0, \tau] \\ 0 & t \in (\tau, t_f] \end{cases} \tag{3.18}$$

[1] $F_s(t)$ is a singular control along the singular arc.

with τ the thus far unknown switching time.

On the first sub-arc, the state equations can be integrated, observing (3.6) and (3.7), to give

$$v_1(t) = \left(\frac{F_m}{M} - K\right)t \tag{3.19}$$

$$s_1(t) = \left(\frac{F_m}{M} - K\right)\frac{t^2}{2}. \tag{3.20}$$

Similar expressions arise for the second sub-arc

$$v_2(t) = K(t_f - t) + v_f \tag{3.21}$$

$$s_2(t) = -\frac{K}{2}(t_f - t)^2 + v_f(t - t_f) + s_f. \tag{3.22}$$

Equating each of the states at the switching time $t = \tau$ gives

$$v_f = \frac{F_m}{M}\tau - Kt_f \tag{3.23}$$

$$0 = F_m\tau^2 - 2M(Kt_f + v_f)\tau + M(Kt_f^2 + 2v_f t_f - 2s_f). \tag{3.24}$$

Eliminating v_f gives

$$0 = F_m\tau^2 - 2F_m t_f\tau + M(2s_f + Kt_f^2). \tag{3.25}$$

As $\tau \leq t_f$, the smaller of the two roots of (3.25) must be taken

$$\tau = t_f - \sqrt{\frac{(F_m - MK)t_f^2 - 2Ms_f}{F_m}}. \tag{3.26}$$

The optimal cost is given by

$$E^* = M\left(Ks_f + \frac{v_f^2}{2}\right), \tag{3.27}$$

which is the car's total energy as it passes the finish line. It consists of a kinetic energy term and a potential energy term. Equation (3.26) contains information about the existence of solutions for the problem. Since $\tau \leq t_f$ must be real

$$(F_m - MK)t_f^2 - 2Ms_f \geq 0, \tag{3.28}$$

which implies that $F_m - MK > 0$. Intuitively, this is correct. The force available for propulsion must be at least as much as required to overcome the vehicle's weight

projected down the incline. The condition, (3.28), also ensures there is enough time for the car to cover the course if driven at the maximum allowable force for the duration. In the case that (3.28) is an equality, $\tau = t_f$ and maximum force is required for the whole length of the course. This means the vehicle arrives in the minimum amount of time

$$t_f^{min} = \sqrt{\frac{2s_f}{F_m/M - K}}. \tag{3.29}$$

Substituting into (3.23) gives the final velocity

$$v_f^{min} = \sqrt{2s_f(F_m/M - K)}, \tag{3.30}$$

which can be substituted into (3.27) to give $(E^*)^{min} = F_m s_f$.

In the case of the original optimal cost (3.27), this can be minimised by realising the condition $v_f = 0$. This results in the switching time

$$\tau = \sqrt{\frac{2MKs_f}{F_m(F_m/M - K)}}, \tag{3.31}$$

and the minimum-energy arrival time

$$t_f^{en} = \sqrt{\frac{2F_m s_f}{K(F_m - MK)}}. \tag{3.32}$$

In this case, the solution to the energy minimisation problem is $(E^*)^{en} = MKs_f$, which is the energy required to raise the vehicle from the initial to the final elevation. These arrival and switching times are only for a bang-bang control scheme. An infinite number of other solutions exist to give the same minimum energy cost. The control profile is ultimately determined by the arrival time. Increasing the arrival time decreases the average velocity, but does not further reduce the cost. This is a result of the lack of inclusion of losses due to drag. This is discussed later when considering the case of a bounded arrival time.

Now consider the case when the velocity state constraint is active. A constant speed limit, $v_m(t) = v_m$, is enforced for the horizon length. When $v(t) = v_m$, $F(t) = MK$ follows from (3.2). Referring to (3.16), an active velocity constraint

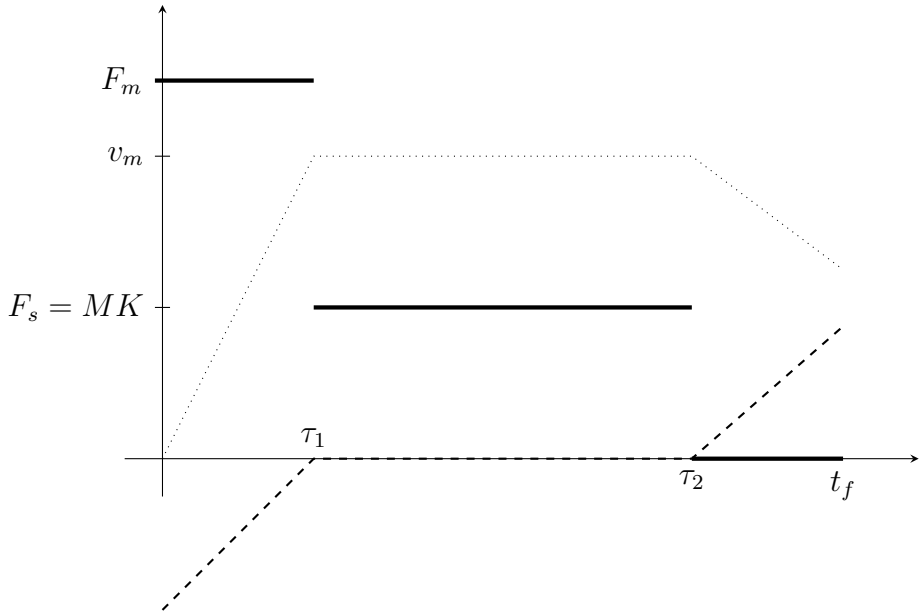


Figure 3.2: Bang-singular-zero control (solid line), the switching function Φ (dashed line) and the velocity (dotted line).

must occur on a singular sub-arc, which implies that for some finite period $\Phi(t) \equiv 0$. The second derivative of the switching function is given by

$$\ddot{\Phi} = -\frac{\dot{\mu}^*}{M} = 0, \tag{3.33}$$

which implies that $\mu^*(t)$ is also constant along the singular sub-arc. The constraint (3.4) is active, which further implies that $\mu^*(t) < 0$ along the arc. It follows, therefore, that the car accelerates up to v_m on the first sub-arc when $\Phi < 0$, the vehicle then progresses at constant speed v_m on a second sub-arc (when $\Phi = 0$), and coasts on a third sub-arc when $\Phi > 0$. The *bang-singular-bang* control strategy is shown in Figure 3.2.

Purely mechanics-based arguments can be used to find the switching times and state trajectories. The states on the first sub-arc are described by

$$v_1(t) = (F_m/M - K)t, \tag{3.34}$$

$$s_1(t) = \frac{(F_m/M - K)}{2}t^2, \tag{3.35}$$

which remains the same as (3.19) and (3.20). For the singular arc the velocity is

constant and the driving force balances the weight projected along the incline. Thus

$$v_2(t) = v_m, \tag{3.36}$$

$$s_2(t) = \frac{(F_m/M - K)}{2}\tau_1^2 + v_m(t - \tau_1), \tag{3.37}$$

where τ_1 is the first switching time. For the third sub-arc the dynamics are given by

$$v_3(t) = K(t_f - t) + v_f \tag{3.38}$$

$$s_3(t) = s_f - \frac{K}{2}(t - t_f)^2 + v_f(t - t_f). \tag{3.39}$$

which remains the same as (3.21) and (3.22). The first switch can be calculated using (3.34) with $v_1 = v_m$ to give

$$\tau_1 = \frac{v_m}{(F_m/M - K)}. \tag{3.40}$$

The car travels the distance

$$\sigma_1 = \frac{(F_m/M - K)}{2}\tau_1^2 \tag{3.41}$$

along this subarc. Equating (3.37) and (3.39) at τ_2 can be used to find the second switching time.

$$\sigma_1 + v_m(\tau_2 - \tau_1) = s_f - \frac{K}{2}(\tau_2 - t_f)^2 + v_f(\tau_2 - t_f) \tag{3.42}$$

It is necessary to eliminate v_f as it depends on the duration spent on the final sub-arc. Equating (3.36) and (3.38), gives

$$v_f = v_m + K(\tau_2 - t_f). \tag{3.43}$$

Therefore

$$\sigma_1 - s_f = -\frac{K}{2}(\tau_2 - t_f)^2 + \Big(v_m + K(\tau_2 - t_f)\Big)(\tau_2 - t_f) - v_m(\tau_2 - \tau_1) \tag{3.44}$$

$$\sigma_1 - s_f = -\frac{K}{2}(\tau_2 - t_f)^2 + v_m(\tau_2 - t_f) + K(\tau_2 - t_f)^2 - v_m(\tau_2 - \tau_1) \tag{3.45}$$

$$(\tau_2 - t_f)^2 = \frac{2}{K}\Big(\sigma_1 - s_f + v_m(t_f - \tau_1)\Big) \tag{3.46}$$

Again, the smaller of the two roots must be taken as $\tau_2 \leq t_f$.

$$\tau_2 = t_f \ - \sqrt{\frac{2}{K}\left((t_f - \tau_1)v_m + \sigma_1 - s_f\right)}. \tag{3.47}$$

If this problem is to be soluble, t_f must also satisfy

$$(t_f - \tau_1)v_m + \sigma_1 - s_f \geq 0, \tag{3.48}$$

which means that the terminal time must be long enough to allow the car to accelerate to v_m under the influence of F_m, and then move at v_m to the finish line at s_f. The final velocity is given by

$$v_f = v_m - K(t_f - \tau_2) \tag{3.49}$$

The optimal energy required to propel the vehicle under this control scheme is

$$E^* = F_m \frac{F_m/M - K}{2}\tau_1^2 + MKv_m(\tau_2 - \tau_1). \tag{3.50}$$

Once again, consider the case where the arrival time is minimised. This is the case whereby, the car is accelerated to the speed limit and subsequently driven at this limit for the remainder of the journey. There is no second switch to zero propulsive force (i.e. $\tau_2 = t_f$). The first switch time is given by (3.40). The minimum journey time is found by considering the case when the quantity under the square root in (3.47) is zero, i.e.

$$(t_f - \tau_1)v_m + \sigma_1 - s_f = 0. \tag{3.51}$$

Substituting for the first switching time (3.40) and distance (3.41) reduces this expression to

$$t_f^{min} = \frac{v_m}{2(F_m/M - K)} + \frac{s_f}{v_m} \tag{3.52}$$

which has corresponding energy expenditure

$$(E^*)^{min} = M(\frac{v_m^2}{2} + Ks_f). \tag{3.53}$$

In the minimum energy case, $v_f = 0$ which implies that $t_f^{en} - \tau_2^{en} = v_m/K$ by (3.49). Combining this result with (3.47) (and the results (3.40) and (3.41)) leads to

$$(t_f^{en} - \tau_2^{en})^2 = \frac{2}{K}\left(v_m t_f^{en} - \frac{v_m^2}{(F_m/M - K)} + \frac{v_m^2}{2(F_m/M - K)} - s_f\right) \tag{3.54}$$

$$\frac{v_m^2}{2K} = v_m t_f^{en} - \frac{v_m^2}{2(F_m/M - K)} - s_f \tag{3.55}$$

$$t_f^{en} = \frac{F_m v_m}{2K(F_m - MK)} + \frac{s_f}{v_m}. \tag{3.56}$$

The corresponding second switching time is

$$\tau_2^{en} = t_f^{en} - v_m/K. \tag{3.57}$$

The energy expenditure for the speed-limited case is $(E^*)^{en} = MKs_f$, which is the same as the unconstrained case. In the case that a boundary condition was enforced such that $v(t_f) = 0$, and the final time was free, infinitely many solutions occur to give the same cost as in the minimum energy case. These results echo those for the unconstrained example.

3.1.1 Arrival time constraint

In the examples examined thus far, the arrival time, t_f, was stipulated. This section examines the effect of upper bounding the arrival time rather than enforcing a predetermined time. It is often the case in driving scenarios, that a journey is completed within a maximum time rather than at a designated time. In order to assess this effect, the Hamiltonian will have to be explicitly computed in order to check terminal boundary conditions. Once again, the case when (3.4) is inactive will be treated first.

Firstly, it is necessary to explicitly derive expressions for the costates (3.15). Substituting for $F^*(t)$ and the known boundary condition $\lambda_v(t_f) = 0$ leads to the following piecewise definition of $\lambda_v^*(t)$,

$$\lambda_v^*(t) = \begin{cases} -(F_m + \lambda_s)t + A, & t \in [0, \tau] \\ (t_f - t)\lambda_s, & t \in (\tau, t_f], \end{cases} \tag{3.58}$$

where A is an unknown constant of integration. Costates are continuous along optimal trajectories, so A can be determined by equating the costate at $t = \tau$,

$$-(F_m + \lambda_s)\tau + A = (t_f - \tau)\lambda_s, \tag{3.59}$$

which results in

$$A = t_f \lambda_s + F_m \tau. \tag{3.60}$$

It was previously determined that $\lambda_s(t)$ remains constant along the whole route. Therefore

$$\lambda_s(t) = \lambda_s. \tag{3.61}$$

A piecewise expression for the Hamiltonian can now be determined. Considering the second control arc, where zero force is applied, and using the expressions (3.21) and (3.22):

$$\mathcal{H}^*(t) = -K(t_f - t)\lambda_s + \lambda_s[K(t_f - t) + v_f]$$

which reduces to

$$\mathcal{H}^*(t) = \lambda_s v_f. \tag{3.62}$$

On the first control arc a similar expression arises:

$$\mathcal{H}^*(t) = (F_m/M - K)[F_m \tau + \lambda_s t_f] \tag{3.63}$$

Importantly neither of these expressions are an explicit function of time. λ_s can be determined by equating the two (after noting 3.23):

$$\lambda_s = -\frac{\tau(F_m - MK)}{t_f - \tau}. \tag{3.64}$$

The optimal Hamiltonian is therefore

$$\mathcal{H}^* = -\frac{\tau}{t_f - \tau}(F_m - MK)v_f. \tag{3.65}$$

As $t_f \to t_f^{min}$, then $\tau \to t_f$ and $\mathcal{H}^* \to -\infty$. When $t_f \geq t_f^{en}$, $v_f = 0$ and $\mathcal{H}^* = 0$. The $\mathcal{H}^*$ is a product of the final velocity, the net force under maximum propulsive effort and the ratio of the time spent on the first sub-arc to the second sub-arc. This results in units of power. The case whereby the final time is upper bounded by the following $t_f \leq T$, can be considered. The upper limit T in this case is enforced to lie in the interval $[t_f^{min}, t_f^{en}]$. This can be written as

$$a(t_f) = (t_f - T) \leq 0. \tag{3.66}$$

A new, constant, Lagrange multiplier (ν) is required which has the following properties:

$$\nu < 0 \quad \text{and} \quad \nu[a(t_f)] = 0 \tag{3.67}$$

Transversality constraints (see Section 2.2.2) at the terminal time imply that

$$\mathcal{H}^*(t_f) = \nu \frac{\partial a(t_f)}{\partial t_f} = \nu.1 \tag{3.68}$$

Therefore, at the optimal solution

$$\mathcal{H}^* = \nu = -\frac{\tau}{t_f - \tau}(F_m - MK)v_f \tag{3.69}$$

The optimal Hamiltonian along the trajectory is always negative, implying $\nu < 0$, and ensuring the terminal constraint is always active. Provided $T \leq t_f^{en}$, this constraint is always met. As the control structure is unknown for terminal times beyond this, it is impossible to say with certainty that an upper bound constraint is met always. However, this is likely to be the case.

Similar analysis can also be carried out for the speed-limited case. The optimal trajectory of the costates $\lambda_s^*(t)$ and $\lambda_v^*(t)$ and the path multiplier $\mu^*(t)$ are required to compute the Hamiltonian for the fixed arrival time case. Integrating the expression for $\lambda_v(t)$, a piecewise definition arises

$$\lambda_v^*(t) = \begin{cases} -(\lambda_s + F_m)t + B, & t \in [0, \tau_1] \\ -(\lambda_s + F_s + \mu)t + C, & t \in (\tau_1, \tau_2], \\ -\lambda_s t + D, & t \in (\tau_2, t_f], \end{cases} \tag{3.70}$$

where B, C, D are three unknown constants of integration. $\mu^*(t)$ is also defined piecewise as

$$\mu^*(t) = \begin{cases} 0, & t \in [0, \tau_1] \\ \mu^*, & t \in (\tau_1, \tau_2], \\ 0, & t \in (\tau_2, t_f]. \end{cases} \tag{3.71}$$

As in the unconstrained case, $\lambda_s(t) = \lambda_s$ is a constant. Noting the terminal boundary condition $\lambda_v^*(t_f) = 0$, and equating at the boundary of each sub-arc, the solution now reads:

$$\lambda_v^*(t) = \begin{cases} -(\lambda_s + F_m)t + (\tau_2 - \tau_1)(MK + \mu) + t_f\lambda_s + F_m\tau_1, & t \in [0, \tau_1] \\ -(\lambda_s + MK + \mu)t + t_f\lambda_s + \tau_2(MK + \mu), & t \in (\tau_1, \tau_2], \\ -\lambda_s(t - t_f), & t \in (\tau_2, t_f], \end{cases} \tag{3.72}$$

Utilising this information, along with the optimal piecewise velocity and force solutions (of Figure 3.2) leads to the optimal Hamiltonian definition:

$$\mathcal{H}^*(t) = \begin{cases} (F_m/M - K)\Big(\lambda_s t_f + (MK + \mu)(\tau_2 - \tau_1) + F_m\tau_1\Big), & t \in [0, \tau_1] \\ v_m\Big(MK + \lambda_s\Big), & t \in (\tau_1, \tau_2], \\ \lambda_s\Big(K(\tau_2 - t_f) + v_m\Big), & t \in (\tau_2, t_f], \end{cases} \tag{3.73}$$

Once again the Hamiltonian is not an explicit function of time. Equating $\mathcal{H}^*$ on the second and third arc leads to

$$MKv_m = \lambda_s^* K(\tau_2 - t_f) \;\rightarrow\; \lambda_s^* = \frac{-Mv_m}{t_f - \tau_2}. \tag{3.74}$$

As before, $t_f > \tau_2$, and M and v_m are positive, implies λ_s^* is always negative. In the time minimising case, $t_f = \tau_2$, and $\lambda_s^* = -\infty$. Similarly, equating the expressions for $\mathcal{H}^*$ on the first and second arcs[2] leads to a definition of μ^*:

$$\mu^* = \frac{\lambda_s(t_f - \tau_1) + MK(\tau_2 - 2\tau_1) + F_m\tau_1}{(\tau_2 - \tau_1)} \tag{3.75}$$

Therefore, in the speed-limited case, the Hamiltonian is simply written as

$$\mathcal{H}^* = -\frac{Mv_m v_f}{t_f - \tau_2} = \lambda_s v_f. \tag{3.76}$$

As in the unlimited case, as $t_f \rightarrow t_f^{min}$ implies that $\mathcal{H}^* \rightarrow -\infty$. Also, after noting (3.57), $t_f \rightarrow t_f^{en}$, implies that $\mathcal{H}^* \rightarrow 0$. A similar physical interpretation can be made for this new expression for $\mathcal{H}^*$. Again the Hamiltonian has units of power and is the product of $\lambda_s v_f$.

The same arguments can be made for the terminal time being upper bounded. Again

$$\mathcal{H}^*(t_f) = \nu, \implies \nu < 0 \tag{3.77}$$

This holds true provided $T \leq t_f^{en}$. In the absence of other terminal state constraints, the maximum arrival time constraint determines the optimal solution. In both cases considered so far, the control structure was bang-bang or bang-singular-bang. If the arrival time is larger than t_f^{en}, the solution can no longer be easily determined. However, the increased time to traverse the course only decreases the average velocity.

[2]Recalling that $\tau_1 = v_m/(F_m/M - K)$.

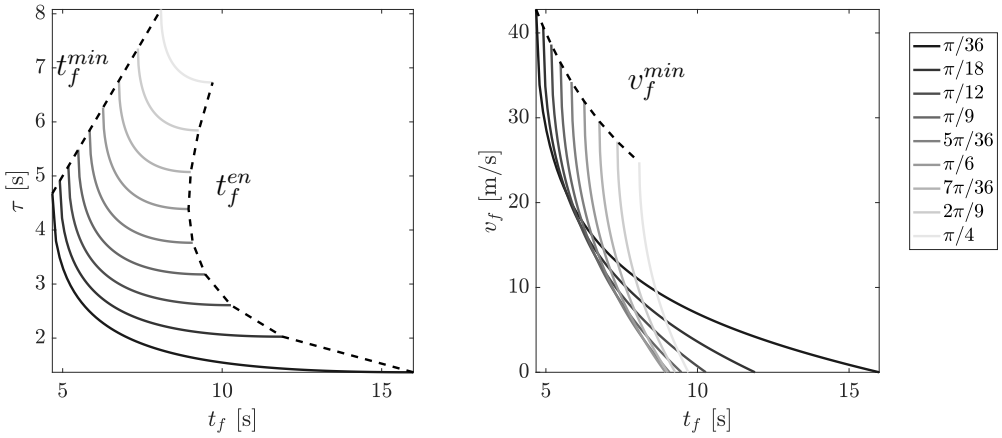


Figure 3.3: Left: This is a figure to show the variation of switching time and arrival time for a sweep of θ. The contours are plots of equation (3.26). Right: figure shows the variation of the final velocity on the same slopes with the minimum traversal time final velocity indicated. The contours are plots of (3.30). These results are for the unconstrained case.

3.1.2 Slope Effects

For both the example in the presence of the speed limit, and the one without, the effects of road grade can be investigated. This is best done by choosing representative numerical values and plotting the effects of the incline on the switching times and final velocity. In this case, the vehicle is required to ascend an incline at a variety of inclinations ($\theta = \frac{k\pi}{36}$ rads, for $k = [1, 2...9]$.). The course length is $s_f = 100$m, the vehicle mass is $M = 1000\,\text{kg}$, the force is limited to 10kN and the velocity limit (when imposed) is 50 km/h, (or $v_m = 13.89\,\text{m/s}$).

In the unconstrained case, the variation of the single switching time, τ for a given θ and t_f is plotted in Figure 3.3. The limiting arrival times, t_f^{min} and t_f^{en} are computed from (3.29) and (3.32). Both are plotted as the dotted lines in LHS of Figure 3.3. The switching time for t_f^{min} in all cases is equal to t_f.

The behaviour of the minimum energy arrival time, t_f^{en}, is a surrogate metric for the average velocity. This is best examined by considering the ways in which the car obtains the zero velocity at low and high values of slope. At low values of slope, the component of weight is low. The car needs to be accelerated for only a short period of time, and decelerates slowly to the end of the course. The average

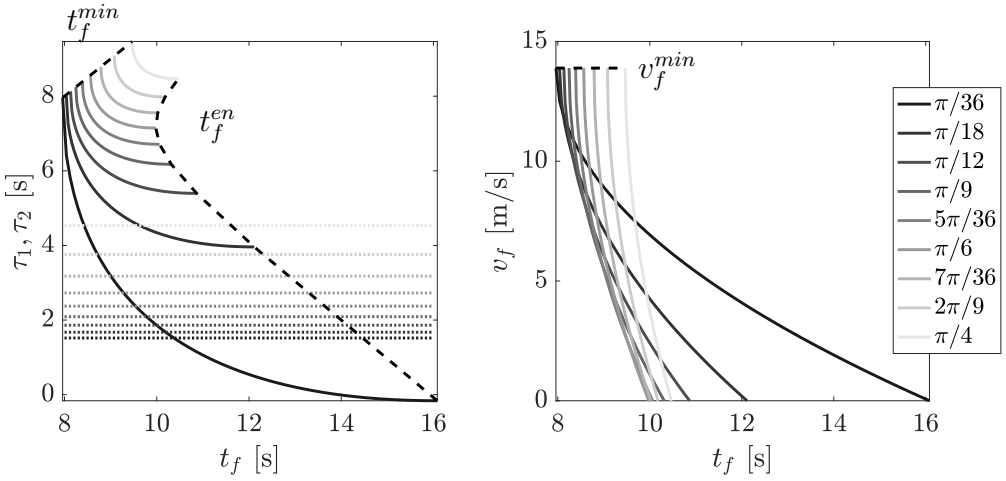


Figure 3.4: Left: Variation of the second switching time τ_2 in the velocity constrained case (solid coloured). The first switching times are dotted lines. Right: Variation of the final velocity v_f for a set of arrival times and road grades.

velocity is low in this case. In the case for very steep slope, the component of weight is harder to overcome. This limits the acceleration available along the first control arc. The subsequent deceleration is rapid, but the vehicle did not reach a high peak speed in this case. This leads again to a low average speed. There is an intermediate inclination for which there is a minimum, t_f^{en} to complete the course. This can be obtained from by minimising the expression (3.32). Doing so leads to the following slope angle

$$\theta^{en} = \sin^{-1}\left(\frac{F_m}{2Mg}\right). \tag{3.78}$$

This angle only depends on the ratio of forces acting on the car.

The range of available switching times also decreases with increasing slope angle (i.e the length of each contour). In the case of a steep slope, the application of the limit force must be longer, decreasing the available time for switching.

For the same set of inclines, the terminal velocity characteristic (3.30) is swept over. The final velocity in the minimum time traversal case is plotted as the dashed line. Obviously, the final velocity in the minimum energy case coincides with the x-axis. The highest minimum time terminal velocity occurs when the course is the flattest, lowering the losses due to the weight component. Interestingly, the spread of available terminal velocities decreases with an increasing slope. Higher

proportions of the journey are required under maximum driving force, removing time for the car to decelerate. The minimum t_f^{en} is highlighted here as the leftmost curve intersecting the x-axis.

For completeness, the same set of inclines is evaluated for the speed limited case. Similar results are obtained. First consider the set of terminal speeds in the right hand side of Figure 3.4. All of the minimum time velocities are at the speed limit. Once again the minimum energy final velocities all intersect the x-axis.

The first switching time, τ_1, is plotted along with τ_2 in Figure 3.4. Again there is an incline which has the lowest t_f^{en}. This is the same as the expression found previously. Interesting behaviour can be noted for shallow slopes; it is possible for the solution to contain no second switch for long arrival times. Instead, the solution does not reach the speed limit, and the control strategy reverts back to a bang-bang control. This behaviour is expected as the resistive effects of the weight are reduced. For example, the lowest grade slope has $\tau_1 = \tau_2$ when $t_f = 10.4$s. At this arrival time, there is a single switch at $\tau_1 = 1.52$s. This situation did not immediately become clear from the analytical expressions derived previously. This analysis is for a single set of vehicle parameters (a specific thrust/weight ratio). It may well be the case that there are further insights to be made with these simple examples by varying these.

3.1.3 Computational Validation

The examples examined so far have been sufficiently simple to enable easy computation by hand. Later examples in this chapter and in this thesis require the numerical optimisation techniques discussed in Chapter 2 to solve the OCP. A brief comparison of the analytical solutions derived for state, costate and control trajectories is discussed here for the velocity limited example. The vehicle parameters are chosen the same as the previous section, along with a slope with $\theta = \pi/12$ and an arrival time of $t_f = 10$s. As a singular arc will occur, the problem will require regularisation (see Section 2.3.5).

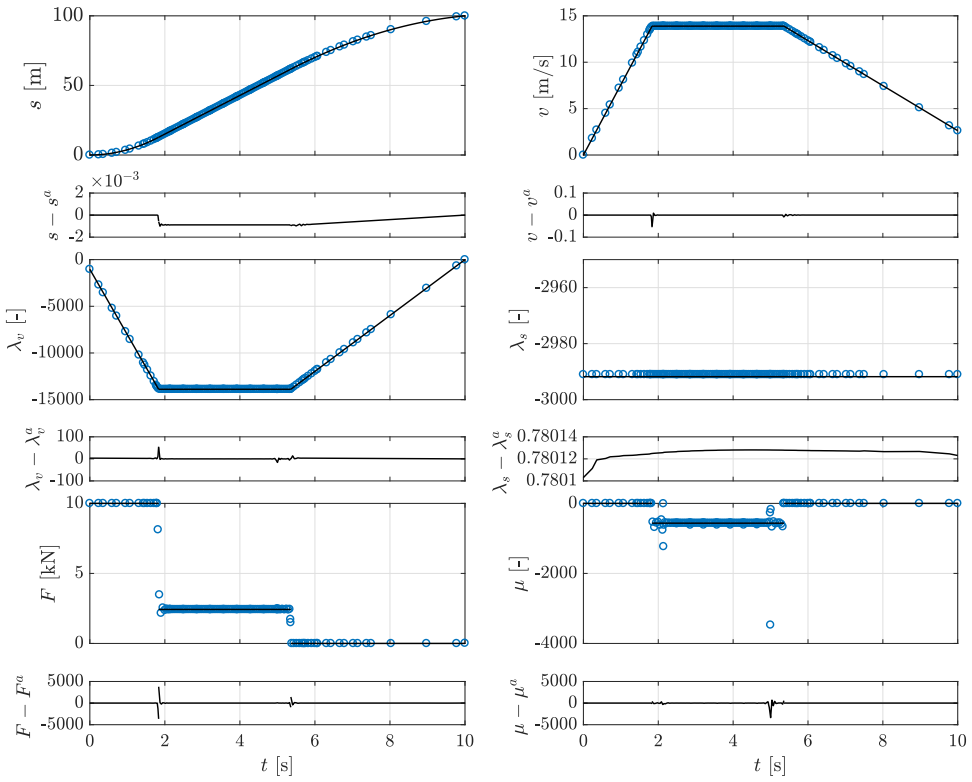


Figure 3.5: A comparison of the numerically computed solution to the problem (circles) compared to the analytically derived solutions of the previous section. The error between each state, control and costate is plotted below each.

Table 3.1: 2-norms of errors in signals on each arc between numerical and analytical solutions.

Trajectory	Sub-arc 1	Sub-arc 2	Sub-arc 3
s [m]	0.001	0.009	0.006
v [m/s]	0.056	0.050	0.011
F [kN]	3.86	3.91	1.48
λ_v [-]	62.35	55.39	24.32
λ_s [-]	7.801	7.801	7.801
μ [-]	2×10^{-6}	0.640	8×10^{-6}

A comparison of the analytically derived states, co-states and path constraint multipliers with their numerical approximation are plotted in Figure 3.5. The circles are the final set of collocation points at which the dynamics are approximated. The OCP is well solved across the horizon. But as the problem has to be regularised, there is a slight error at the beginning and end of the singular sub-arc. This accounts

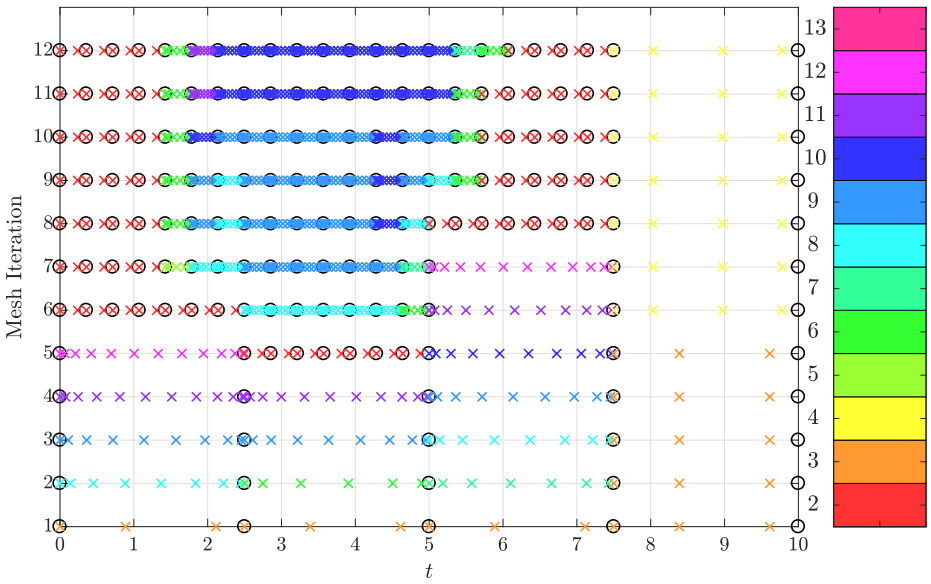


Figure 3.6: Collocation point locations (crosses) and mesh points for each mesh iteration (circles). The legend states the order of each mesh interval (and thus number of collocation points).

for the fairly large error in the recovery of the control and the path constraint multiplier. Table 3.1 displays the errors in states and some numerical metrics for the problem based on the analytical solutions. It is clear once again that the errors are small in magnitude compared to the signals themselves. Adjustment of the regularisation weighting can improve the transitioning between arcs. Reducing the penalty on the gradient of the control will enable faster switching. Increasing the weight too far will result in slow transitions between control arcs, and increased error. However, if the weighting is too small, the problem will remain too close to the original singular problem and the optimisation program will have to move through an increased number of mesh iterations and may not converge effectively.

The mesh refinement is detailed in Figure 3.6. The OCP is initialised with 4 mesh intervals, each containing 3 collocation points. The algorithm for mesh refinement favours using several mesh intervals on problem exit (in the periods before and after the singular arc). Each of these has the minimum number of collocation points (set to 2). The initial mesh has 4×3 collocation points which increases to 141 on the final mesh.

During the singular arc between τ_1 and τ_2 the problem requires a large number of collocation points as the problem is regularised here. In the case that the program was un-regularised, the singular control would have been unrecoverable. The program would exit with an incredibly large number of mesh intervals in which the highest order polynomial was used in each. The solution will oscillate uncontrollably. This simple exercise has enabled validation for the algebraic solutions of Section 3.1. It has also ensured that the solution is still solvable despite its singular nature. Later problems will have an unknown switching structure, but with these numerical methods it is feasible to compute and appraise them.

3.2 Variable Velocity Limit

The ideas presented for the example of constant state limit can be extended to the case where the speed constraint is dependent on the elapsed distance along the course. One example of this appears as a lateral acceleration limit for the car due to turns in the road. This limit will depend on the road curvature, $\mathcal{C}$, in the lateral plane. This measures how much the course turns left/right, as appears to the driver.

Algebraically, the curvature is related to the forward velocity $v(t)$ and the acceleration limit ζ via:

$$\mathcal{C}(s(t))v(t)^2 \leq \zeta \tag{3.79}$$

with $\mathcal{C}(s(t))$ the distance dependent curvature. The rest of the optimal control problem remains as (3.1)-(3.7). The slightly modified Hamiltonian for the problem can be formed

$$\mathcal{H}(t) = F(t)v(t) + \lambda_v(t)\Big(F(t)/M - K\Big) + \lambda_s(t)v(t) + \mu(t)\Big(\mathcal{C}(s(t))v(t)^2 - \zeta\Big). \tag{3.80}$$

The new co-state dynamics are:

$$\dot{\lambda}_v(t) = -(F(t) + \lambda_s(t) + 2\mu(t)\mathcal{C}(s(t))v(t)) \tag{3.81}$$

$$\dot{\lambda}_s(t) = -\mu(t)\frac{d\mathcal{C}(s)}{ds}v(t)^2 \tag{3.82}$$

The switching function remains as in the first example, $\Phi(t) = v(t) + \lambda_v(t)/M$, and again a *bang-bang* control structure will arise in the case that the velocity constraint is never active. The vehicle ascends the hill in the same manner as in the unconstrained example. To the driver, it appears that the road is not curved. In the case that the constraint is met, $v^*(t) = \sqrt{\frac{\zeta}{\mathcal{C}(s)}}$ along the resulting singular arc. The singular control will therefore be given by

$$F_s(t) = MK + M\dot{v}^*(t) \tag{3.83}$$

This is simply the state equation for the velocity. The longitudinal acceleration of the vehicle now depends on the gradient of the curvature.

$$\dot{v}^*(t) = -\frac{ds}{dt}\sqrt{\zeta}\frac{1}{2\mathcal{C}(s)^{3/2}}\frac{d\mathcal{C}(s)}{ds} \tag{3.84}$$

$$= -\frac{\zeta}{2\mathcal{C}(s)^2}\frac{d\mathcal{C}(s)}{ds} \tag{3.85}$$

The switching structure is unknown and cannot be derived analytically. Any number of switches, and thus control arcs could occur. From PMP, it is known that the maximum number of switches that one can solve for analytically is determined by the number of states the system has [31]. In order to see how the solution behaves in this scenario, the vehicle is driven up the curve in Figure 3.7.

The slope is again set to a constant value ($\theta = 0.1$rad) and the lateral shape is defined by the lemniscate of Bernoulli [60].

This curve is parametrised as follows

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} \frac{a\sqrt{2}\cos(2\pi s/s_f)}{\sin^2(2\pi s/s_f)} \\ \frac{a\sqrt{2}\cos(2\pi s/s_f)\sin(2\pi s/s_f)}{\sin^2(2\pi s/s_f)} \\ s\sin\theta \end{bmatrix} \tag{3.86}$$

The curvature of the lemniscate in the tangent plane can be expressed analytically as

$$\mathcal{C}(s) = \frac{3\sqrt{2}\cos(2\pi s/s_f)}{a\sqrt{3-\cos(4\pi s/s_f)}}. \tag{3.87}$$

with the final distance defined as $s_f = 5.244a$. The curve parameter a is chosen such that the course is approximately 100m in length. The lemniscate is a useful

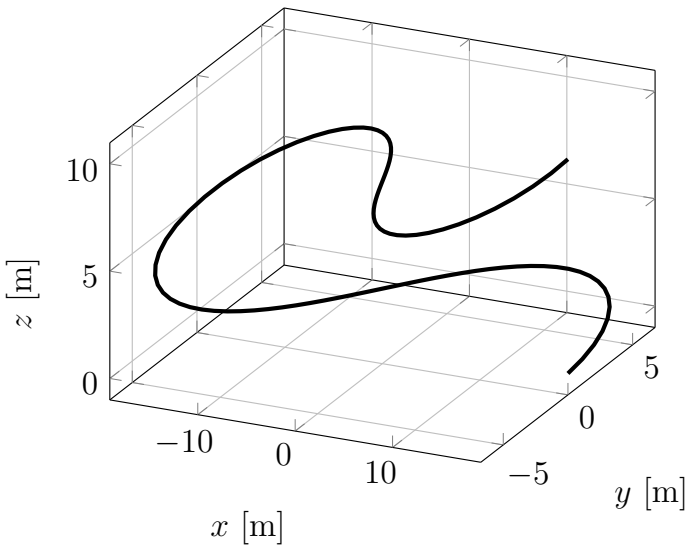


Figure 3.7: The lemniscate for the lateral acceleration limited problem. The start point is at the bottom of the track at the point $(a, 0, 0)$ and ends directly above at $(a, 0, 10)$. The track is 100m long.

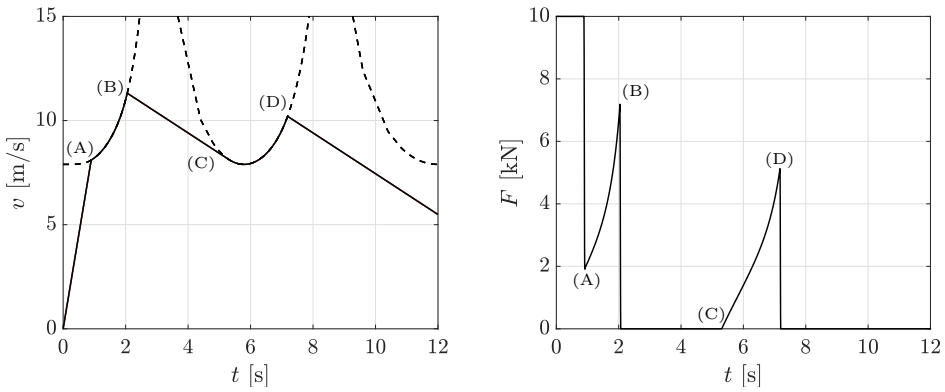


Figure 3.8: Left: Optimal velocity profile for the curvature constrained case in solid with the velocity limit plotted as the black dashed line. Right: Optimal force trajectory.

curve to consider as there are instances where the curvature is zero (i.e the road is straight) at $s = \frac{s_f}{4}$ and $s = \frac{3s_f}{4}$. This represents instances where the lateral acceleration constraint will not be active..

It is reasonable to assume that a similar optimal control structure will arise: a propulsive arc accelerates the car from rest to the speed limit, once at this limit the vehicle is propelled by the singular force (3.83). The car may then be coasted and propelled again until it reaches the end of the course. This behaviour is determined by the arrival time.

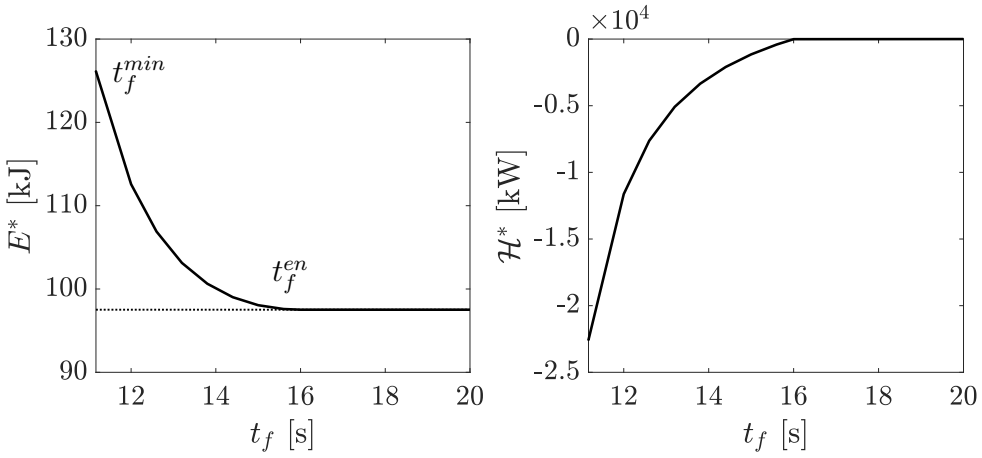


Figure 3.9: Left: Variation of the optimal energy input for the ascent based on the final arrival time. Minimum energy E^* is dotted. Right : Computed Hamiltonian for the ascent under variable velocity limits.

The problem is solved for an arrival time of $t_f = 12$s in the GPOPS-II suite. The optimal velocity and force profiles of which are plotted in Figure 3.8. The solution structure remains similar to that for the constant constraint case - namely through the use of full-propulsive arcs followed by adherence to the speed limit.

First the car is accelerated at maximum propulsive force until the velocity limit is met at (A). The force then continues to accelerate the vehicle along the singular arc (the dotted line in the RHS of Figure 3.8).

The point along the singular arc which determines the point at which the car is no longer accelerated (B), is the point at which the deceleration coincides with the newly positive singular force at point (C). If the vehicle was accelerated further past the point (B) then a braking force would be required to satisfy the limit at the point (C). This coincidence begins a second phase of acceleration, from (C) to (D), which ensures the vehicle completes the course in the allotted time. As it is no longer feasible to solve the system dynamics analytically, it is not possible to derive expressions for the minimum arrival time with the corresponding switching times and energy usage. However, the minimum energy case still corresponds to the situation where the mass comes to rest at the end of the route; $(E^*)^{min} = MKs_f$.

It is possible to investigate the effects of the arrival time on metrics of the problem. Figure 3.9 examines the variation of the optimal cost, E^* and the

control Hamiltonian, $\mathcal{H}^*$. As in the previous two examples, the minimum energy is discovered as the arrival time limit is increased. The second propulsive arc shortens as the arrival time increases. The minimum energy arrival time which preserves the solution structure is approximately 16.1 seconds. Beyond this, the average velocity further decreases, and the lateral acceleration limit is not met at any point along the course. The terminal velocity is always 0 and the optimal force profile can take any shape. The Hamiltonian is zero when the final time approaches (and surpasses the minimum energy time). Also as the time approaches the minimum route time, the Hamiltonian approaches $-\infty$. These results are consistent with the constant velocity limit and unlimited cases of the previous section.

3.3 Variable Road Grade

So far, the effects of velocity limits on the cars ascending of roads of constant incline have been discussed. One of the more important features of optimal driving is longer routes with variable slopes. Still in the absence of resistive forces, the car is driven along the variable slope road in Figure 3.10. This road profile can be written as the cubic

$$y(x) = a_3x^3 + a_2x^2 + a_1x. \tag{3.88}$$

The course parameters are chosen such that there is a peak at an elevation of 10m at 100m in the $x-$direction, followed by a dip 5m above the datum at 250m. The route ends at 350m in the $x-$direction, with a final height of 15m. The total spine length of the course is $s_f = 352$m.

The original dynamics were posed in the car's co-ordinate frame; the velocity and elapsed distance metrics are related to the tangential distance along the road. In order to effectively use the undulating road definition, a transformation must take place which measures the cars progress along the x-direction[3]. Geometry dictates

[3]It is possible to compute the inclination angle for this course. Clearly, $\theta = \arctan\frac{dy}{dx}$. This can be indexed as a function of the elapsed distance along the curve which can be calculated with $s = \int_{x_0}^{x_f}\sqrt{1+\frac{dy}{dx}^2}dx$. As the velocity is always positive, this is a one-to-one mapping. This is a non-trivial integral to compute analytically, so to avoid numerical integral error, a coordinate transformation is utilised.

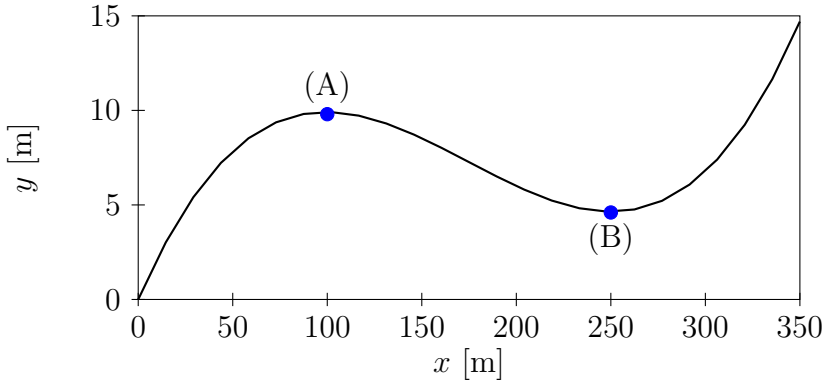


Figure 3.10: Example of an undulating road with elevation $y = 15$m at $x = 350$m. The peak at (A) has elevation $y = 10$m at $x = 100$m, while the dip at (B) has elevation $y = 5$m at $x = 250$m. The road length is $s_f = 352$m.

that the road angle θ is related to the gradient of the curve

$$\sin\theta = \frac{\frac{dy}{dx}}{\sqrt{1+\left(\frac{dy}{dx}\right)^2}} \quad \text{and} \tag{3.89}$$

$$\cos\theta = \frac{1}{\sqrt{1+\left(\frac{dy}{dx}\right)^2}}. \tag{3.90}$$

This enables the existing system dynamics (3.2) and (3.3) to be modified. The dynamic equation describing the velocity now takes the form:

$$\dot{v}(t) = \frac{F(t)}{M} - g\frac{\frac{dy}{dx}}{\sqrt{1+\left(\frac{dy}{dx}\right)^2}} \tag{3.91}$$

Velocity in the direction parallel with the x-axis is then given by

$$\dot{x}(t) = \frac{v(t)}{\sqrt{1+\left(\frac{dy}{dx}\right)^2}}. \tag{3.92}$$

The elapsed distance along the course state s, is now replaced with the elapsed distance along the x-axis, x. The car is likely to need to brake during the descent of the hill from (A) to (B) to prevent violation of the speed constraint. The lower bound on the propulsive force (3.5) is relaxed to the maximum braking effort $F_b = -10$kN. The control force is now:

$$F_b \leq F(t) \leq F_m \tag{3.93}$$

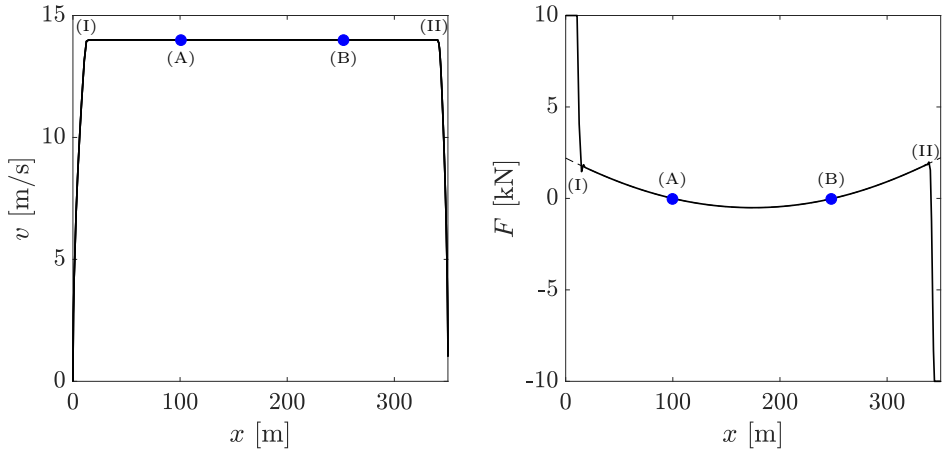


Figure 3.11: Left: The optimal velocity profile for an arrival time of 26.5s. Right: The optimal force control profile with the singular force (dashed).

The vehicle parameters and speed limit for this example are the same as the example in Section 3.1.2. The final modification for the variable road grade case offers insight into the effects of regenerative braking vs conventional friction brakes. This comes in the form of a different cost function to be minimised. For regenerative braking, energy recovered in braking reduces the cost. Therefore the cost will remain the same as (3.1). In the case of conventional friction brakes this becomes

$$E = \int_0^{t_f} \max(F(t), 0) v(t) dt. \tag{3.94}$$

Braking no longer reduces the cost in this case. This final comparison considers the effects of differing powertrains on the obtained solution. As the OCP is now considerably more complex and non-linear, the solution cannot be derived by hand easily, so the solution will be obtained using GPOPS-II.

Regenerative Braking

In the case of regenerative braking, there is no energy store, and therefore no losses associated with braking and deploying power. The optimal propulsion trajectory for the regenerative braking case is plotted in the RHS of Figure 3.11 for an arrival time of 26.5s. The solution takes a similar form to that for the constant slope angle case; bang-singular-bang. The car is accelerated to the speed limit (I) on the first control subarc. Upon reaching the speed limit, the driving force is modulated to

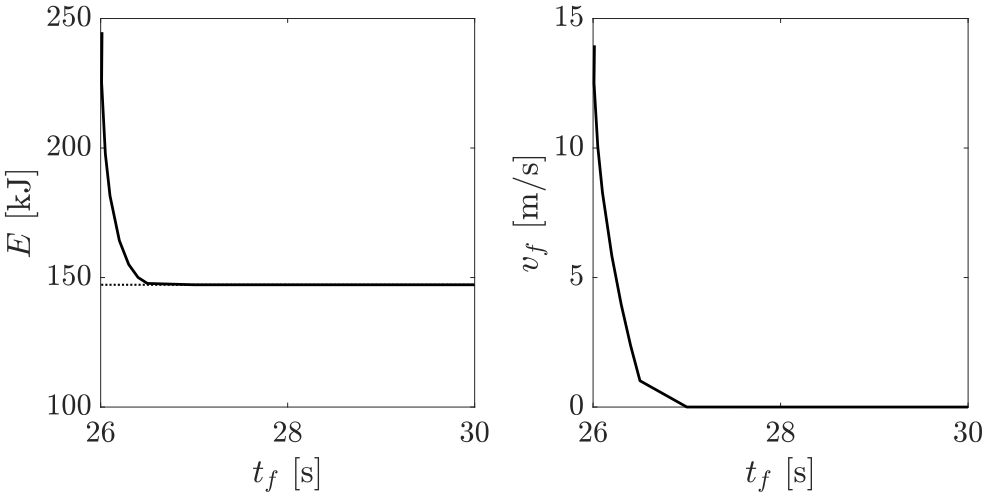


Figure 3.12: Left: The effect of maximum arrival time against the optimal energy input cost. The minimum cost, E^*, is plotted as a dotted line. Right: The terminal velocity variation of the vehicle against the arrival time.

balance the varying vehicles weight along the singular subarc between (I) and (II). At the peak at point (A) the driving force is instantaneously zero. After the peak, the vehicle is braked to avoid violating the speed limit. Upon reaching the dip at (B), the vehicle is propelled with a positive driving force once again. The vehicle is finally braked to rest at the end of the course (II).

The optimal speed profile is plotted in the LHS of Figure 3.11. The vehicle is accelerated under the maximum propulsive force to the speed limit, at which it is held for most of the course (between (I) and (II)). Along the final sub-arc, the vehicle is braked to zero velocity. This reduces the cost by removing all of the car's kinetic energy. The optimal energy expenditure is plotted in Figure 3.12 for varying the arrival time. The familiar bang-bang structure and non-zero terminal velocity conditions are only valid for a small interval of arrival times. The course minimum traversal time is $26s$ and in the case that the arrival time is greater than $27s$ then the velocity constraint is no longer active at all. The steepness of the road at the end of the course means that the velocity can be reduced to 0 effectively, almost without requiring braking. However, it is always beneficial to brake the vehicle at the end of the course to achieve the zero velocity terminal condition synonymous with the minimum energy traversal.

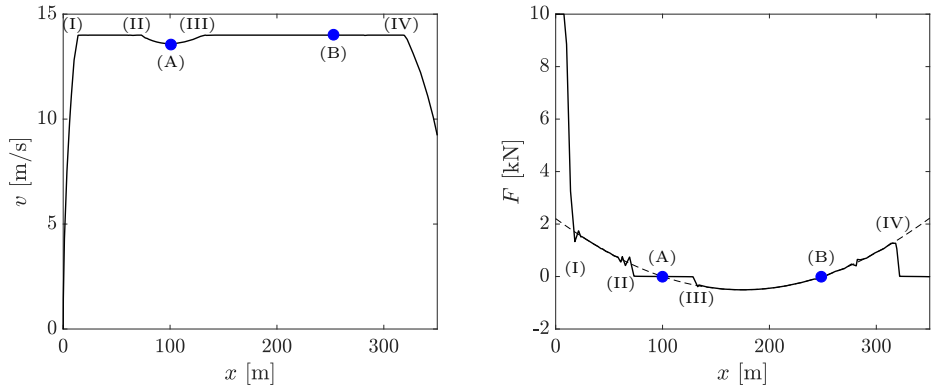


Figure 3.13: Driving and braking force with optimal dissipative braking. The component of the weight parallel to the slope is dashed.

Dissipative Braking

Arcs of zero propulsion will replace braking arcs as the cost cannot be decreased by braking in this case. The optimal propulsive law once again consists of bang switches and multiple sub-arcs[4]. An example solution is plotted in Figure 3.13 for an arrival time of 26.5 s.

The vehicle is once again accelerated up the speed limit (at (I)) under the maximum driving force. Between (I) and (II) the force is singular: matching the component of the vehicles weight parallel to the slope. As the vehicle approaches the top of the hill before (A), it is coasted on the subarc between (II) and (III). This coasting arc ensures that braking required during the descent from (A) to (B) is minimised. During this arc, the car's speed decreases away from the limit.

This period of coast reduces the journeys average speed, meaning its duration is determined by the arrival time chosen. During the descent of the hill, between (III) and (B), the vehicle is braked. The driving force is instantaneously zero at (B) and the car is re-accelerated up the final incline to (IV). The car is then coasted over the line and not braked to rest. It is clear that coasting is favoured over braking, but braking cannot be avoided due to the tight arrival time selected. Immediately this situation differs from the regenerative braking one.

[4]As the cost function requires the evaluation of the non-smooth $\max(x, 0)$ function, the smoothing laid out in Section 2.3.5 is implemented

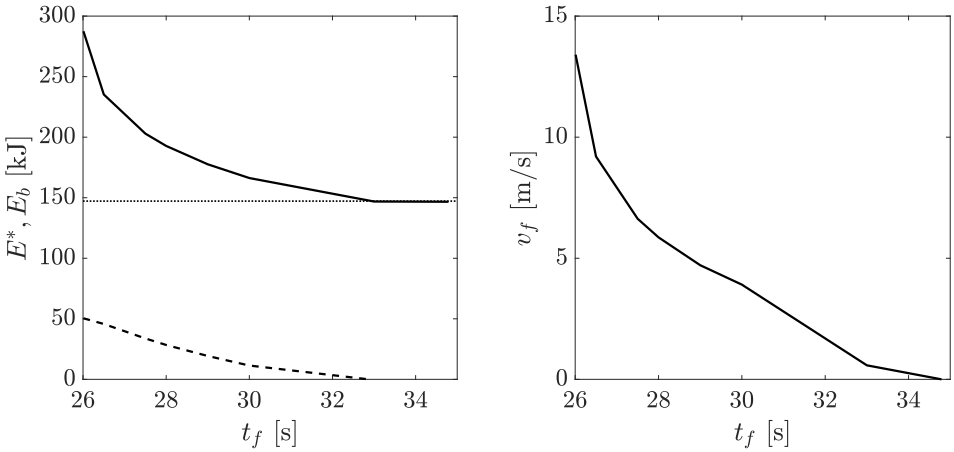


Figure 3.14: Left: The effect of maximum arrival time against the optimal energy input cost. The minimum cost, E^*, is plotted as a dotted line. The braking energy is plotted as a dashed line. Right: The terminal velocity variation of the vehicle against the arrival time.

The effect of changing the arrival time on the optimal cost is examined in Figure 3.14. A much broader series of arrival times are possible compared to the regenerative braking case. The minimum arrival time is broadly the same as the regenerative braking case (approximately 26 seconds). However, there is an increase in the cost, highlighting the usefulness of the regenerative braking powertrain. Increasing the arrival time, increases the zero propulsion arc from (II) to (III). As the coasting arc increases in length, the braking segment from (A) to (B) decreases accordingly. This decreases the total braking energy E_b, dashed line in LHS of Figure 3.14. At the minimum energy solution there is no braking. Instead the car coasts from (II) to (B) and the car does not reach the speed limit.

The simple, *bang-singular-bang* solution found in the constant road grade examples no longer carries straight over to the case where the gradient is not constant. Even in this simple undulating road example, the solution structure was impossible to infer from the road profile alone. This effect would have been further exacerbated if the data was not analytical but rather taken from GPS data. The effect of the arrival time has also been examined - highlighting the requirement for its careful selection. The example of the undulating road

also highlights the difference in solution structure which can arise from differing powertrain configurations traversing the same route.

3.4 Conclusions

This chapter considered the solution of minimum-energy optimal control problems set in a vehicular context. Initially, a simple example of a vehicle ascending a slope of constant incline is presented and analytical results were derived and appraised. Conditions for tractability along with the extreme values for arrival time and total energy input are considered. This example can be extended to the case when a speed limitation is active. The new control structure can still be derived, and the corresponding expressions for minimum time and energy obtained.

Further analysis took place which facilitated the understanding of upper bounding, rather than fixing, the time of arrival. In either case, it is worth noting that arrival time is one of the more important selectable parameters for generating driving data. In the case that the arrival time is relaxed too far, the solution of such a problem ceases to be practically relevant. In the context of driving in the real world, this would be deemed either unsafe due to other traffic or unsuitable in the context of arriving at the destination on time.

Simply extending the problem to the case where the road is no longer of constant slope requires solution of the OCP via numerical methods. Even in this simple example the effects of regenerative braking vs dissipative braking can be appraised. Control strategies of multiple propulsion, coasting and braking arcs appear in the solution of these problems. The difficulty of solution by analytical means limits the capability of closed form solution for the variable speed limit case. Both the variable slope and speed limit scenarios are solved for a variety of arrival time constraints. To ensure that the solutions are reasonable, the speed limited case for constant incline is solved using GPOPS-II and verified with the analytical results.

Despite the analysis for arrival time constraint fulfilment taking place in the absence of losses, it can be reasoned that the arrival time will be fulfilled. In order to minimise the total energy input, one must avoid unnecessarily high speeds which

increase losses due to drag. Based on this, the total time to traverse the route will be increased as the average velocity reduces. In the case that the arrival time is upper bounded, this constraint will always be met.

A number of key insights are noted here to take forward to the subsequent chapters:

1. Due to the nature of energy optimal control problems, singular arcs are likely to occur especially in the case of aggressive arrival times close to the minimum time to complete the course. It is important therefore to use some regularisation or other methods to ensure that the solution can be recovered.

2. Different solutions may arise due to the powertrain used for propulsion in the OCP. This suggests that there is no simple power/velocity profile optimisation problem that can be solved and applied to all road vehicles. The nature of the propulsion method will have an effect on the optimal solution.

3. The selection of arrival times is crucial to the solution tractability and altering this can result in different results. Additional constraints on the problem (such as total integral constraints on emissions) can effect the arrival time constraint satisfaction and also change the optimal solution anyway.

In the following two chapters, the results found in this chapter are extended to more realistic driving scenarios. Two vehicles, one equipped with a diesel internal combustion engine, the other with a hybrid powertrain, are driven on routes that are predetermined and are described by their geometry and velocity limits. The optimal control problems are solved for a number of cases including the implementation of emissions limits.

Everything in life is somewhere else, and you get there in a car.

— E.B. White, *Fro-Joy*, January 1940

4 Fuel-Optimal Driving

Contents

4.1 Car Model 90
4.1.1 Vehicle Dynamics 90
4.1.2 Powertrain Model 91
4.2 Driving Optimal Control Problem 97
4.2.1 Multi–phase Numerical Optimal Control 101
4.2.2 Discontinuous Functions 102
4.3 Results 103
4.3.1 Power Transmission Comparison 103
4.3.2 Emissions Sensitivity 107
4.3.3 Enforced Stops 110
4.4 Conclusions 114

Fixed drive cycles have remained the norm for engine testing and vehicle homologation for a number of years as discussed in Chapter 1. More recently it has become apparent that there are shortcomings in these testing procedures. As a result, the road car industry is under scrutiny to ensure that emissions testing methods represent real-world driving and thus real-world emissions production. Engine calibration makes use of static lookup maps that determine engine operation under a prescribed set of torque-speed conditions [10]. These maps are typically generated with high-dimensional static engine models over fixed drive cycle scenarios. This testing procedure does not represent real world driving. New tests are now

in place which require portable emissions equipment to be installed on the vehicle to measures the production of oxides of nitrogen (NO_x), carbon monoxide (CO), unburnt hydrocarbons (HC) and particulate matter (PM — HC and soot). This ensures that the real world emissions production falls within the legislative guidelines. As there is an element of real world driving incorporated into the vehicle testing and homologation, it is beneficial to be able to evaluate vehicle performance offline.

Optimal control studies, using dynamic engine and vehicle models, can be used for the offline evaluation of a number of different control strategies. The solution of these problems usually derives from Pontryagin's Maximum Principle (PMP) [31]. Alternatively, solutions for simplified problems can be derived analytically; see for example [58, 61]. Optimal driving can be split into segments that comprise such things as maximum power acceleration, constant speed, coasting, and limit braking. However, it is not necessarily optimal to drive in the same manner when emissions constraints are introduced, and indeed faster route traversals may be used to ensure that the emissions constraints are met [62].

Dynamic programming (DP) can also be used to compute the optimal control solution, and is one of the most widely used solution techniques [31]. While this method is well suited to problems with few state and control variables, increasing the complexity of the model leads to the 'curse of dimensionality'. The inclusion of terminal constraints in the cost function [63] facilitates the efficient use of DP to investigate the optimal trade-off between NO_x and fuelling. Optimal eco-driving cycles for hybrids was solved with DP in [64], where it is noted that the optimality of the solution depends heavily on the driver accepting the driving profile given. An online driver assistance system could be based on this concept. The optimal driving solution is computed online, and the driver is advised on the most efficient course of action (in terms of the engine torque and gearing) [65]. The optimisation of driving style and powertrain control for a vehicle equipped with a hybrid powertrain is studied in [66]. A short countryside route was considered in order to suppress the influence of traffic. This study highlighted the benefits of optimising both the power usage and the driving profile.



Figure 4.1: The closed driving route (denoted Route A) around Oxford, UK, utilised in this chapter.

The results of the optimal control could form the basis of eco-driving guidelines that depend on the scenario under consideration. In this context a number of different optimisation techniques may be used to advantage [67]. A pseudospectral method was employed to solve the optimal fuel economy optimal power split problem for a series hybrid electric bus in [68]. This was shown to be computationally more efficient than a standard DP approach. Improvements in solution accuracy are possible [69]. The authors of [70] presented a formulation of the eco-driving problem with particular emphasis on the high-fidelity modelling of both the vehicle and the road. The solution method enabled the use of vehicle model information which would otherwise be neglected. In the context of drive cycle generation the method does not scale well.

This chapter outlines a novel method for generating fuel-optimal engine usage trajectories which can be utilised in the creation of optimal engine calibrations. These profiles will be the optimal driving solution based on the route in Figure 4.1 (which is dubbed Route A), which includes both urban and fast motoring sections. Simultaneous driving and powertrain control are solved for using a direct method outlined in Chapter 2. The route is traversed with a vehicle equipped with both a continuously variable transmission and a discrete gearbox with a shifting law.

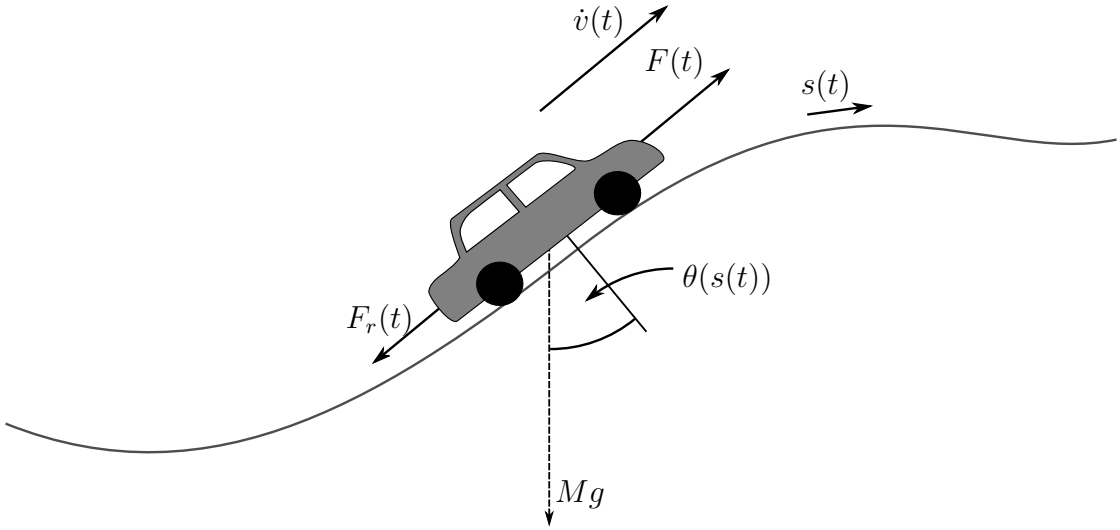


Figure 4.2: Forces acting on the vehicle on a general undulating road.

Additionally, the sensitivity of the resulting solution to emissions constraints, as well as the need to stop at junctions, will be examined.

In order to check the validity of the optimisation result, the route is driven by a real vehicle, and velocity data measured. The results of the comparison are discussed in Appendix C.

Material from this chapter appears in [36, Chap. 9.2] and [59]. An early version of this work was prepared in [62].

4.1 Car Model

This section presents the vehicle and powertrain model that will be used in the remainder of the chapter.

4.1.1 Vehicle Dynamics

Prior work that addresses the optimisation of a hybrid vehicle on a 3D track, in combination with a non-linear tyre model, is [70]. As the route length is much greater than the effects of roll and pitch can be ignored as the dynamic effects are considered 'fast'. As this work is mainly focused on the effects of powertrain limitations rather than vehicle dynamics, a simple single-degree-of-freedom model of the vehicle longitudinal dynamics is used.

A force balance between the tractive force $F(t)$ and resistive forces $F_r(t)$ (in Figure 4.2) can be expressed as

$$\dot{v}(t) = \frac{1}{M}\Big(F(t) - F_r(t)\Big), \tag{4.1}$$

where M is the vehicle mass. The resistive force is made up of three components: aerodynamic drag, rolling resistance and the component of the vehicle's weight projection along the slope:

$$F_r(t) = Mg(\mu \cos\theta(s(t)) + \sin\theta(s(t))) + \frac{1}{2}\rho C_d A v(t)^2. \tag{4.2}$$

The aerodynamic drag coefficient is C_d, A is the vehicle frontal area and ρ the density of air. The coefficient of rolling resistance is μ and g is the acceleration due to gravity. The elapsed distance is $s(t)$, while the road inclination angle is $\theta(s(t))$. Equation 4.1 represents the only vehicle system dynamics used here.

4.1.2 Powertrain Model

The engine dynamics will be considered 'fast', and will thus be neglected in the model used here; the engine data represents a prototypical mid-size diesel car engine.

Fuel and Emissions Models

In order to predict the engine's fuel consumption and emissions production, two-dimensional functional maps based on test bed data are used[1]. The brake-specific fuel consumption (BSFC) is a useful measure of the engine operating efficiency and is given by the ratio of the fuel mass flow rate ($\dot{m}_f$) to the engine output power

$$\text{BSFC} = \frac{\dot{m}_f(\omega_e, \tau_e)}{\omega_e \tau_e}. \tag{4.3}$$

The engine brake power is defined as the product of engine speed (ω_e) and engine torque (τ_e). The map given in Figure 4.3 includes the thermal efficiency of combustion, frictional losses and pumping losses. In similar vein, the brake-specific (BS) emissions, which include the efficiencies relating to production of emissions, are shown in Figures 4.4 and 4.5.

[1]The engine data used in this study was provided by Ricardo PLC.

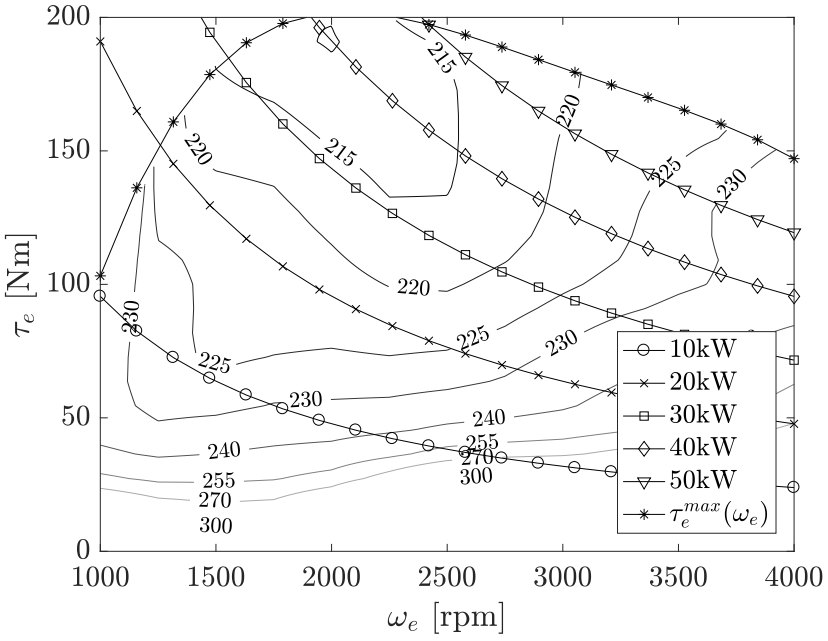


Figure 4.3: Brake specific FC (BSFC) map [g/kWh] including maximum torque line and lines of constant power.

For optimal control purposes, smooth maps for the prediction of fuel and emissions flows are required. The fuel mass flow rate can be represented by a biquadratic map [62]. The engine operating range is constrained by the idling speed ω_e^{idle}, the maximum engine speed ω_e^{max}, and zero torque and the maximum torque line $\tau_e^{max}(\omega_e)$. This limit is given by

$$\tau_e^{max}(\omega_e) = a_4\omega_e^4 + a_3\omega_e^3 + a_2\omega_e^2 + a_1\omega_e + a_0. \tag{4.4}$$

The coefficients a_i are computed using a curve fitting tool; the results of which were used to produce the τ_e^{max} curve in Figure 4.3.

The nonlinear response of emissions production to the engine operating condition is represented by Gaussian Process models of the type described in [71, 72]; further detail is available in [73]. This modelling method has the advantage that no black-box model is required, instead samples from a distribution of functions are used.

In general, Gaussian process models take the form

$$\mathbf{y} = f(\mathbf{x}) + \epsilon, \tag{4.5}$$

which relate the vector of inputs $\mathbf{x}$ to the vector of outputs $\mathbf{y}$ via the unknown, underlying stochastic process $f(\mathbf{x})$. The output is measured, and some zero-mean Gaussian noise ϵ, is introduced. The variance of the noise is σ_n^2. The function $f(\mathbf{x})$ is a static input-output relationship and is assumed to have zero-mean. The covariance for outputs $\mathbf{y}$ and $\mathbf{y}'$, corresponding to the inputs $\mathbf{x}$, $\mathbf{x}'$, is given by

$$cov(f(\mathbf{x}), f(\mathbf{x}')) = k(\mathbf{x}, \mathbf{x}'). \tag{4.6}$$

The *covariance* function $k(\mathbf{x}, \mathbf{x})$ is selected by the user. Consider a set of training data inputs $\mathbf{x}_i \in \mathbb{R}^{1\times d}$ with corresponding output response, $\mathbf{y}_i \in \mathbb{R}^{1\times p}$. This can be written succinctly in terms of matrices: $\mathbf{X} \in \mathbb{R}^{n\times d}$ and $\mathbf{Y} \in \mathbb{R}^{n\times p}$. Therefore, based on (4.5), the output data can be thought of as samples at each $\mathbf{x}_i$ taken from the distribution:

$$\mathbf{Y} \sim \mathcal{N}(\mathbf{0}, \mathbf{K}(\mathbf{X}, \mathbf{X}) + \sigma_n^2 I). \tag{4.7}$$

The matrix $\mathbf{K}(\mathbf{X}, \mathbf{X})$ is an $n \times n$ matrix of the covariance between each of the training input points.

The joint distribution between the training data outputs $\mathbf{Y}$ and a new, noise-free test output $\mathbf{y}^* = f(\mathbf{x}^*)$ is

$$\begin{bmatrix} \mathbf{Y} \\ \mathbf{y}^* \end{bmatrix} \sim \mathcal{N}\left(\mathbf{0}, \begin{bmatrix} \mathbf{K}(\mathbf{X}, \mathbf{X}) + \sigma_n^2 I & \mathbf{k}(\mathbf{x}^*, \mathbf{X})^{\mathrm{T}} \\ \mathbf{k}(\mathbf{x}^*, \mathbf{X}) & \mathbf{k}(\mathbf{x}^*, \mathbf{x}^*) \end{bmatrix}\right). \tag{4.8}$$

The entries of $\mathbf{k}(\mathbf{x}^*, \mathbf{X})$ are the covariance function between each of the training points and the test point $\mathbf{x}^*$. In order to predict from the distribution of functions, the joint Gaussian prior distribution is conditioned on the observations (Appendix A.2 in [73]):

$$\mathbf{y}^*|\mathbf{y}, \mathbf{X}, \mathbf{x} \sim \mathcal{N}(\mathbf{k}(\mathbf{x}^*, \mathbf{X})^{\mathrm{T}}(\mathbf{K}(\mathbf{X}, \mathbf{X}) + \sigma_n^2 I)^{-1}\mathbf{y}\ , \\ \mathbf{k}(\mathbf{x}^*, \mathbf{x}^*) - \mathbf{k}(\mathbf{x}^*, \mathbf{X})(\mathbf{K}(\mathbf{X}, \mathbf{X}) + \sigma_n^2 I)^{-1}\mathbf{k}(\mathbf{x}^*, \mathbf{X})^{\mathrm{T}}). \tag{4.9}$$

The expected value of the output, given the training data and the new test point is therefore the mean of this distribution and is given by

$$\mathbf{y}^* = \mathbf{k}(\mathbf{x}^*, \mathbf{X})^{\mathrm{T}}(\mathbf{K}(\mathbf{X}, \mathbf{X}) + \sigma_n^2 I)^{-1}\mathbf{y}. \tag{4.10}$$

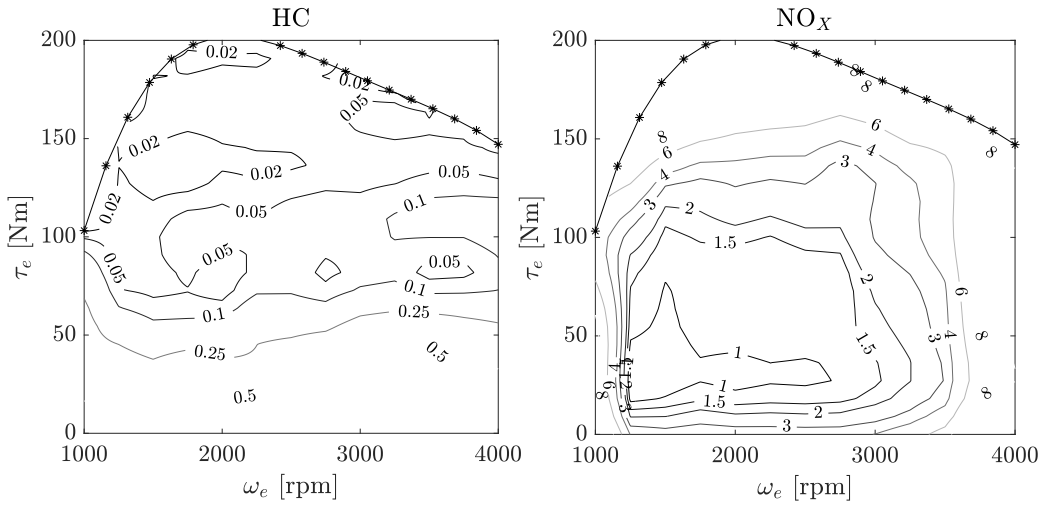


Figure 4.4: Brake specific emissions maps [g/kWh] (Left - unburnt hydrocarbons (HC), Right - NO_x; the maximum torque line is overlaid.

Training data with fixed nonzero mean can be treated in the same way. The zero-mean Gaussian Process is applied to the difference between the observations and their mean (β). The predictive mean (4.10) becomes:

$$\mathbf{y}^* = \beta + \mathbf{k}(\mathbf{x}^*, \mathbf{X})^{\mathrm{T}}(\mathbf{K}(\mathbf{X}, \mathbf{X}) + \sigma_n^2 I)^{-1}(\mathbf{y} - \beta). \tag{4.11}$$

The covariance function used for this work is the *squared exponential*:

$$k(\mathbf{x}, \mathbf{x}') = \sigma_F^2 \exp\left(-\frac{|\mathbf{x} - \mathbf{x}'|^2}{2\sigma_L^2} \right). \tag{4.12}$$

The characteristic length scale σ_L^2 and overall variance of the function σ_F^2 are estimated by maximising the log-likelihood of the predictions [73]. In the case for the emissions models used in this study, the input vector is made up of the engine speed and torque, $\mathbf{x} = [\omega_e, \tau_e]$. The now-scalar output y, is brake-specific (BS) NO_x HC, soot or CO. The regression models are generated based on steady-state engine testbed data and are trained with the Statistics and Machine Learning Toolbox in MATLAB. The corresponding maps are those plotted in Figures 4.4 and 4.5.

Power Transmission

Power flows from the engine to the wheels via a transmission. This enables the engine and road wheels to be rotated at different speeds. A transmission is also

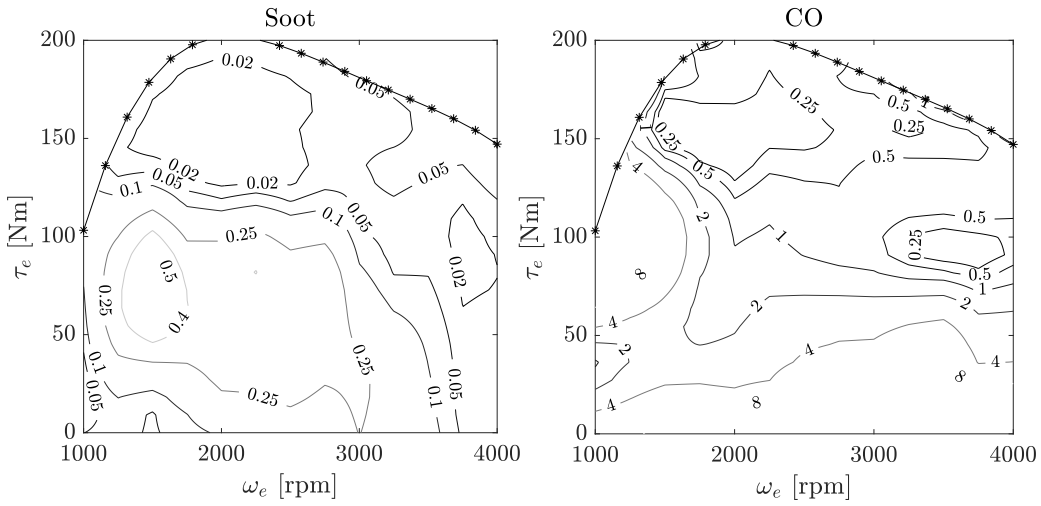
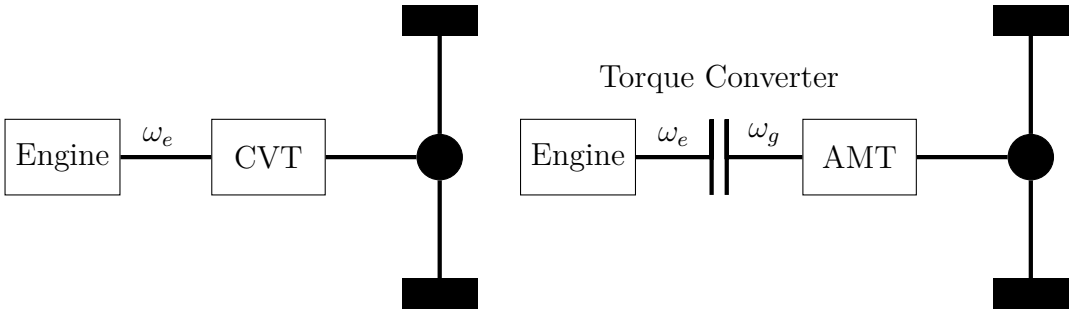


Figure 4.5: Brake specific emissions maps [g/kWh] (Left - soot, Right - CO); the maximum torque line is overlaid.



(a) The layout for the CVT optimisation case. **(b)** The AMT powertrain transmission layout for optimisation indicating the locations of ω_e and ω_g.

Figure 4.6: Transmission layouts for both optimisation scenarios.

required to enable the vehicle to move away from rest. Continuously variable transmissions (CVTs) are mechanical devices with no pre-determined gear ratios which, from an engine operation standpoint, allows any engine speed to be realised [67]. The major drawback to this type of power transmission are the high mechanical losses, for a belt driven CVT this is around 15%. It is not uncommon to make a simplification in the evaluation of optimal control strategies by assuming a CVT rather than a discrete set of gears (see for example [61]). Automated manual transmissions (AMT) are implemented with velocity dependent shifting laws—a major improvement over CVTs are the lower power losses of approximately 5%. The layout of both of these transmissions are plotted in Figure 4.6.

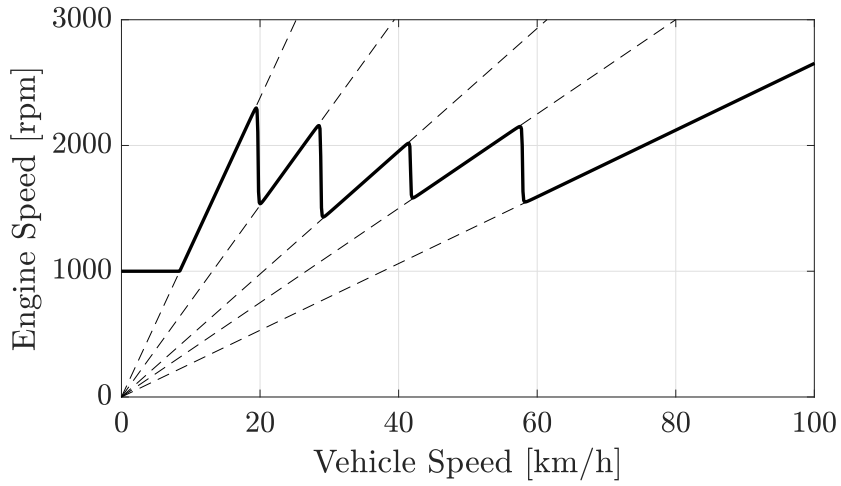


Figure 4.7: Engine speed as a function of forward velocity; the overall drive ratios are shown as dotted lines. The condition (4.15) is clear as the horizontal line up to the dashed line corresponding to first gear.

In order to evaluate the loss of information arising from a 'CVT simplification', the optimal driving strategy will be computed for both types of transmission. The 5-speed AMT with a forward velocity ($v(t)$) dependent shifting law takes the following form (see [74]):

$$\begin{aligned} N_g(v(t)) \approx N_g^1 + 0.5(N_g^2 - N_g^1)(1 + \sin(\arctan(\alpha_1(v(t) - u_1)))) + \\ 0.5(N_g^3 - N_g^2)(1 + \sin(\arctan(\alpha_2(v(t) - u_2)))) + \\ 0.5(N_g^4 - N_g^3)(1 + \sin(\arctan(\alpha_3(v(t) - u_3)))) + \\ 0.5(N_g^5 - N_g^4)(1 + \sin(\arctan(\alpha_4(v(t) - u_4)))). \end{aligned} \tag{4.13}$$

The N_g^is for $i = [1, 2, 3, 4, 5]$ are the fixed gear ratios and the u_is for $i = [1, 2, 3, 4]$ are the gear-shift speeds. The parameters α_i for $i = [1, 2, 3, 4]$ are chosen to make the gear change approximation sufficiently accurate and smooth for optimal control purposes. An example shifting law is plotted in Figure 4.7, the shift speeds u_i are chosen initially so that the shift occurs at approximately $\omega_e = 2000$rpm. These will later be optimised.

The rotational speed of the gearbox on the driveline side of the torque converter (indicated in Figure 4.6b) is given by

$$\omega_g(t) = \frac{N_g(t) N_f v(t)}{r_w}, \tag{4.14}$$

where N_f is the final drive ratio and r_w is the radius of the rear wheel. To model the disengaging of the engine from the wheels, the following constraint is enforced

$$\omega_e(t) - \omega_e^{idle} - \max(\omega_g(t) - \omega_e^{idle}, 0) = 0. \tag{4.15}$$

From rest, when $w_g(t) = 0$, the engine speed is constrained to be equal to ω_e^{idle} until $\omega_g(t) = \omega_e^{idle}$ in first gear (Figure 4.7). During this time, the torque converter is 'slipping'.

4.2 Driving Optimal Control Problem

The minimum fuel control problem can now be assembled. It is convenient to transform the problem from the time domain to the distance travelled domain (s) so that route-specific information can be straightforwardly encoded in the problem setup. There is a one–to–one correspondence between time and the distance travelled provided $v > 0$, since $s = vt$. In differential form there holds:

$$\frac{d}{ds} = \frac{1}{v(s)}\frac{d}{dt}. \tag{4.16}$$

The s-domain formulation is equivalent in size, in terms of the number of states, as the t-domain. However, to properly implement distance based phenomenon such as discrete changes in state limits, or additional event constraints such as enforced stopping, the s-domain formulation is superior.

The OCP can be posed as a fuel minimisation problem by seeking to minimise the following integral

$$m_{fuel} = \int_0^{s_f} \dot{m}_f(\omega_e(s), \tau_e(s))\frac{\mathrm{d}s}{v(s)}, \tag{4.17}$$

in which $\dot{m}_f(\cdot,\cdot)$ is the fuel mass flow rate.

The state dynamics are given by Equation (4.1) following transformation into the s-domain:

$$\frac{dv(s)}{ds} = \frac{1}{Mv(s)}\Big(\frac{P_t(s)}{v(s)} - Mg(\mu\cos\theta(s) + \sin\theta(s)) - \frac{\rho C_d A v(s)^2}{2}\Big). \tag{4.18}$$

The tractive power $P_t(s)$, given by $F(s)v(s)$, along with the engine control inputs ω_e and τ_e, make up the control inputs for the problem. The following inequality constraints must hold:

$$\begin{bmatrix} \omega_e^{idle} \\ 0 \\ P_b \end{bmatrix} \leq \begin{bmatrix} \omega_e(s) \\ \tau_e(s) \\ P_t(s) \end{bmatrix} \leq \begin{bmatrix} \omega_e^{max} \\ \tau_e^{max}(\omega_e(s)) \\ P_e^{max} \end{bmatrix}. \tag{4.19}$$

The engine speed and torque are constrained to remain on the engine maps (see Figures 4.3, 4.4 and 4.5). The tractive power is bounded by the maximum engine power P_e^{max} and the maximum braking power P_b. The power balance between the powertrain and the car is modelled with the following inequality

$$\eta_X \omega_e(s)\tau_e(s) - P_t(s) \geq 0, \tag{4.20}$$

where η_X is the transmission efficiency for either the CVT or AMT.

Driver comfort and tyre usage constraints are captured by bounds on the maximum longitudinal and lateral accelerations; commonly referred to as a *g-g* diagram. A typical driver will not exploit fully the available acceleration limits [75]. An alternative description of the operating region is:

$$\left|\frac{a_x(s)}{a_x^{max}}\right|^n + \left|\frac{a_y(s)}{a_y^{max}}\right|^n \leq 1 \tag{4.21}$$

where n takes values between 1 and 2—corresponding to a rhombus and circle respectively. The lower the value of n, the more conservative the driver is with the available accelerations—especially under cornering conditions. The longitudinal acceleration $a_x(s)$ is given by the right-hand side of (4.1). The lateral acceleration is given by

$$a_y(s) = \mathcal{C}(s)v(s)^2. \tag{4.22}$$

Finally it is necessary to impose a number of integral constraints. The first of which is the upper bound on the arrival time

$$\int_0^{s_f} \frac{1}{v(s)}\, ds \leq T. \tag{4.23}$$

Table 4.1: Table of the vehicle parameters used in the optimisation.

Symbol	Value
M	$1750kg$
C_dA	$0.6m^2$
ρ	$1.2kg/m^3$
a_x^{max}, a_y^{max}	$0.1gm/s^2$
μ	0.01
n	1
η_{CVT}	0.85
η_{AMT}	0.95
P_b	$-30kW$
P_e^{max}	$80kW$
ω_e^{idle}	$1000rpm$
ω_e^{max}	$4000rpm$
N_g	$[4.48, 2.87, 1.84, 1.41, 1]$

The bound on T acts as a surrogate for the driving style, a lower arrival constraint implies more aggressive driving. Particulate matter (PM) is made up of unburnt (HC) and soot according to the following mixing law [76]:

$$\dot{m}_{PM} = 1.024\dot{m}_{Soot} + 0.277\dot{m}_{HC} \tag{4.24}$$

The total quantity of PM created, along with total NO_x, HC and CO form the emissions constraints:

$$\int_0^{s_f} \omega_e(s)\tau_e(s)\ \mathrm{BSNO_x}(\omega_e(s),\tau_e(s))\frac{\mathrm{d}s}{v(s)} \leq \overline{m}_{\mathrm{NO_x}} \tag{4.25}$$

$$\int_0^{s_f} \omega_e(s)\tau_e(s)\ \mathrm{BSHC}(\omega_e(s),\tau_e(s))\frac{\mathrm{d}s}{v(s)} \leq \overline{m}_{HC} \tag{4.26}$$

$$\int_0^{s_f} \omega_e(s)\tau_e(s)\ \mathrm{BSCO}(\omega_e(s),\tau_e(s))\frac{\mathrm{d}s}{v(s)} \leq \overline{m}_{CO} \tag{4.27}$$

$$\begin{aligned}\int_0^{s_f} \omega_e(s)\tau_e(s)\ \Big(&1.024\mathrm{BSSoot}(\omega_e(s),\tau_e(s)) \\ &+0.277\mathrm{BSHC}(\omega_e(s),\tau_e(s))\Big)\frac{\mathrm{d}s}{v(s)} \leq \overline{m}_{PM}.\end{aligned} \tag{4.28}$$

In short, the whole optimal control problem for the CVT optimisation is as follows: minimise the following integral

$$\int_0^{s_f} \dot{m}_f(\omega_e(s),\tau_e(s)) + r_1\dot{P}_t(s)^2 + r_2\dot{\omega}_e(s)^2 + r_3\dot{\tau}_e(s)^2\frac{\mathrm{d}s}{v(s)}. \tag{4.29}$$

The state vector is bounded in the following way

$$\begin{bmatrix} P_b \\ \epsilon \\ \omega_e^{idle} \\ 0 \end{bmatrix} \leq \mathbf{x} = \begin{bmatrix} P_t \\ v \\ \omega_e \\ \tau_e \end{bmatrix} \leq \begin{bmatrix} P_e^{max} \\ v_m \\ \omega_e^{max} \\ \tau_e^{max}(\omega_e) \end{bmatrix}. \tag{4.30}$$

The corresponding state dynamics are P_t, (4.1),ω_e and τ_e. The control inputs are

$$\begin{bmatrix} \dot{P}_t \\ \dot{\omega}_e \\ \dot{\tau}_e \end{bmatrix} \tag{4.31}$$

which is effectively unbounded. The set of inequality constraints is

$$\mathbf{u} = \begin{bmatrix} 0 \\ -1 \\ -1 \end{bmatrix} \leq \begin{bmatrix} \eta_{CVT}\omega_e(s)\tau_e(s) - P_t(s) \\ \frac{a_x(s)}{a_x^{max}} + \frac{a_y(s)}{a_y^{max}} \\ \frac{a_x(s)}{a_x^{max}} - \frac{a_y(s)}{a_y^{max}} \end{bmatrix} \leq \begin{bmatrix} \infty \\ +1 \\ +1 \end{bmatrix}. \tag{4.32}$$

When the AMT powertrain configuration is used, the engine speed and engine speed rate control and state are removed, and an additional parameter set is optimised (namely those constant shifting speeds in (4.13)): The cost function reads:

$$\int_0^{s_f} \dot{m}_f(\omega_e(s), \tau_e(s)) + r_4\dot{P}_t(s)^2 + r_5\dot{\tau}_e(s)^2 \frac{\mathrm{d}s}{v(s)}. \tag{4.33}$$

The state vector is bounded in the following way

$$\begin{bmatrix} P_b \\ \epsilon \\ 0 \end{bmatrix} \leq \mathbf{x} = \begin{bmatrix} P_t \\ v \\ \tau_e \end{bmatrix} \leq \begin{bmatrix} P_e^{max} \\ v_m \\ \tau_e^{max}(\omega_e) \end{bmatrix}. \tag{4.34}$$

The corresponding state dynamics are P_t, (4.1) and τ_e. The control inputs are

$$\mathbf{u} = \begin{bmatrix} \dot{P}_t \\ \dot{\tau}_e \end{bmatrix} \tag{4.35}$$

which is effectively unbounded. The static parameter set is

$$\mathbf{p} = \begin{bmatrix} u_1 \\ u_2 \\ u_3 \\ u_4 \end{bmatrix} \tag{4.36}$$

The following constraints must be enforced to ensure gears are shifted in the correct order, but $\mathbf{p}$ is unconstrained otherwise

$$u_1 < u_2, \quad u_2 < u_3, \quad u_3 < u_4. \tag{4.37}$$

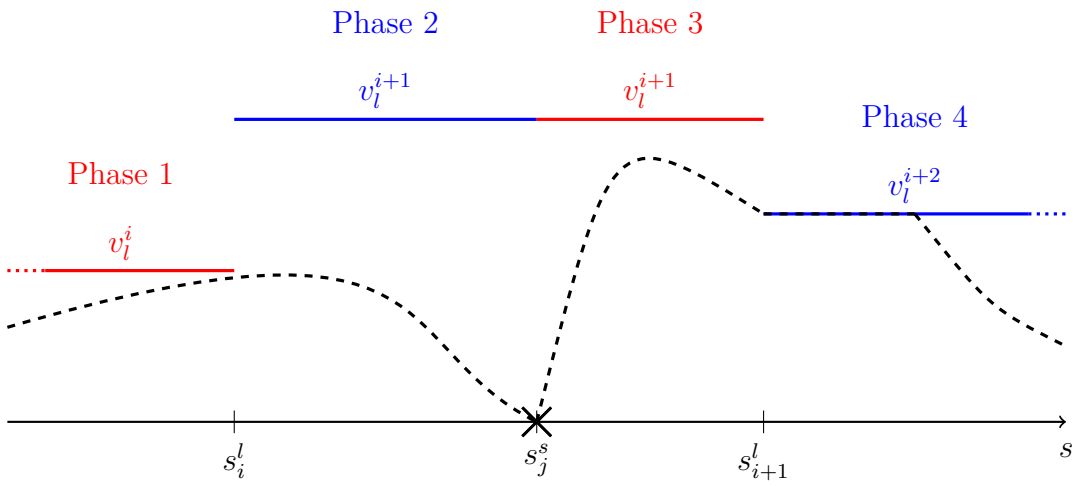


Figure 4.8: A diagram of how the distance based speed limit information is shown here. Each speed limit is separated by a boundary where either the speed is continuous or required to be zero. An admissible speed trajectory is shown as the dashed black curve.

The set of inequality constraints is

$$\mathbf{u} = \begin{bmatrix} 0 \\ -1 \\ -1 \end{bmatrix} \leq \begin{bmatrix} \eta_{AMT}\omega_e(s)\tau_e(s) - P_t(s) \\ \frac{a_x(s)}{a_x^{max}} + \frac{a_y(s)}{a_y^{max}} \\ \frac{a_x(s)}{a_x^{max}} - \frac{a_y(s)}{a_y^{max}} \end{bmatrix} \leq \begin{bmatrix} \infty \\ +1 \\ +1 \end{bmatrix}. \tag{4.38}$$

The complete set of integral constraints (when imposed) is formed from (4.23) and (4.25)–(4.28). Problem setup parameters are tabulated in Table 4.1. The regularisation parameters r_1, r_2, r_3, r_4 and r_5 are chosen to avoid a large regularisation part of the cost. All of these r_i parameters of the order 10^{-6} to 10^{-3}.

4.2.1 Multi–phase Numerical Optimal Control

The OCP of the previous section is solved numerically with a direct transcription method based on Legendre-Gauss-Radau (LGR) collocation and Radau's integration formula [39].

In this study the OCP is discretized and solved using GPOPS-II [39], which uses *hp*-adaptive mesh refinement [51]. The large, sparse NLP is solved with the interior point algorithm such as IPOPT [35]. In order to solve the problem effectively in the presence of discrete changes in the speed limit, the OCP is broken down into multiple phases. An example multi-phase velocity state trajectory is plotted in Figure 4.8. Phases 1 and 2, and, 3 and 4 have differing speed limits and require

Table 4.2: Route A speed limit and stops information.

(a) Speed limit and distance information.

s [km]	v_l^p [mph]
$[0, 5.97]$	30
$[5.97, 11.4]$	50
$[11.4, 11.53]$	30
$[11.53, 12.36]$	50
$[12.36, 16.58]$	70
$[16.58, 20.08]$	50
$[20.08, 23.62]$	70
$[23.62, 24.42]$	30
$[24.42, 25]$	40
$[25, 26.11]$	30
$[26.11, 26.53]$	20
$[26.53, 28.12]$	30

(b) Stopping instances.

s_s [km]	
0.57	12.2
1.84	13.43
2.18	13.97
2.93	23.84
4.01	24.35
5.66	24.91
8.81	26.59
11.44	

continuity of the state at the boundary between them. Phases 2 and 3 are separated by a stop and have the same speed limit. This stopping event can be seen as a splitting of an already existing phase into two.

At the stopping points $s = s_s$, the condition

$$v(s_s) = \epsilon, \tag{4.39}$$

is enforced. The parameter ϵ is a small positive number to ensure the vehicle can come to rest, but avoids the problem breaking down if $v(s) = 0$.

The boundary conditions on the velocity at the beginning and end of the route are

$$v(0) = v(s_f) = \epsilon. \tag{4.40}$$

The legislative speed limits for Route A, are listed in Table 4.2a. The fifteen potential enforced stopping points are listed in Table 4.2b. This leads to a total of twenty-two potential phases.

4.2.2 Discontinuous Functions

Discontinuous functions pose a difficulty for the solution of the NLP due to the loss of gradient information. One way to overcome this problem is to break it down into multiple phases—distance-dependent speed limits are an example of relevance here.

Another method is to 'smooth' the problem by replacing discontinuous functions with smooth approximations using the methods detailed in Chapter 2. For example, the $\max(x, 0)$ function in (4.15) has an undefined gradient at $x = 0$. In this chapter, the hyperbolic tangent function is chosen to approximate the $\max(x, 0)$ function.

4.3 Results

A number of fuel minimization results that represent 'real' driving situations will be presented. These will begin with a comparison of gearbox types and the optimization of the gearbox shifting speeds. The sensitivity of the solution to changes in emissions constraints will then be studied. Finally the influence of enforced on-route stops will be analysed. The solutions use optimized driver and engine controls.

4.3.1 Power Transmission Comparison

In the first study, fuel consumption is minimized for a maximum arrival time of $T = 1800\,\text{s}$. The emissions constraints given in (4.25)–(4.28) are omitted. This problem is solved for both CVT- and AMT-equipped vehicles. In the CVT case the engine speed and torque are optimised. In the AMT case the engine torque and gearbox shifting speeds (u_i in (4.13)) are optimised.

The optimal speed and tractive power for each transmission is plotted in Figure 4.9. Engine power, and thus fuel, is required to ascend the hills, accelerate out of corners, and increase the vehicle's speed when the legislated speed limit increases. Braking is needed to slow the car to negotiate corners at 7km and 25km for example. Noticeable is the minimal use of the highest speed limits, which are avoided due to the low-speed turns at the end of each segment.

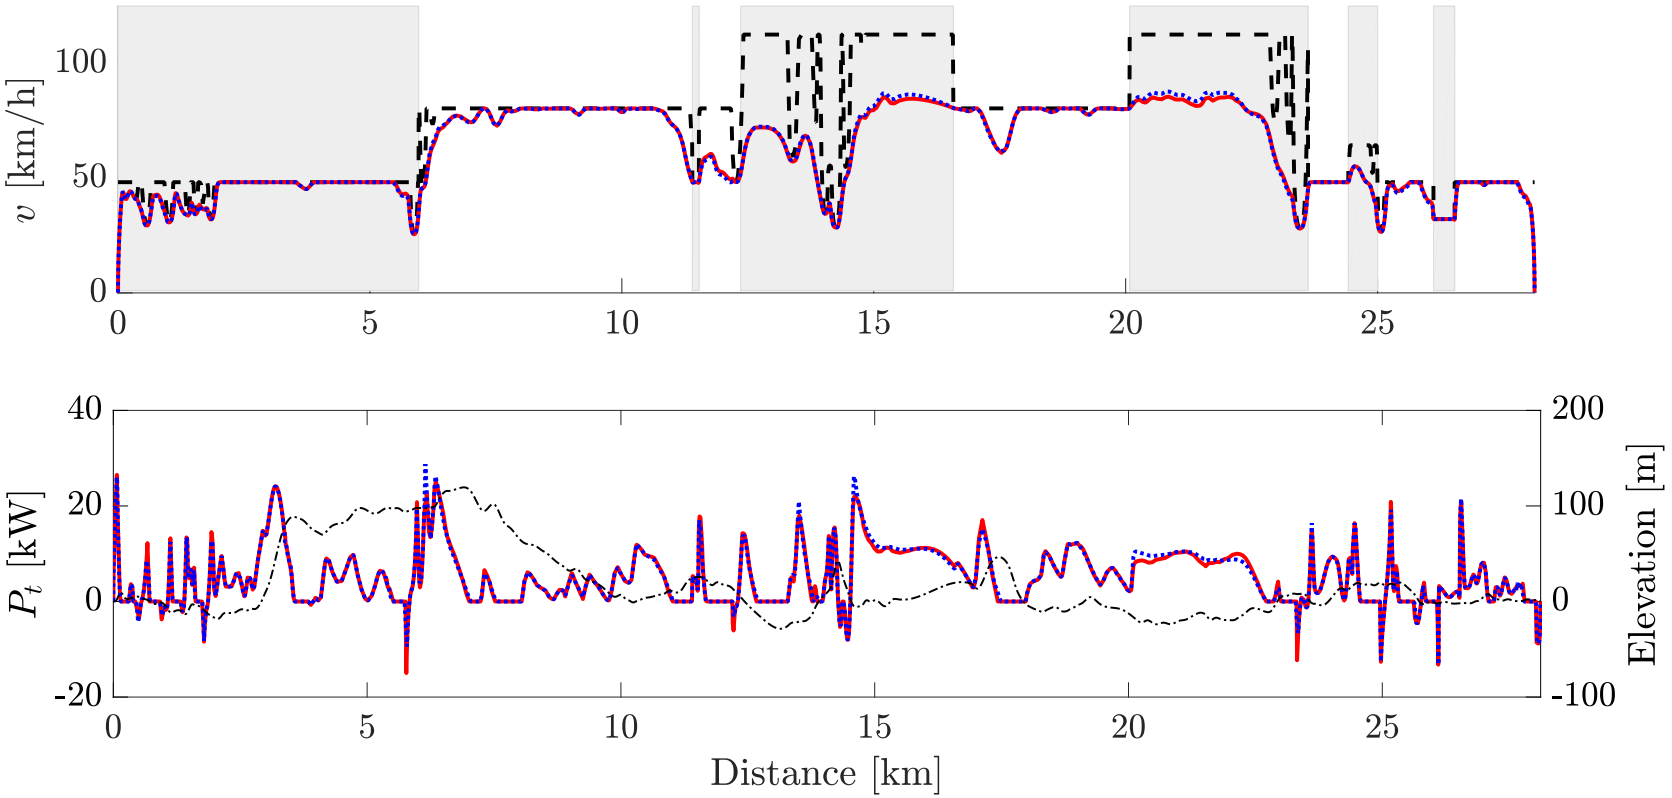


Figure 4.9: Optimal velocity and power profiles for CVT (blue dotted) and AMT (red solid) transmission systems. Combined legislative and curvature-related speed limits are plotted as a dashed black line on the velocity figure. The route elevation is plotted as a dot-dash curve on the power figure. Individual phases are represented by alternating light grey regions.

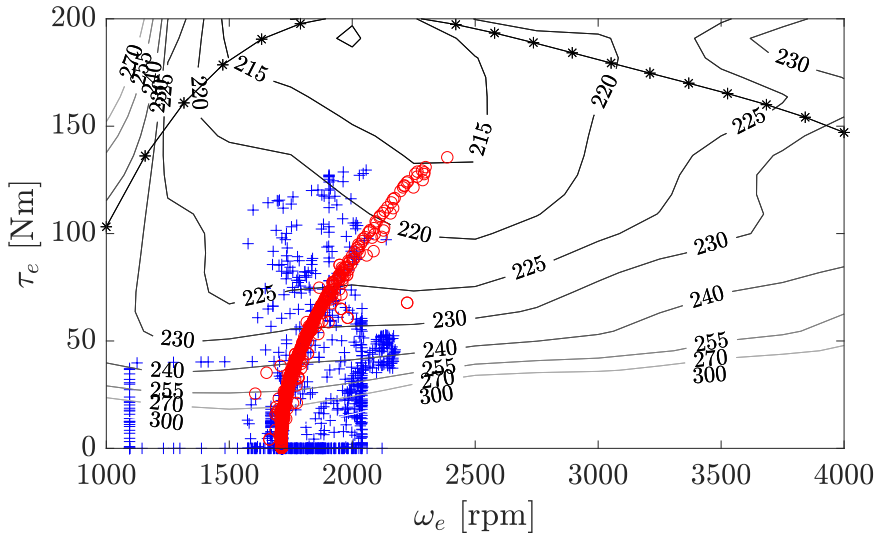
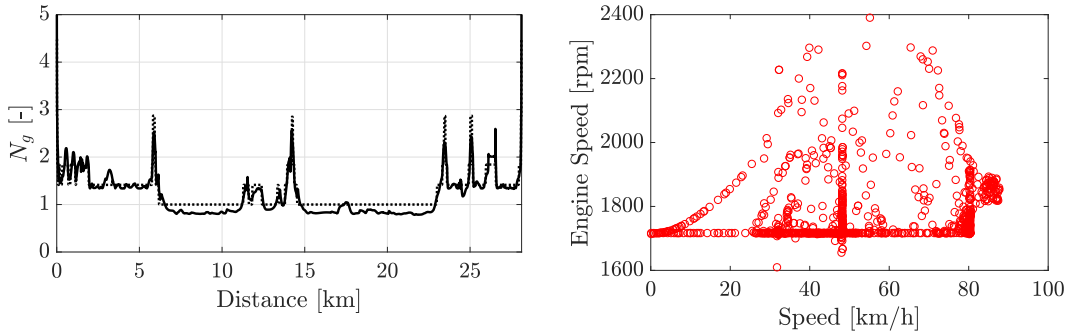


Figure 4.10: Scatter of optimal engine operating points on the BSFC [g/kWh] map. The CVT results are plotted as circles with the AMT plotted as crosses.

In contrast to the example given in Chapter 3, the maximum available power is not utilised on the route. Constraints on the maximum longitudinal acceleration and the relatively low speed limits obviate the need for maximum power.

Optimal engine operating points are compared in Figure 4.10. This clearly highlights the difference in the engine usage for the different transmission types. In the case of a CVT, the line associated with minimum fuelling are shown as circular scattered points. For the AMT, the engine speed is constrained by the gear selected and the forward speed of the vehicle. This leaves the torque as the only free variable—this results in a more scattered set of engine operating points; see the crosses in Figure 4.10. The engine operating region is determined by the (optimized) gearbox switching speeds. The total FC for the CVT is 657 g compared to the AMT which used a total of 601 g. Despite the CVT operating along the locus of highest fuel efficiency, the increased mechanical efficiency of the AMT results in superior fuel economy.

Figure 4.11a compares the optimized gearing for each transmission. At the beginning and end of the route the CVT ratio is considerable higher that the AMT. This gives rise to increased acceleration for the CVT equipped vehicle. Also, the improved gearing offered by the CVT at higher speeds, ensures that lower



(a) Gearbox ratios for CVT (solid) and the AMT (dotted).

(b) The mapping between the vehicle forward speed and the engine speed for the CVT case.

Figure 4.11: Optimisation gear ratio results

engine speeds, in combination with higher torques, can be used to reduce the fuel consumption. The CVT mapping between the vehicle forward speed and the engine speed is plotted in Figure 4.11b. This plot highlights the lack of one-to-one mapping which is present for the AMT equipped vehicle.

The regularisation cost accounts for approximately 0.05% of the total cost in both cases. The computation time for computation on an 8 core 3.5GHz desktop PC for the CVT case is approximately 550 s, whereas for the AMT, this increases to approximately 2800 s. Additional processing time is needed due to the discontinuous nature of the engine idle speed constraint, and optimising the gear-speed shifting parameters given in (4.13). Both cases are initialised using the same mesh, consisting of 10 mesh intervals per phase, each with 2 collocation points. The initial guess for each is simply the midpoint of each state and control. The CVT problem is solved with 21 mesh iterations, increasing the total number of mesh intervals to 186 (and 1054 collocation points). Coincidentally, the AMT problem also requires 21 mesh iterations, but ultimately exits with a total of mesh 213 intervals with 1205 collocation points. While the solution of the optimal control problem is simplified by neglecting the discrete nature of the gearbox, most road vehicles are equipped with discrete gearboxes. As such the resulting engine trajectories from the CVT optimisation are unrepresentative of reality. For the remainder of this study the vehicle is considered equipped with a discrete AMT, which makes use of the optimised gear-change speeds found here.

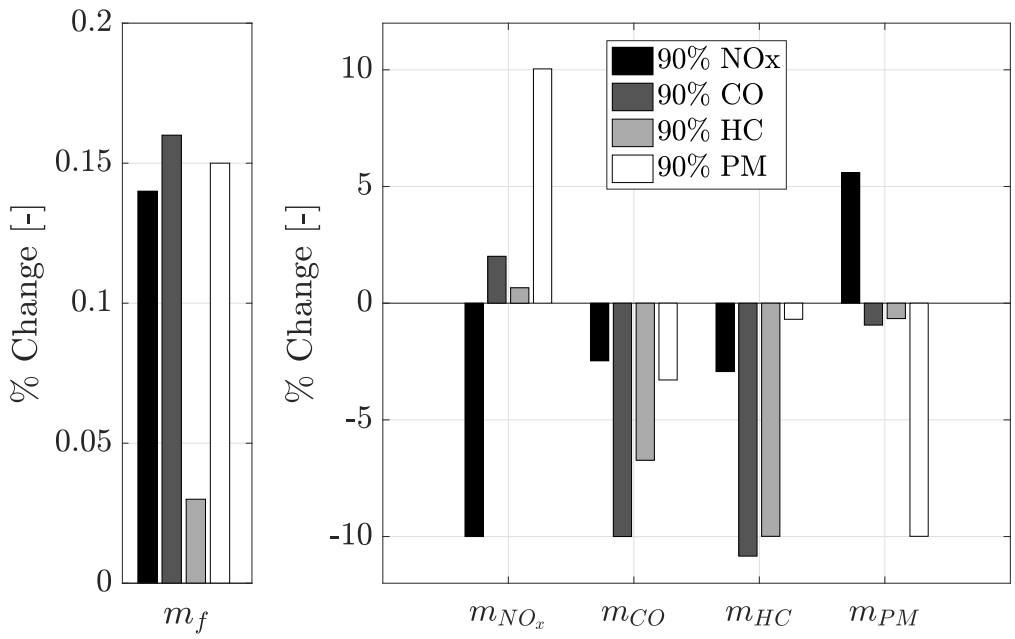


Figure 4.12: Total percentage changes in each of: fuel, NO_x CO HC and PM for each of the constrained optimisations.

4.3.2 Emissions Sensitivity

This section considers the impact of enforcing the constraints on the emissions generated en route (see inequalities (4.25)–(4.28)). The gear-shift speed-change law is fixed to the minimum fuel solution found in the previous section. This helps reduce the problem complexity and increase the solution speed. In this study each emissions type is constrained in turn to evaluate the sensitivity of the FC and optimal driving strategy for each case. The optimisation is initialised with the solution to the unconstrained solution of the previous section.

Figure 4.12 shows the effect of limiting the emissions components one-by-one to 90% of their unconstrained values. It is important to note that further restrictions on the emissions are likely to make the optimisation problem infeasible in some cases. With the 'tight' arrival time constraint being imposed, there is little scope for altering the power requirements. Consequentially, it is expected that only small changes in FC are achievable (also shown in Figure 4.12).

The first optimisation considers a limit of 90% of the unconstrained NO_x (black

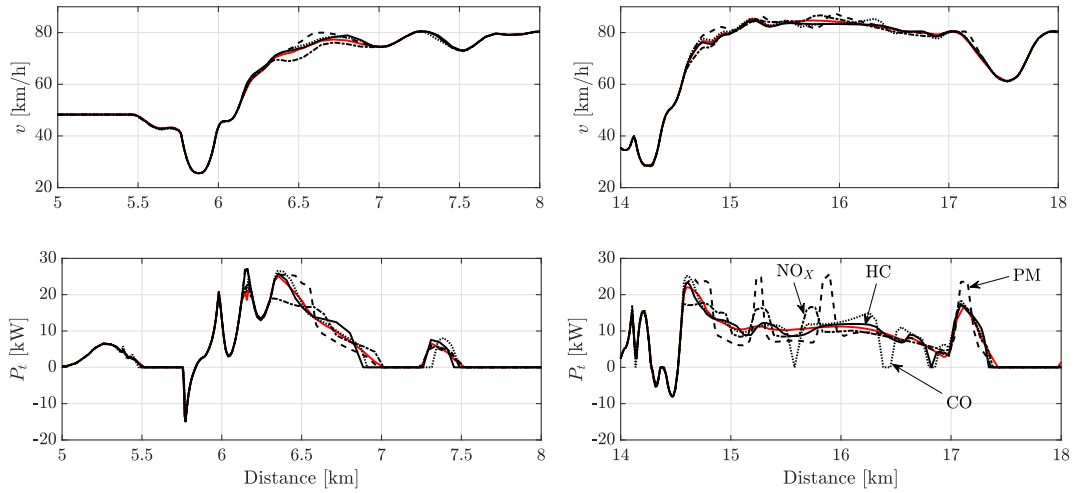


Figure 4.13: Segments of the emissions-constrained optimal driving solution for the 5km—8km section, and the highway driving section 14km—18km. The baseline fuel-optimal solution is plotted in red. The 90% constraints cases on NO_x CO, HC and PM are shown, respectively, as dot-dash, dotted, solid, and dashed.

bars in Figure 4.12). A negligible increase in the FC is noted (as it is for the other cases in this section). The total CO and HC decreases by 2.5% and 2.9% respectively, while the PM increases by 5.6%. The increase in PM follows from the fact that the BS-NO_x and BS-soot maps have complementary minimum regions (see Figures 4.4 and 4.5), and the generation of PM is dominated by the soot produced (4.24).

In the second optimisation case, a limit of 90% of the unconstrained CO is enforced (dark grey in Figure 4.12). Total NO_x is increased by 2.0%, while the PM decreases by 0.9%, and HC decreases by 10.8%. The coincidence of the regions of low BS-HC and BS-CO accounts for the large decrease in HC seen here.

Imposing a limit on the total HC (light grey in Figure 4.12) has the least effect on the FC (an increase of only 0.03%) implying that the region in which the engine is operated changes very little. This small difference results in changes of -0.6% in the total PM produced, and a change in the NO_x produced of +0.6%. This slight shift in engine operating point is enough to reduce the CO by 6.7%, which is consistent with the limited CO case.

The final case constrains the PM to 90% of that obtained in Section 4.3.1 (the white bars in Figure 4.12). The large 10.0% increase in the NO_x is expected

as discussed earlier. Abatements of 3.3% in CO, and 0.7% in HC, are also observed, which correlates with a move in engine operation to the higher power areas in the BS-maps.

In the results presented so far, behavioural changes in the route emissions have been explained by examination of BS- emissions maps. To complete the analysis it is necessary to consider not only the regions in which the engine is operated, but also how the power is utilised in these regions during driving. A comparison of the effects of the emissions constraints on driving, in terms of optimal tractive power and speed profiles, for each constrained optimisation are shown in Figure 4.13. In order to reduce the total NO_x output, the tractive power is kept below 20 kW. This reinforces the idea of avoiding the highest regions of BS-NO_x that occur at high power. To ensure that the speed is increased during highway driving, power is applied in small bursts in the 15 km to 16 km stretch.

When the PM constraint is introduced the application of power changes significantly. As with the restricted NO_x case, the optimal tractive power includes rapid changes between high and low engine powers. High power engine operation is beneficial to soot reduction, but the speed limits have to be adhered to.

During the 6.25 km to 7 km period of driving, the opposing requirements of the PM and NO_x solutions can be observed. The application of high power is followed by a change to lower power in the PM constrained case. Whereas a longer and smoother application of power occurs in the NO_x limited case. Despite the limited HC and CO being similar in terms of their total outputs, the method of achieving this outcome requires different driving strategies. The limited HC case requires a driving strategy which is almost identical to the baseline. This corresponds to the negligible change in FC (see Figure 4.12). The CO-constrained case includes periods of zero engine power during which no emissions are produced. In this case it is preferable to coast for periods, rather than drive the engine hard as in the PM-limited case.

The results in this section illustrate the need to optimise simultaneously the driving strategy, and the engine operation to abate emissions production. Further reductions may be possible by relaxing the arrival time constraint, thereby allowing

Table 4.3: The stops used in each of the comparison runs.

	Stop No.															
Run	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Total Stops
1	*		*							*	*	*	*	*		7
2						*	*		*			*		*		5
3	*		*			*	*	*		*	*	*		*	*	10
4					*		*	*	*			*	*		*	7
5				*					*	*				*		4
6	*	*	*	*	*				*	*			*	*	*	10
7			*		*	*		*		*	*	*	*	*		9
8		*		*		*		*								4
9	*		*	*	*	*					*		*		*	8
10	*		*	*		*		*		*		*		*		8
11	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	15

for more varied power and velocity profiles. Due to the opposing means of production of emissions types, combined constraints are infeasible to implement.

4.3.3 Enforced Stops

One can expect to encounter enforced stops along the route at intersections and give ways, which will impact on the FC and the production of emissions. These additional boundary conditions can be included in the problem by dividing the route into additional phases. In the stop-start study case adjacent phases may have the same speed limit, but are necessarily separated by the zero velocity boundary condition (4.39).

The potential stopping places around the Oxford route are listed in Table 4.2b (the start–finish point is not counted). Traffic conditions, and accidents of timing, mean that there is a degree of uncertainty associated with negotiating an urban road system. To investigate these influences on the FC and emissions, eleven different sets of stops will now be studied as shown in Table 4.3. Each set of stops is generated randomly and is known to the optimal control problem beforehand.

Enforced stops can influence the achievable route transit time and consequently the arrival time constraint must thus be relaxed to ensure that the route can be covered. The minimum arrival time for the route in the case that each enforced

stop is respected is approximately 1850 seconds. The arrival time constraint of the previous example is therefore relaxed to 2000 s (rather than the previous 1800 s).

Once again only the AMT will be considered as it reflects the more realistic driving scenario. As in the non-stopping unconstrained case with an AMT, the arrival time constraint is active, but increased to 2000 s. Clearly, even if the number of stops is increased, a minimum average velocity must be maintained. This results in increased tractive power effort and increased fuel consumption.

The car must be accelerated more frequently to counteract the increased number of braking phases. In the case that there are no stops, the car does not reach the speed limit between 6km and 11.5km, and between 14km and 23km. Power is applied more smoothly and rapid acceleration can be avoided. The results of three optimisations are presented in Figure 4.14.

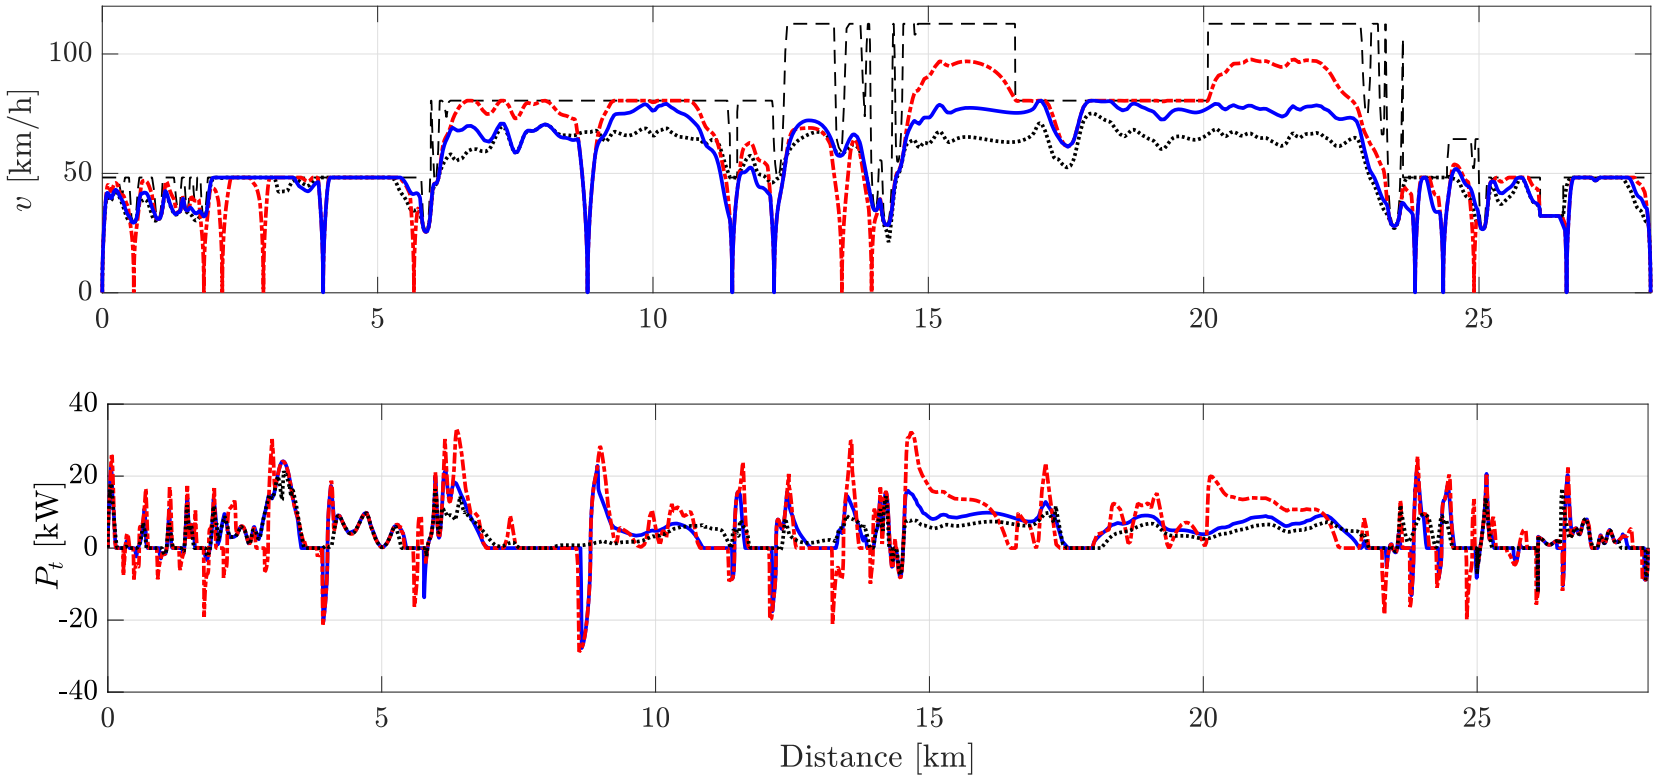


Figure 4.14: Optimal velocity and power profiles for zero stops (black dotted), 7 random stops - Run 4 in Table 4.3 (blue solid) and all stops (red dot-dashed). Velocity limits are plotted as dashed black lines.

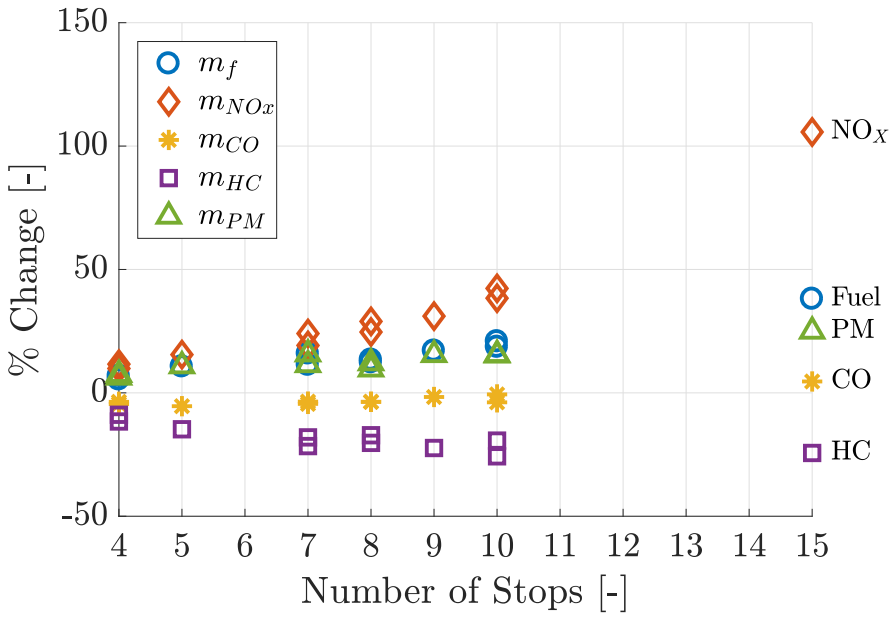


Figure 4.15: Percentage changes in the total masses of fuel, NO_x, CO, HC and PM relative to driving without stopping.

The effect of including stopping events on the total route emissions and fuel consumption is shown in Figure 4.15. The inclusion of stopping events necessarily increases the number of braking events, as well as the need for higher peak speeds. This, in turn, leads to greater fuelling requirements due to increased numbers of prolonged applications of power. Increased power usage results in reduced HC—higher combustion temperatures remove some HC products. Total NO_x increases with the number of stops, this is once again linked to the higher engine power usage.

Furthermore, in order to effectively increase the speed of the vehicle, while remaining in the fuel efficient region of Figure 4.3, then the shifting speeds must change (see Figure 4.16). The optimisation is once again allowed to select the gear shift speeds in these cases. The lowest shifting speed, between first and second gear, is reduced in all stopping cases. This ensures that the engine operates within the optimal fuelling region. The highest shifting speed is increased as the number of stops increases. The requirement for higher peak speeds leads to this result.

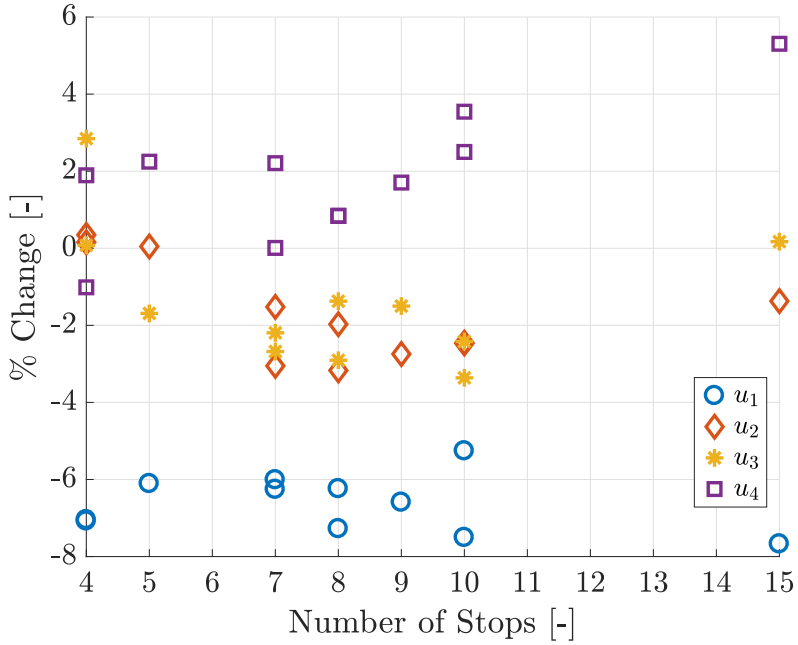


Figure 4.16: Percentage changes in the shifting speeds relative to driving without stopping.

4.4 Conclusions

This chapter details an effective method for utilising route information in combination with state of the art engine modelling techniques to simultaneously optimise engine usage and driving style. Building on the knowledge of Chapter 3, driving cycle data can be created based on a prototypical vehicle model with representative powertrain characteristics.

The vehicle is driven on a 3D route, based on information obtained from GPS. Gaussian process models for engine flow rate quantities form a set of integral constraints. Models of this form have the advantage of being accurate under interpolation and are good at capturing non-linear behaviours. A multi-phase optimal control problem is posed and solved, which represents real world driving situations for vehicles equipped with either continuously variable, or automated manual transmissions. Driving style changes are evident when new constraints are posed on the system.

The methodology presented here is to overcome some of the poor cycle predictions

gained from fixed drive cycle testing which lacks the flexibility to properly optimise FC and emissions predictions. The results of using fixed drive cycles could result in poor engine calibration. While there is little flexibility for changing the total fuel consumption when the arrival time is tight, the production of emissions can change significantly with driving style variations. If the arrival-time constraint is relaxed, greater flexibility in driving strategy can reduce emissions still further. It is not sufficient to ignore on-route stopping, because stopping causes a significant increase in FC, NO_x, CO, and PM. Validation of the resulting velocity profiles is presented in Appendix C.

The complex shape of the engine emissions maps, and the discrete nature of the gearbox, makes these studies computationally demanding. While this work is not directed specifically at real-time applications, it does have a role in providing 'what's possible' benchmarking information in the development of real-time schemes such as receding horizon controllers. Powertrain configurations, which reflect new technologies, will be the subject of further work. The next chapter considers the ability for flexible power delivery from a hybrid powertrain to effectively heat a catalytic converter. An optimal control problem of similar structure to that in this chapter is presented and solved.

. . . the Stone Age did not end because we ran out of stones; we transitioned to better solutions. The same opportunity lies before us with energy efficiency and clean energy.

— Steven Chu, In letter to Energy Department employees, 1st February 2013

5

Optimal Hybrid Vehicle Driving and Aftertreatment Warmup

Contents

5.1 Model 120
5.1.1 Vehicle Model 120
5.1.2 Powertrain Model 122
5.1.3 Engine Model 125
5.1.4 Energy Store 129
5.1.5 Three-way Catalytic Converter 130
5.1.6 Route Data 131
5.1.7 Optimal Control Problem Setup 132
5.2 Results 134
5.2.1 Baseline Study: Fuel-Optimal 135
5.2.2 Case Study 1- Emissions Constraints 141
5.2.3 Case Study 2 - Route B2 148
5.3 Conclusions 153

CAPABILITY for a hybrid powertrain to improve fuel economy has been studied extensively. Regenerative braking recovers energy that would otherwise have been dissipated with friction brakes via conversion of kinetic to potential energy. Furthermore, the fuel economy of the ICE can be improved by operating it in its most fuel efficient manner. Typically, this involves the use of an electrical machine working in conjunction with a battery. This stored energy can then be

redeployed replacing the burning of further fuel.

Rule-based architectures for power-split control rely on estimation of the state of charge of the energy store and make use of instantaneously available power demands. By utilising future information, an optimisation minimising fuel consumption can be formed. For real-time applications, the problem must remain simple in terms of the number of dynamic states and the model itself.

In conventional power-splitting control, a single dynamic state relating to the state-of-charge (SOC) of the energy store (ES) is required [18, 77]. This is then optimised over a fixed drive cycle, which features as a power demand which must be met by the powertrain. This drive cycle is known *a priori*. Analytical powersplitting results, based on Pontryagin's Maximum Principle (PMP), can also be derived [58, 61]. This can incorporate optimisation of the velocity and the powertrain operation based on simple road information such as constant velocity limits and constant road grade. Increasing the complexity of the model makes solutions more cumbersome to compute (for example, by including a variable road grade). Utilising these results, online controllers can be derived called Equivalent consumption minimisation strategies (ECMS) [18]. This methodology employs the solution of the powersplit problem with PMP for the single dynamic SOC. The costate associated with this dynamic state can be computed explicitly and solved for online. A considerable amount of further research has further developed ECMS, for example, to include route emissions [19].

Numerical optimisation via dynamic programming (DP) has been studied extensively and remains a standard for solving the power split problem due to its simplicity [78]. One key issue is the inclusion of switched dynamics relating to engine on/off operation. This requires the use of a discrete decision variable. A problem formulation of this type can be solved with mixed integer nonlinear programming [79], but again the problem can only be solved over a fixed power and velocity profile.

While this is suitable for small scale problems, in the case of additional constraints and an increased number of states, solution time becomes a hindrance. Inclusion of the temperature dynamics of the TWC increases model order and relies on

simplifications to the model in order to maintain tractability. Optimising the power splitting in a mild-hybrid powertrain was the subject of [80]. A high fidelity model of a TWC was combined with a gasoline engine to evaluate *smart-heating* strategies. In charge-sustaining scenarios, this showed promise for large reductions in emissions over conventional power-splitting control. The warmup of the engine itself can be optimised with a hybrid powertrain to help alleviate inefficiencies due to cold lubricants [81]. Better power split optimisation can be achieved by including the power required for cooling [82].

Most of these studies rely on computation of the solution offline. Despite the fact that the optimisation routines are close to real-time, either the solution is far from ideally computed via model simplifications or the computation time is slow.

One of the key drawbacks for offline powersplit control is the limitations enforced by the solution method. An increased number of dynamic states with terminal constraints becomes infeasible for solution via DP. Equally, analytical solutions based on PMP are similarly limited - this time by the complexity of any non-linear features. If the cases considered in Chapter 3 were extended to include non-linear roads and a complex powertrain model, solutions will be infeasible to compute analytically.

This chapter considers the combined problem of optimal driving and powertrain control for a parallel-series equipped vehicle and the solution of the problem via direct collocation.

A simple thermal model for the conversion efficiency and dynamic thermal characteristics enables the appraisal of different heating strategies which is crucial to curtail harmful emissions. Additionally, no fixed control laws are applied which govern engine on/off. This is intended to ensure that the optimiser has complete control over the powertrain usage.

It is shown that, depending on the route, different, combined power and velocity trajectories will emerge for the minimisation of fuel consumption with emissions constraints. This chapter initially presents a model suitable for offline dynamic optimal control. The optimal control problem is formed and solved for a number of different scenarios and the results discussed. Results of this chapter

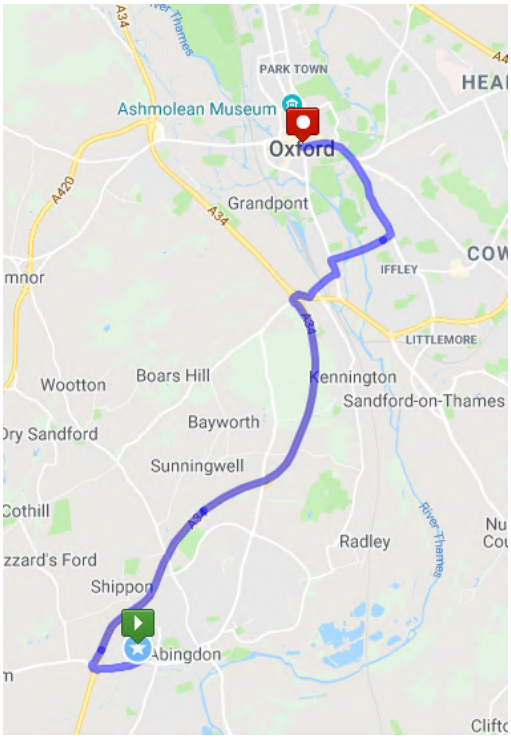


Figure 5.1: Driving route for this chapter, denoted Route B.

are in preparation for publication. The route that the vehicle is driven along is plotted in Figure 5.1.

5.1 Model

Firstly, a model for the vehicle and powertrain is formed. In order to effectively use numerical optimal control to compute the solution to the powersplit plus driving problem, all required component models need to be smooth and differentiable.

5.1.1 Vehicle Model

At the timescales considered, dynamics relating to, weight transfer, driveline inertia and tyre properties can be neglected. The vehicle model follows that in Chapter 4 and is repeated here for completeness. In order to properly utilise the route information regarding speed limits, curvature and road-grade the vehicle model is posed once again in the elapsed distance domain, denoted s. This transformation is allowable provided the vehicle speed condition, $v(s) > 0$, holds. All of the

time-based derivatives can be rewritten using the following:

$$\frac{d}{ds} = \frac{1}{v(s)}\frac{d}{dt} \tag{5.1}$$

The car's longitudinal acceleration is governed by:

$$\dot{v}(s) = a_x(s) = \frac{1}{Mv(s)}\left(\frac{P_t(s)}{v(s)} - F_r(s)\right) \tag{5.2}$$

where M is the vehicle mass. The input variable $P_t(s)$ is the tractive power. The resistive force $F_r(s)$ is made up of three components: aerodynamic drag, the rolling resistance and the component of the vehicle's weight parallel to the slope

$$F_r(s) = Mg(\mu \cos\theta(s) + \sin\theta(s)) + \frac{1}{2}\rho C_d A v(s)^2. \tag{5.3}$$

The distance dependent local road gradient is $\theta(s)$. The rolling resistance is denoted μ, A is the car frontal area, the drag coefficient is C_d, ρ is the density of air and g is the acceleration due to gravity.

The lateral acceleration allows for inclusion of the road curvature $\mathcal{C}(s)$ and is given by:

$$a_y(s) = \mathcal{C}(s)v(s)^2. \tag{5.4}$$

This expression along with (5.2), forms the feasible space for normalised vehicle accelerations

$$\left|\frac{a_x(s)}{a_x^{max}}\right| + \left|\frac{a_y(s)}{a_y^{max}}\right| \leq 1 \tag{5.5}$$

This region of allowable accelerations limits the optimisation (which represents our driver) to remain within comfort bounds, below the limits of adhesion from the tyres. Experimentally, regions of this type have been studied extensively for racing drivers (known as g-g diagrams [36]). For representation of normal road driving, the region shape is shrunk as drivers are more cautious when cornering and braking/accelerating [66].

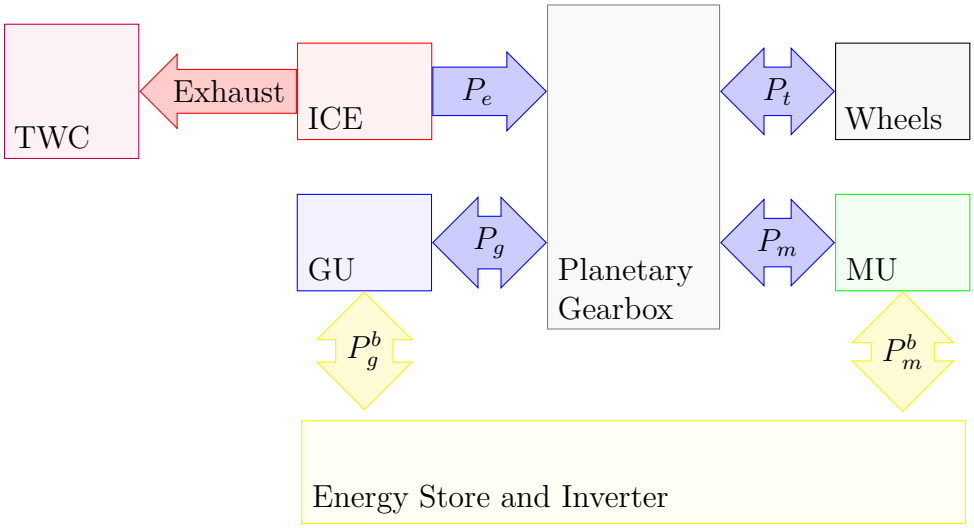


Figure 5.2: A model diagram of the available power flows in the hybrid powertrain under consideration. Mechanical power links are blue arrows, electrical links are yellow arrows.

5.1.2 Powertrain Model

A parallel-series hybrid powertrain comprises three power sources: an internal combustion engine (ICE); a large motor unit (MU) and a smaller motor-generator unit (GU). The layout of which is shown in Figure 5.2.

Power flows are blue arrows. The exhaust gas flow is red, and electrical power flows are yellow. Power is allowed to flow in either direction through the MU and GU from the energy store (ES) via an inverter. The ES in this study is a NiMH battery; however, other devices such as ultracapacitors and flywheels can be used [70].

The three power generating devices are configured in an *epicyclic gearing* arrangement, sometimes referred to as a planetary gearbox. In this configuration, each of the devices lies on parallel shafts. The sun gear, sits in the center and is orbited by several planet gears attached to a carrier. This set of gears is surrounded by an internal ring gear which encloses the whole device. The GU shaft sits inside the ICE shaft which is in turn connected to the planet carrier (red in Figure 5.3). These enter the device from the left hand side, and the MU shaft leaves the device on the right hand side. The MU is attached to the ring gear (shown in green in Figure 5.3). Finally, the GU is attached to the sun gear in the centre of the configuration. As the MU is coupled to the wheels, its rotational speed ($\omega_m(s)$), is fixed by the forward

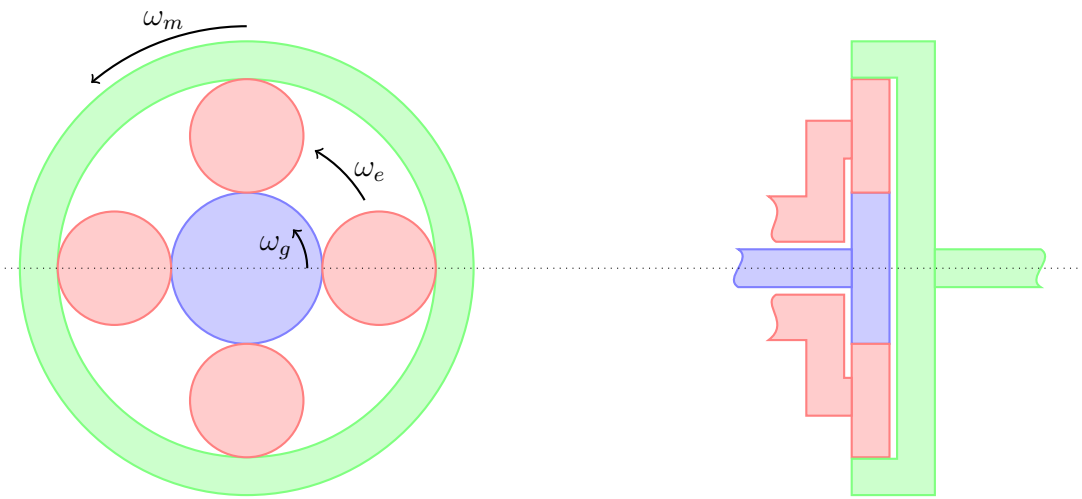


Figure 5.3: The layout of the electrical machines and internal combustion engine gearing. The ICE is coupled to the planetary carrier and planet gears (shown in red), the MU is attached to the ring gear (shown in green) and the GU is attached to the sun gear (shown in blue).

velocity ($v(s)$) of the car: $\omega_m(s) = v(s)\frac{N_f}{r_w}$ (where N_f and r_w are the final drive ratio and the wheel radius). In this arrangement, a smoothly varying gear ratio can be realised between the engine speed $\omega_e(s)$ and the wheels - behaving like a CVT. The ratio of which is determined by the GU rotational speed $\omega_g(s)$. A complete analysis of epicyclic gearboxes of this nature can be found in texts such as [83, p. 248].

In short the rotational speeds are governed by:

$$\omega_g(s) + \frac{N_r}{N_s}\omega_m(s) = \frac{(N_r + N_s)}{N_s}\omega_e(s). \tag{5.6}$$

N_s and N_r, are the number of teeth on the sun and the ring gear [1].

The limits on speeds on the electrical machines and engine speed, in combination with (5.6), creates the operating space in Figure 5.4.

The blue region in the diagram indicates the *all-electric* operating region which is determined by the maximum negative rotational speed of the generator. The lever in this case is fixed to the point where $\omega_e = 0$. The maximum forward speed under all electric propulsion is determined from the intersection of this lever with the ω_m axis and minimum ω_g (plotted as the blue dashed line). To enable greater forward vehicle speeds, the ICE must now start to rotate. As the speed of the GU controls

[1] By geometric arguments, $N_r = 2N_p + N_s$, where N_p is the number of teeth on the planet gear.

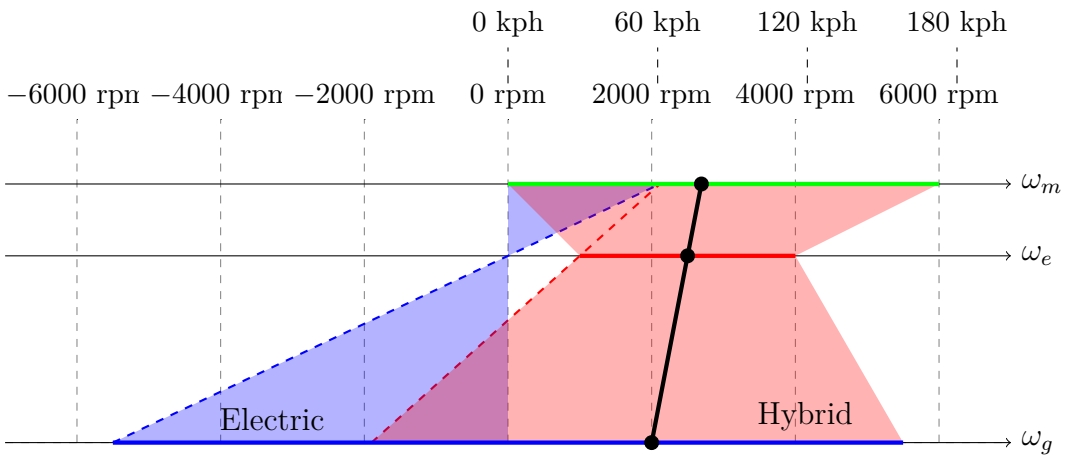


Figure 5.4: The relation between the rotational speeds of the three power devices for power generation. The motor speed is dictated by the vehicle speed, and the selection of either the generator speed fixes the engine speed or vice versa. The black line represents an example solution to (5.6). The region in which the powertrain can operate in all electric mode is shown in blue. The red region shows the available operating points for the hybrid operating mode. The engine is required to start to move from the blue to the red dashed line. The maximum all-electric speed is approximately 60km/h.

the rotational speed of the ICE it behaves like a starter motor. In hybrid power delivery mode, the ICE is on, ω_e has to remain above the idling speed of 1000rpm. An example operating point is plotted as the black line in Figure 5.4. During stationary operation ($\omega_m = 0$), the ICE can be idled by balancing the rotational speed with positive ω_g - this removes the need for a clutch. This independence of ω_e and ω_m enables more efficient operation of the ICE by decoupling it from the wheels. By comparison to other more conventional parallel hybrid architectures, the freedom to choose the engine torque and speed is clear. In a typical parallel hybrid, the rotational speed of the ICE and electric machine are constrained at a fixed ratio. A gearbox is required to connect the power sources to the wheels which may be discrete or continuous.

Electrical Machines

It is now required that suitable smooth and reasonable models for the efficiency of each of the machines in this study are created.

The MU is coupled to the wheels via the differential, thereby enforcing unidirectional rotation limited to 6000rpm (628 rad/s). The shaft torque is lim-

ited to ±300Nm. The maximum available power for deployment or recovery is limited to ±30kW.

The efficiency of the MU (η_m) is determined by the operating point - which consists the output torque τ_m and rotational speed, ω_m. This incorporates several sources of power losses: Joule heating, inverter and mechanical.

The efficiency of the GU, η_g, is related to its shaft torque, τ_g and rotational speed, ω_g. The GU is much smaller than the MU leading to lower efficiency and lower power limits- ±15kW. It is also constrained to rotate between ±5500rpm and torque transfer is limited to ±55Nm. Efficiency data is taken from the Advisor data set [84] and approximated with a polynomial map. The efficiency maps for both machines are plotted in Figure 5.5.

Finally, the electrical power requirement for the MU $P_m^b(s)$, is related to the shaft power $P_m(s) = \tau_m(s)\omega_m(s)$ by

$$P_m^b(s) = \eta_m(\omega_m(s), \tau_m(s))^{-\text{sign}P_m(s)} P_m(s) \tag{5.7}$$

and similarly for the generator unit

$$P_g^b(s) = \eta_g(\omega_g(s), \tau_g(s))^{-\text{sign}P_g(s)} P_g(s). \tag{5.8}$$

Both sign() functions are approximated with a hyperbolic tangent as was discussed in Chapter 2.

5.1.3 Engine Model

The representative engine model in this study is a typical gasoline ICE. Chapter 4 made use of a 2D polynomial map to predict fuel flow based on the engine operating point pair ω_e, τ_e. This was suitable for the driving situations in that chapter due to the lack of requirement for engine operation below the idling speed ω_e^i. It is now required that the engine speed be allowed to go to 0rpm. It could be that the problem is formulated with a discrete variable which enables switching between the two modes of operation in Figure 5.4. Also a ruleset could be imposed which stated that below the threshold forward speed, the vehicle is operated in electric mode.

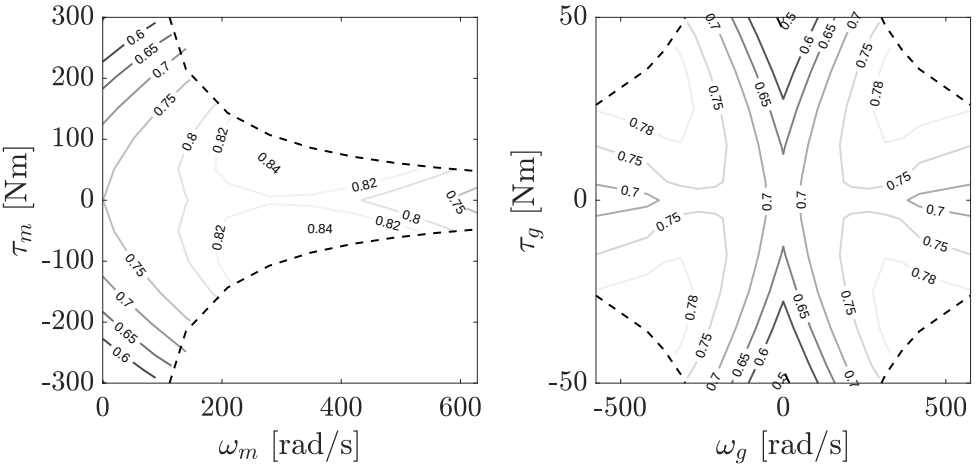


Figure 5.5: Polynomial efficiency maps for MU (left) and GU (right). Contours are levels of constant efficiency. Power limits for both machines are plotted as black dashed lines.

Once it reaches the maximum forward speed, the engine is actuated and used for propulsion. This thesis presents a different, novel approach to the problem. Instead of enforcing a presupposed rule regarding engine switching, a drag torque consisting of losses in friction and pumping work is introduced. This represents the motoring losses which are normally overcome by combustion and starter motor torques. The drag torque takes ranges from zero engine speed to the engine operating range. For simplicity, the model is assumed linear

$$\tau_{drag} = k_f \omega_e. \tag{5.9}$$

It is reasonable to assume that the fuel flow rate $\dot{m}_f$ is linear with engine crank power (referred to as the Willans' line approximation).

$$\dot{m}_f = k_1 \tau_e \omega_e + k_0. \tag{5.10}$$

In the engine operating space, there is a fuel-minimal operating line which forms a one-to-one mapping between fuel flow rate and engine power. This phenomenon arises in the solution of optimal control problems for series hybrids, vehicles equipped with CVTs [62, 70] or is used to form a simplified cost function for optimisation [61, 81]. This thesis computes this mapping offline. In this study, the operating line is formed of the torque axis and the maximum torque curve. A new control

variable can be introduced which is the torque developed by the engine provided no losses occur $\bar{\tau} = \tau_e + \tau_{drag}$.

The Willans' line approximation now reads

$$\dot{m}_f = k_1(\bar{\tau} - k_f\omega_e)\omega_e + k_0. \tag{5.11}$$

To select the parameter k_f, the drag losses at idle are chosen equal to the fuel required for idling. This retains the correct prediction of fuel requirements at low engine power requirements, which is the likely engine operating region. In other words,

$$k_f = \frac{k_0}{(\omega_e^i)^2} \tag{5.12}$$

The crank power is the difference between this torque, and the drag torque multiplied by the engine speed:

$$P_e = \omega_e(\bar{\tau} - k_f\omega_e). \tag{5.13}$$

The required fuel flow is

$$\dot{m}_f = k_1\bar{\tau}\omega_e. \tag{5.14}$$

There is clearly a limited range of validity for the model, fuel flow rates above 8kW are over estimated. However, the added benefit of not imposing the engine switching cannot be ignored. This formulation avoids requirement to switch the constant term in a cost function consisting of (5.10). A constraint is required to ensure that the engine does not produce power below the engine idle-speed ω_e^i:

$$0 \leq \bar{\tau}(s)(\omega_e(s) - \omega_e^i). \tag{5.15}$$

Finally, the following constraints on the engine speed ω_e and the fuel torque $\bar{\tau}$ must be respected:

$$0 \leq \omega_e(s) \leq \omega_e^{max}, \tag{5.16}$$

$$0 \leq \bar{\tau}(s) \leq \bar{\tau}^{max}(\omega_e). \tag{5.17}$$

As the engine operates at stoichiometric air-fuel-ratio ($\text{AFR}_s = 14.7$), the mass flow of exhaust gas out of the engine is given as:

$$\dot{m}_e(s) = (1 + \text{AFR}_s)\dot{m}_f(s). \tag{5.18}$$

In order to ensure the models effectively predict engine outlet gas conditions, a switch is required which smoothly transitions between the engine on and off states.

$$\phi(\bar{\tau}) = \frac{1 + \tanh(\bar{\tau})}{2} \tag{5.19}$$

The exhaust gas temperature T_{exh} at manifold exit, is described by a bi-quadratic map based on the engine speed $\omega_e(s)$ and the net crank torque τ_e,

$$T_{exh}(\omega_e, \tau_e) = \phi(\bar{\tau})T_e\Big(\omega_e, \tau_e\Big). \tag{5.20}$$

The map $T_e(\omega_e, \tau_e)$ is a polynomial of the following form.

$$T_e(\omega_e, \tau_e) \;=\; 340 \;+\; 1.2\omega_e \;+\; 12\tau_e \;-\; 0.0013\omega_e^2 \;-\; 0.055\tau_e^2 \;-\; 0.0034\omega_e\tau_e \tag{5.21}$$

Instantaneous flow rates of carbon monoxide (CO) ($\dot{m}^e_{CO}(s)$), unburnt hydrocarbons (HC) ($\dot{m}^e_{HC}(s)$) and oxides of nitrogen (NOx) ($\dot{m}^e_{NOx}(s)$) from the ICE can be generated based on the crank pairing $\omega_e, \bar{\tau}$. Engine emissions are much more non-linear than fuel flow and as such cannot be computed linearly based on engine power data. Two-dimensional polynomial maps are accurate enough for this work. For reasons of computational speed the Gaussian process models of the previous chapter are too slow when combined with the rest of the powertrain model to be effective. The engine out flow rates are the following:

$$\begin{aligned}\dot{m}^e_{NOx}(s) = \phi(\bar{\tau}(s))\Big[&- 0.0098 + 0.00012\omega_e - 0.00026\tau_e \\ &- 1.8 \times 10^{-7}\omega_e^2 + 2 \times 10^{-6}\tau_e^2 + 2 \times 10^{-6}\omega_e\tau_e\Big],\end{aligned} \tag{5.22}$$

$$\begin{aligned}\dot{m}^e_{HC}(s) = \phi(\bar{\tau}(s))\Big[&0.0044 + 4.3 \times 10^{-5}\omega_e - 9.1 \times 10^{-5}\tau_e \\ &-6.6 \times 10^{-8}\omega_e^2 + 7 \times 10^{-7}\tau_e^2 + 7.1 \times -7\omega_e\tau_e\Big],\end{aligned} \tag{5.23}$$

$$\begin{aligned}\dot{m}^e_{CO}(s) = \phi(\bar{\tau}(s))\Big[&- 0.0084 + 0.00022\omega_e - 0.00047\tau_e \\ &-3.4 \times 10^{-7}\omega_e^2 + 3.6 \times 10^{-6}\tau_e^2 + 3.7 \times -6\omega_e\tau_e\Big].\end{aligned} \tag{5.24}$$

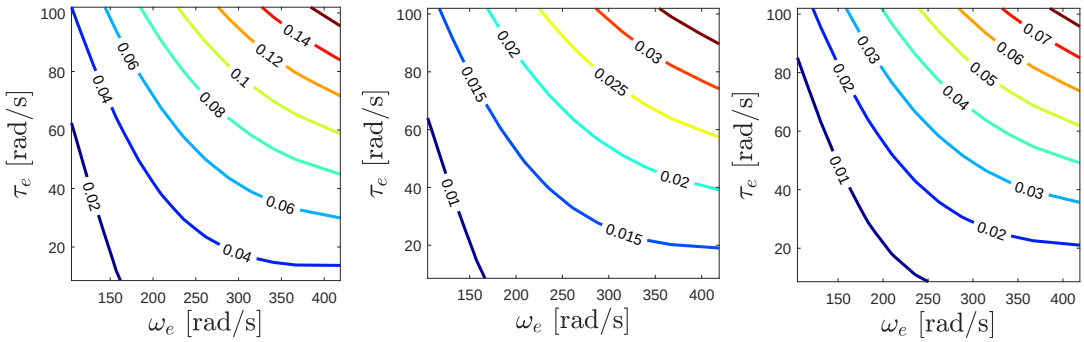


Figure 5.6: Emissions mass rate flow maps [g/s]. Left-to-right: CO, HC and NOx.

The response surfaces for each are plotted in Figure 5.6. Data for both temperature and emissions flow rates is taken from the Advisor database [84].

5.1.4 Energy Store

The ES is modelled by the Thévenin equivalent circuit model (with internal resistance R_b) where the ideal source voltage $V_s(s)$ is a function of the charge stored in the battery, $Q_b(s)$ [70]. The terminal voltage $V_b(s)$ and current $I(s)$ are related by

$$V_b(s) = V_s(Q_b(s)) - R_b I(s). \tag{5.25}$$

The source voltage is an affine function of the stored charge:

$$V_s(Q_b(s)) = V_b^{min} + (V_b^{max} - V_b^{min})\frac{Q_b(s)}{Q_b^{max}}. \tag{5.26}$$

A useful definition is the state-of-charge (SOC),

$$\text{SOC(s)} = \frac{Q_b(s)}{Q_b^{max}}. \tag{5.27}$$

The maximum battery charge is defined by

$$Q_b^{max} = \frac{2E_b^{max}}{V_b^{max} + V_b^{min}}. \tag{5.28}$$

The battery output power is the product of the terminal voltage and current, $P_b(s) = V_b(s)I(s)$. A quadratic equation relating the battery current and power can be formed

$$I(s)^2 - \frac{V_s(Q_b(s))}{R_b} I(s) + \frac{P_b(s)}{R_b} = 0. \tag{5.29}$$

This sign convention implies that positive $I(s)$ corresponds to discharging the battery. The determinant of (5.29) ensures that the maximum power available from the battery at any instance is greater than the power demanded from it. In order to avoid calculation of the explicit solution to (5.29), $I(s)$ can be introduced as an algebraic variable and (5.29) becomes an algebraic (constraint) equation. This avoids the instance where solutions may become complex. The dynamic equation describing the SOC in the battery is

$$\frac{dQ_b(s)}{ds} = -\frac{I(s)}{v(s)}. \tag{5.30}$$

The battery power is balanced by the demands from each of the electrical machines

$$P_b(s) = P_m^b(s) + P_g^b(s). \tag{5.31}$$

Finally the power balance between the required tractive power and the mechanical output power from the prime movers is as follows:

$$P_t(s) - P_e(s) - P_g(s) - P_m(s) = 0. \tag{5.32}$$

5.1.5 Three-way Catalytic Converter

The TWC reduces the tailpipe amounts of the three emissions species. It achieves this by oxidising CO and HC into CO_2 and H_2O and reducing NO_x to N_2 [76]. The reactions occur on catalytic bed which has a dependence on the temperature. Each of the reactions is most favourable at high substrate temperatures. This relies on the substrate body reaching the *light-off* temperature. This is the point at which above 50% emissions reduction can occur. The reduction efficiency characteristics can be represented by the functions

$$\eta_{CO}(s) = 0.475(1 + \tanh(0.016(T_c(s) - 499))) \tag{5.33}$$

$$\eta_{HC}(s) = 0.490(1 + \tanh(0.014(T_c(s) - 521))) \tag{5.34}$$

$$\eta_{NOx}(s) = 0.490(1 + \tanh(0.016(T_c(s) - 488))). \tag{5.35}$$

These characteristics are plotted in Figure 5.7.

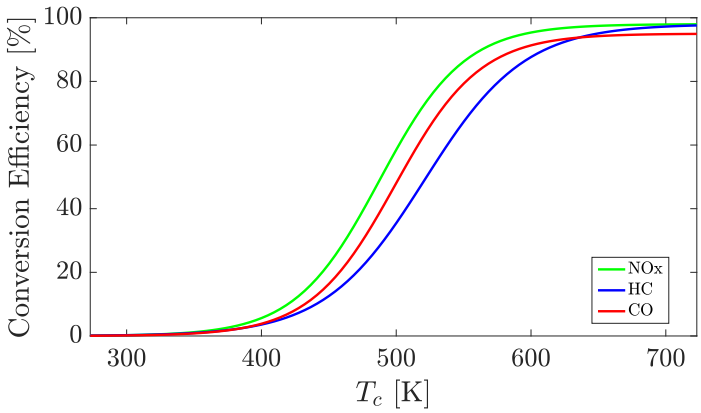


Figure 5.7: Catalytic converter temperature dependencies.

The TWC block temperature, T_c is modelled based on the incoming heat and the convective cooling [85],

$$\dot{T}_c(s) = \frac{\dot{m}_e(s)c_p}{C_c}(T_{exh}(s) - T_c(s)) - \frac{h_a}{C_c}(T_c(s) - T_a). \tag{5.36}$$

where T_a is the ambient air temperature, C_c is the specific heat capacity of the block itself, c_p is the heat capacity of the exhaust gas and h_a is the heat transfer coefficient to the surrounding environment. The total tailpipe emissions created along the route are

$$m^t_{CO}(s) = \int_0^{s_f} (1 - \eta_{CO}(s))\dot{m}_{CO}(s)\frac{ds}{v(s)} \tag{5.37}$$

$$m^t_{HC}(s) = \int_0^{s_f} (1 - \eta_{HC}(s))\dot{m}_{HC}(s)\frac{ds}{v(s)} \tag{5.38}$$

$$m^t_{NOx}(s) = \int_0^{s_f} (1 - \eta_{NOx}(s))\dot{m}_{NOx}(s)\frac{ds}{v(s)}. \tag{5.39}$$

5.1.6 Route Data

Once again the route data for the OCP is estimated based on the method discussed in Appendix B. For reference, the curvature and elevation are plotted in Figure 5.8.

The distance dependent speed limits, $v^l(s)$, are listed in Table 5.1a. These form the independent phases for the multi-phase formulation. Between each of these phases, continuity of the vehicle state is required. The exception to this constraint is at instances where the vehicle has to come to rest. At each of the

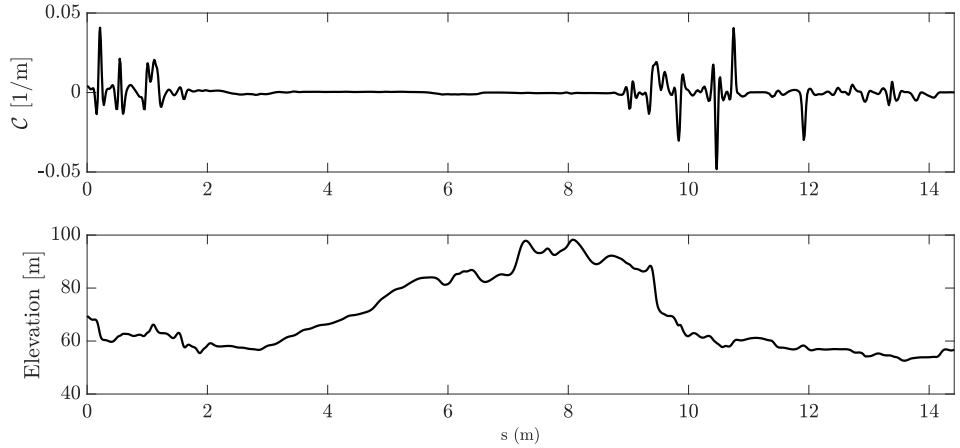


Figure 5.8: Curvature and elevation for Route B.

Table 5.1: Route B speed limit and stops information.

(a) Phase structure of optimal control problem.

s k[m]	v^l [mph]
$[0, 0.95]$	30
$[0.95, 0.98]$	30
$[0.98, 9.34]$	70
$[9.34, 10.47]$	70
$[10.47, 10.52]$	70
$[10.52, 10.75]$	30
$[10.75, 13.16]$	30
$[13.16, 13.33]$	20
$[13.33, 14.42]$	20

(b) Stop distances.

s_s [km]
0
0.95
9.34
10.47
10.75
13.33
14.42

stopping distances s_s, the additional constraint that $v(s_s) = \epsilon$ is enforced. The stopping distances are listed in Table 5.1b.

Within each phase, the following optimal control problem is solved. Overall boundary conditions are imposed at the beginning of the first and end of the last phases. Total integrals are calculated by summing contributions from each phase.

5.1.7 Optimal Control Problem Setup

The optimal powersplit and driving control problem can be formulated into the following. Implicit dependence on elapsed distance is omitted for clarity. Minimisation of the total fuel consumption over the route is given by the integral:

$$J = \int_0^{s_f} k_0 \omega_e(s) \bar{\tau}(s) \, \frac{ds}{v(s)}. \tag{5.40}$$

This cost function is minimised subject to the set of dynamics consisting the vehicle speed, battery SOC, the TWC temperature, the cumulative tailpipe emissions and time. These form a state vector, which is bounded in the following way

$$\begin{bmatrix} \epsilon \\ 0.2Q_b^{max} \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \leq \begin{bmatrix} v \\ Q_b \\ T_c \\ m^t_{NOx} \\ m^t_{CO} \\ m^t_{HC} \\ t \end{bmatrix} \leq \begin{bmatrix} v_m \\ 0.8Q_b^{max} \\ 1000 \\ \overline{m}^t_{NO_x} \\ \overline{m}^t_{CO} \\ \overline{m}^t_{HC} \\ T \end{bmatrix}. \tag{5.41}$$

The corresponding dynamic equations are formed from (5.2), (5.30), (5.36), (5.37), (5.38) and (5.39) [2]. Constraints on the total emissions are placed here as overbar quantities, while the maximum allowable route time is T. The cars top speed is v_m. Initial conditions on the state are as follows

$$\mathbf{x}(0) = \begin{bmatrix} \epsilon & Q_b(0) & T_a & 0 & 0 & 0 & 0 \end{bmatrix}^\top. \tag{5.42}$$

The control vector is made up of the tractive power; torques for each of the engine, motor and generator; the engine speed and the battery current:

$$\mathbf{u} = \begin{bmatrix} P_t & \bar{\tau} & \tau_g & \tau_m & \omega_e & I \end{bmatrix}^\top \tag{5.43}$$

The equality constraint on the system is

$$R_b I^2 - IV_s + P_b = 0 \tag{5.44}$$

and inequality constraints take the form

$$\begin{bmatrix} -\infty \\ P_m^{min} \\ P_g^{min} \\ \omega_g^{min} \\ -\infty \\ 0 \end{bmatrix} \leq \begin{bmatrix} P_t - (P_e + P_m + P_g) \\ \tau_m \omega_m \\ \tau_g \omega_g \\ \omega_g \\ \bar{\tau} \\ (\omega_e - \omega_e^i)\bar{\tau} \end{bmatrix} \leq \begin{bmatrix} 0 \\ P_m^{max} \\ P_g^{max} \\ \omega_g^{max} \\ \bar{\tau}^{max}(\omega_e) \\ \infty \end{bmatrix}. \tag{5.45}$$

Admissible accelerations are bounded by the following inequalities :

$$-1 \leq \frac{a_x(s)}{a_x^{max}} - \frac{a_y(s)}{a_y^{max}} \leq 1 \tag{5.46}$$

$$-1 \leq \frac{a_x(s)}{a_x^{max}} + \frac{a_y(s)}{a_y^{max}} \leq 1 \tag{5.47}$$

[2] The time state equation is $\dot{t}(s) = \frac{1}{v(s)}$.

Table 5.2: Parameter values and descriptions.

Parameter	Value	Description [Units]
M	1750	Vehicle mass [kg]
c_p	1003	Heat capacity of exhaust gas [J/K/kg]
T_a	298	Ambient temperature [K]
h_a	12	Ambient heat transfer coefficient [W/kg]
C_c	2000	Heat capacity of the catalyst [J/Kg]
$T_{l,HC}$	521	Light off temperature HC [K]
$T_{l,CO}$	499	Light off temperature CO [K]
T_{l,NO_x}	488	Light off temperature NO_x [K]
r_w	0.3	Wheel radius [m]
N_f	3.93	Final drive ratio [-]
V_b^{max}	240	Maximum battery voltage [V]
V_b^{min}	210	Minimum battery voltage [V]
E_b	7	Maximum battery energy [MJ]
R_b	0.5	Internal resistance [Ω]
N_s	30	Number of teeth on sun [-]
N_r	78	Number of teeth on ring [-]
v_m	55	Vehicle maximum speed [m/s]

Parameter values for the model are listed in Table 5.2. Powertrain data is taken from the open-source Advisor database [84].

5.2 Results

The optimal control problem is solved with an adaptive Legendre-Gauss-Radau (LGR) direct collocation method which is implemented in GPOPS-II [39]. Derivative information is provided by automatic differentiation software (ADiGator [50]) and the arising non-linear program is solved with the interior point method using IPOPT [35]. The problem is scaled following the method using the non-dimensionalising method detailed in Chapter 2.

As has been established in the previous chapters, regularisation is required to ensure that solutions can be computed when singular arcs occur. Therefore, the cost is modified again to include derivatives on the control efforts, and the rate of control now becomes the input for some of the variables:

$$J' = \int_0^{s_f} k_1 \bar{\tau}\omega_e + r_1 \dot{\tau_e}^2 + r_2 \dot{\tau_g}^2 + r_3 \dot{\tau_m}^2 + r_4 \dot{\omega_e}^2 \frac{ds}{v(s)} \tag{5.48}$$

A series of test cases are now considered. The first computes the fuel-optimal route traversal solution and examines its structure. This forms a baseline control strategy, which can be used as a reference when emissions constraints are imposed. This result is also used to consider the effect that the regularisation has on the problem solution. Total emissions constraints are then imposed relative to baseline quantities. This will enable differences in the heating strategies that arise for specific emissions types to be examined. Furthermore, the optimal solutions are computed and differences for limiting emissions cases considered on the tractive power and velocity pair. These form the driving part of the solution. Finally the route is traversed in the opposite direction to appraise the differences in road geometry and speed limits. Emissions constraints are once again introduced. The comparison of routes ensures any findings for emissions reductions are explained in terms of species and the power requirements.

5.2.1 Baseline Study: Fuel-Optimal

A fuel minimum solution is first computed for the traversal of Route B1. The SOC is constrained to begin and end at 60% (i.e $Q_b(0) = Q_b(s_f) = 0.6Q_b^{max}$). The arrival time is stipulated to be less than 1000 seconds.

The optimal velocity and tractive power profiles are plotted in Figure 5.9. In the same manner as in Chapter 4, the tractive power is required to accelerate from rest, after cornering and during the ascent of hills. Braking is required to slow the car during speed limit changes, cornering events and to bring the car to rest. The highest speed limits are not used, as the arrival time constraint is chosen to be longer than the minimum time to complete the journey[3]. The upper bounded arrival time is met again aligning with the conclusions of Chapter 3.

The optimal powersplit trajectories are plotted in the middle of Figure 5.9. The few available braking events mean that there is a lack of opportunity to increase the stored energy in the ES. Therefore the ICE is required predominantly to fulfil tractive power demands. Unlike conventional use of an ICE, there is no clutched

[3]In the minimum time case, the car is driven as fast as possible, reaching all of the speed limits.

gearbox to overcome the lack of torque at low speeds, so the vehicle is driven from corners and stopping events via the electrical part of the powertrain. During braking, the electrical part of the powertrain recovers energy. In fuel-optimal powertrain operation, the ES is maintained at approximately 60% charge for the duration of the journey (the bottom plot in Figure 5.9). The terminal SOC constraint is met by "over-powering" the engine during the 11km to 13km section. All of the electrical energy that is deployed consists the energy recovered in braking and any periods of engine generator behaviour. The latter only occurs during periods where the demanded tractive power is below that of the ICE power at the intersection of the maximum torque line and the torque axis $\tilde{P}_e \approx 8$kW (see the left hand plot in Figure 5.10). This avoids losses in the engine drag torque. ICE, MU and GU operating points are scattered in Figure 5.10. Engine powers are all achieved along the operating line determined by the idle speed and maximum torque line. This is a direct consequence of the model formulation. This phenomenon is observed in other literature where the engine operating points are chosen in the absence of a speed constraint [36, 70] (i.e when the powertrain utilises a continuously variable transmission (CVT)). Due to the higher efficiency of the MU, torque (and thus power) is deployed through this mechanical route to the wheels. The GU serves to decouple the engine speed from the wheel speed and so is rarely torqued. The MU is predominantly used for electric propulsive efforts, while the GU aids in braking when the power demand is higher than the most efficient part of the MU map.

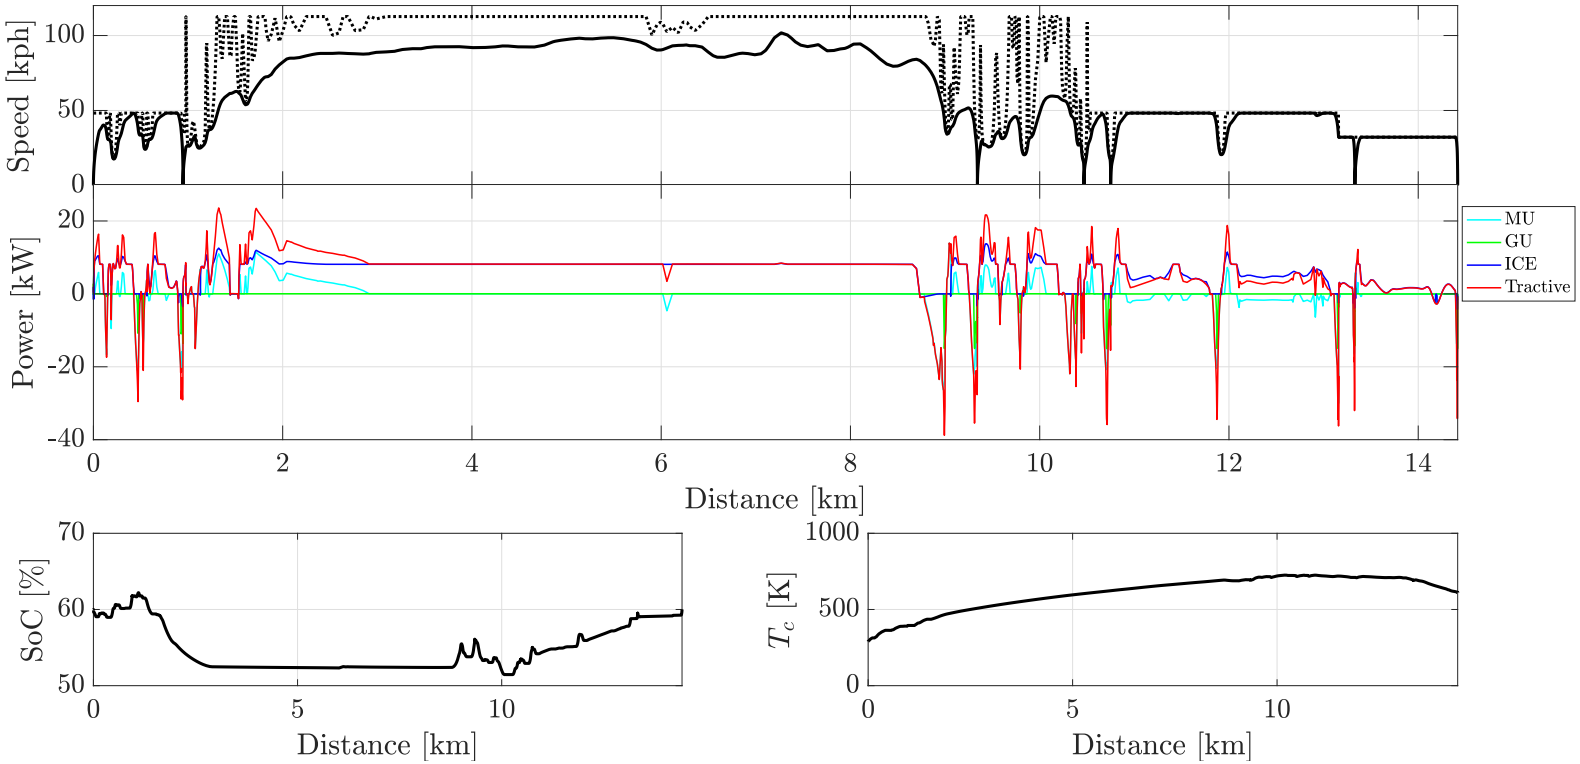


Figure 5.9: Top: Optimal speed (km/h) profile (solid). The route speed limits are dotted, taking the minimum of the legislative speed limit and the lateral acceleration limit. Middle: Tractive power (kW) is plotted along with the power delivery from each device. Bottom left: The SOC [%] of the energy store is plotted. Bottom right: TWC temperature.

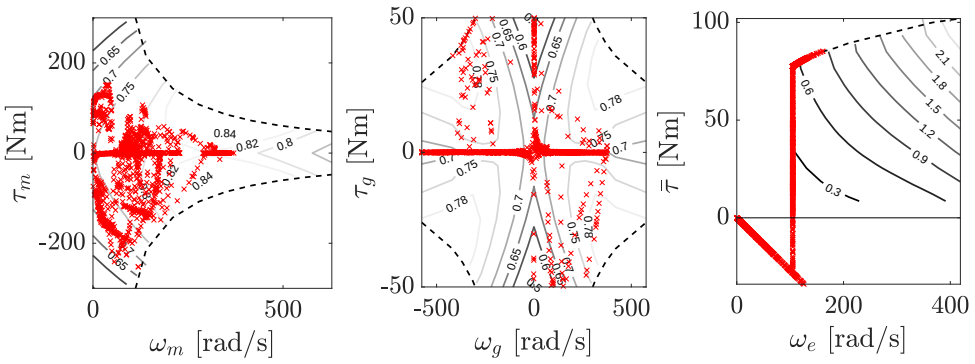


Figure 5.10: Left: Motor unit operating points on contours of constant efficiency. Centre: Generator unit operating points on contours of constant efficiency. Right: Engine operating points overlaid on fuel map. The engine torque is net torque - whenever it is negative it corresponds to that required to start the engine.

The TWC temperature profile is increasing for the majority of the journey, peaking just before the end (see Figure 5.9). The constant engine power drives the TWC to heat, until it has reached a temperature of around 700K. The temperature starts to fall again as the engine power requirement decreases towards the end of the drive.

The point at which light-off is reached is clearly seen in the emissions flow rates plots in Figure 5.11. It takes approximately 6km of driving before the point at which the TWC is most effective for the abatement of all emissions species. The lower ICE powers near the end of the trip mean a decrease in the engine output temperature, but as the substrate has already reached its operating regime, emissions are still effectively curtailed. It is also clear to see that the majority of emission production occurs within the first 4km of this route. The time taken to drive to 6km is approximately 344s.

The baseline fuel-optimal problems take the average of the state boundaries as their initial guesses. The later emissions constrained cases use this solution as the initialisation point. All of the computations were completed on an 8 core 3.5 GHz desktop PC. Due to the multiple phase formulation of the problem, solution speed can vary as each phase converges at different rates. This effect, coupled with the fact that there are rapid changes in the engine speed trajectory impede the solution time (for example the fuel-optimal solution computes in approximately 7000s).

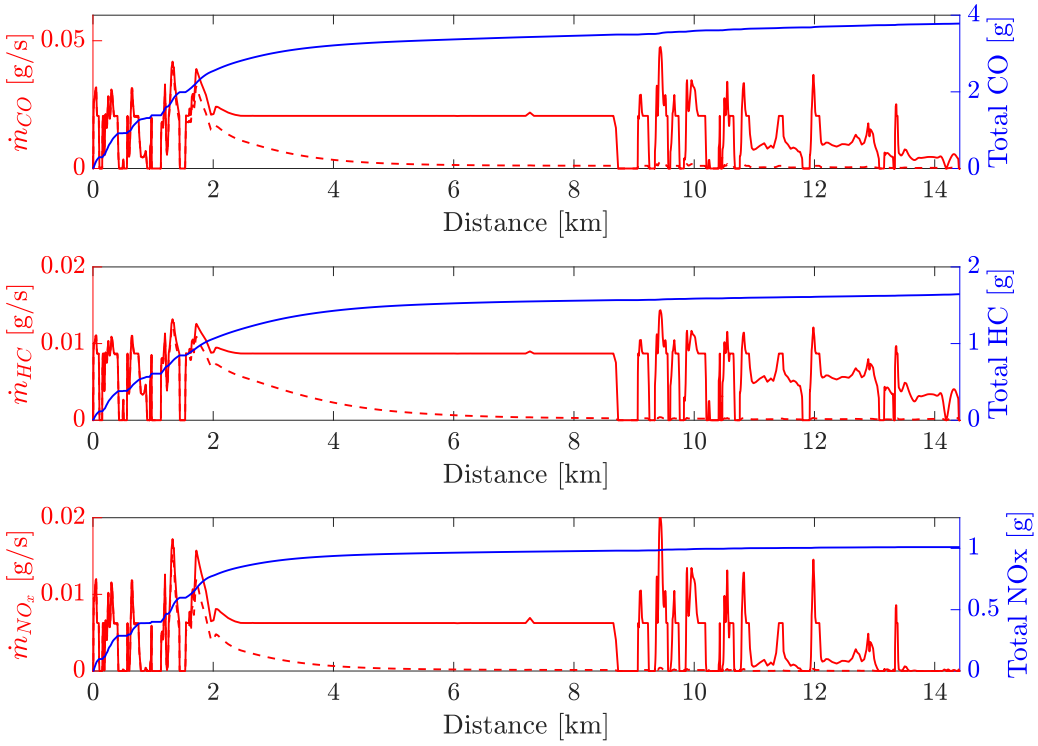


Figure 5.11: Red shows the flow rates of each emissions species: solid is engine out (Equations (5.22),(5.23) and (5.24)), and dashed is tailpipe (Equations (5.37), (5.38) and (5.39)). The blue trace is the cumulative tailpipe emissions.

Regularisation Weight

Before the effects of emissions constraints can be examined, the selection of the regularisation weight needs consideration. In the case of this optimisation, the final powertrain switching strategy is determined by the size of the engine speed weighting (r_4) in (5.48). The sensitivity of the solution to changes any of the torque weightings (r_1, r_2 or r_3) is negligible. Typically, when computing numerical optimal control solutions, one would start with a relatively large regularisation weighting which is slowly decreased until the percentage of the total cost is insignificant [54]. However, in the case of the solution to this problem, r_4 also doubles as a means to avoid unnecessary switching of the engine and remains a reasonably large amount of the total integral cost as a result. As the regularisation weight r_4 increases, the engine speed is penalised for changing rapidly. This decreases the number of ICE

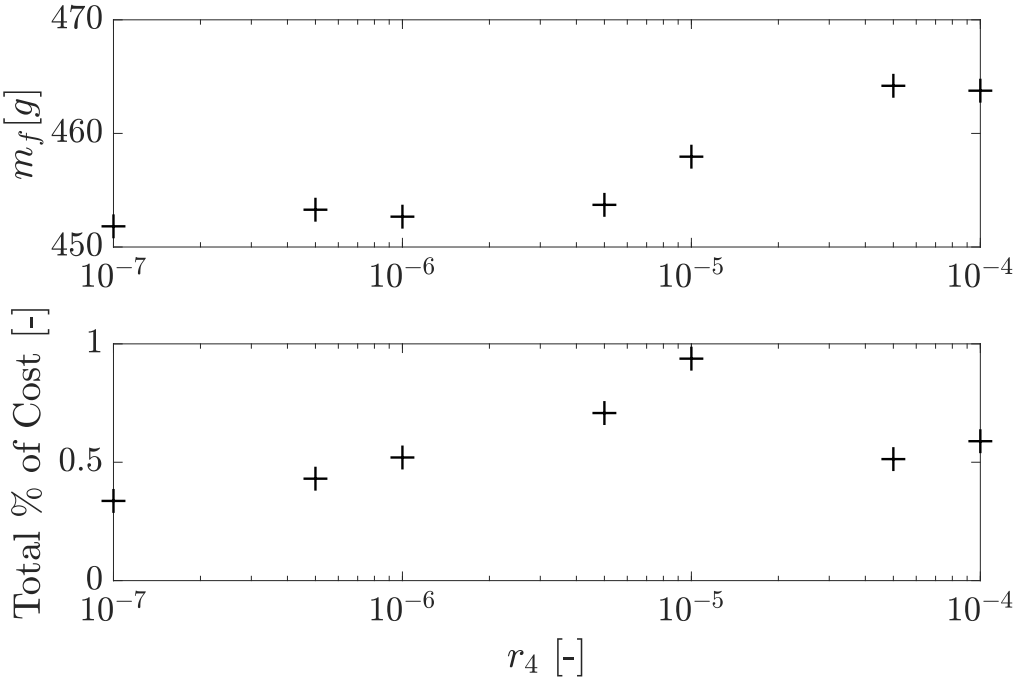


Figure 5.12: Total fuel consumption and total regularisation percentage against the regularisation weight on the engine speed.

switches and leads to an increased fuel consumption (see Figure 5.12). Initially, the total percentage of the cost increases, but decreases once the engine strategy changes (at $r_4 = 5 \times 10^{-5}$). At this point, the optimiser no longer tries to enforce complete engine switch-offs during the drive. Despite the increase in regularisation, it is still limited to 1% of the total cost. An inspection of the resulting trajectories for the engine speed reveals that the solution does not change when the tractive power is singular at velocity limits. Propulsive power during these periods is produced solely by the engine and, as there is an optimal one to one mapping for engine power, speed and torque, the solution remains unchanged.

For four of the regularisation weights, sections of interesting behaviour for the resulting ω_e trajectories are plotted in Figure 5.13.

For low regularisation weights switches occur quickly. This minimises losses in the drag torque during the switching period. This also leads to an additional engine off-on switch at 1.5km. While this is conducive to reducing the fuel consumption, it slows the problem computation down.

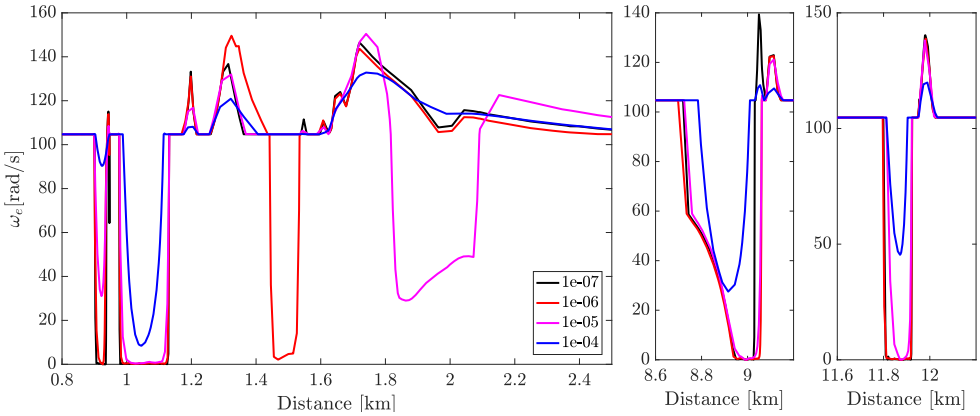


Figure 5.13: Changes in the engine switching strategy (by comparison of the optimal engine speed trajectories) for different cases of r_4 in(5.48).

At the other extreme, the highest regularisation smooths the trajectories considerably. Incomplete switches now occur resulting in increased losses in the engine drag torque. Furthermore, peak speeds are lowered which correspond to lower ICE power. Altering the weighting can also lead to differing switch timing. For example, $r_4 = 1 \times 10^{-5}$ leads to the additional half switch at 2km and $r_4 = 1 \times 10^{-6}$ has a switch by itself at 1.5km.

In the case of changes to the regularisation weighting, changes to the developed drive-cycle are not unreasonable. These occur when the velocity limits are not met, as there is the ability to alter the speed profile here. Any changes that do occur are still smaller than the alterations to the driving that enforcing emissions constraints will impose.

Increasing the regularisation weight beyond $r_4 > 5 \times 10^{-4}$ prevents the engine from shutting off entirely. Therefore, for the remainder of this chapter, a weighting of $r_4 = 1 \times 10^{-6}$ is chosen as the trade-off between smoothed trajectories and the ability for the optimisation to turn the engine on and off.

5.2.2 Case Study 1- Emissions Constraints

In the case that the emissions production is unconstrained, the engine is mainly used for propulsion, with few instances of generator behaviour. Recharging the ES via the ICE is only completed once the catalyst substrate is above the light-off

temperature and the demanded tractive power is below $\tilde{P}_e$. This fulfilment of the terminal SOC constraint using generator behaviour late in the journey may not remain the optimal means for powersplit when emissions constraints are introduced.

For cases where emissions production is limited, differing strategies could occur. In one case, the ICE could be used to heat the TWC more promptly during the initial parts of the drive. This power will exceed the tractive power demand and excess energy will be siphoned off into the ES. The terminal SOC constraint is fulfilled by discharging the ES. Another viable strategy is to reduce engine usage during the heating period, and discharge the ES more heavily than a fuel-optimal strategy. The battery could then be refilled later in the route when emissions are more heavily reduced.

To determine the method by which individual emissions species can be reduced, state constraints (5.41) are enforced on the total route emissions. These limits are taken relative to the total emissions produced in the baseline solution and are reduced until the OCP becomes infeasible. As there are three emissions types, it is necessary to check the solution sensitivity to constraints on each of these in turn. Results will be considered in terms of changes to fuel consumption and quantities of other emissions species. Trajectory changes in both engine operation and in the driving style are expected. The latter is the difference in the tractive power and developed speed traces compared to the unconstrained baseline.

Unburnt Hydrocarbon Constraints

Constraints on the total HC will be considered first, the results of which are plotted in Figure 5.14. Imposing the limit on total HC leads to the overpowering of the engine during the first 2km of the route. This is confirmed by the increase in the SOC during this period and rapidly increasing T_c trajectories. In all these cases, the engine power is higher than the demanded tractive power. The excess is stored in the ES. During the highway section the powersplit is the same as the baseline - tractive effort is equal to $\tilde{P}_e$. Finally once the car reaches 2.5km from the end, the ES is discharged in an all-electric mode. Decreasing the constraint leads to higher

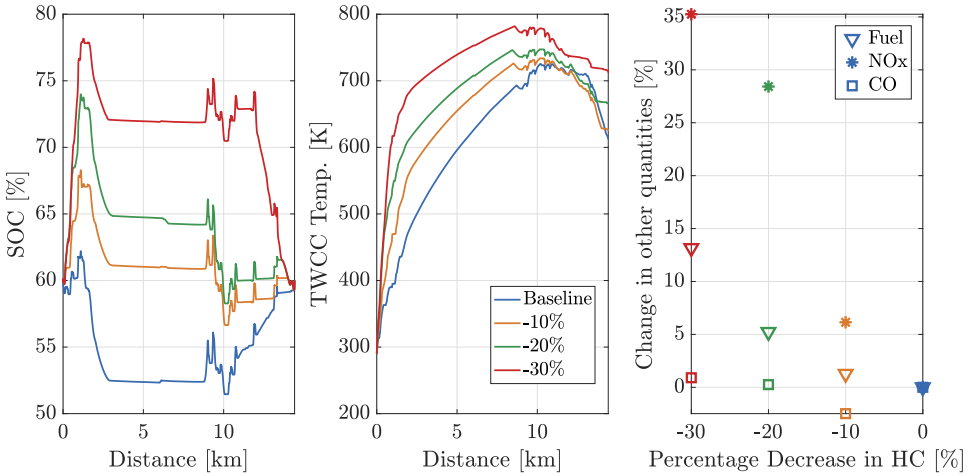


Figure 5.14: HC Constraint Results

peak stored energy. Therefore limit of abatement of HC is given by the maximum available space for storage at the beginning of the route. The time to reach light-off (for HC) in the baseline case is approximately 228 seconds, this is reduced, in the limit case, to 68s. This rapid heating strategy increases the FC for all cases.

In the limiting case, a 30% reduction in the total tailpipe HC can be achieved, but at the cost of 13% extra fuel and a 35% increase in NO_x. The increase in NO_x during the heating period is not offset by the increased reduction later in the route. The total CO is largely unchanged for any constraint on HC (see RHS of Figure 5.14).

Carbon Monoxide Constraints

As the CO is largely unchanged during HC constraints, this suggests the rapid heating strategy may not be optimal when CO constraints are introduced. Inspecting the SOC and TWC temperature trajectories in Figure 5.15 appears to confirm this.

It can be seen that, compared to the baseline, small constraints on CO favour marginally faster heating of the TWC. This ultimately leads to a 10s reduction in the heating time. However, the largest reduction in CO is achieved in a different manner. It is no longer the case that the ICE power is demanded higher than the tractive power. But rather the electric part of the powertrain is used more heavily at the start of the drive. This slower heating is synonymous with lower overall exhaust flow rates and a lower production efficiency of CO. It is likely, as will

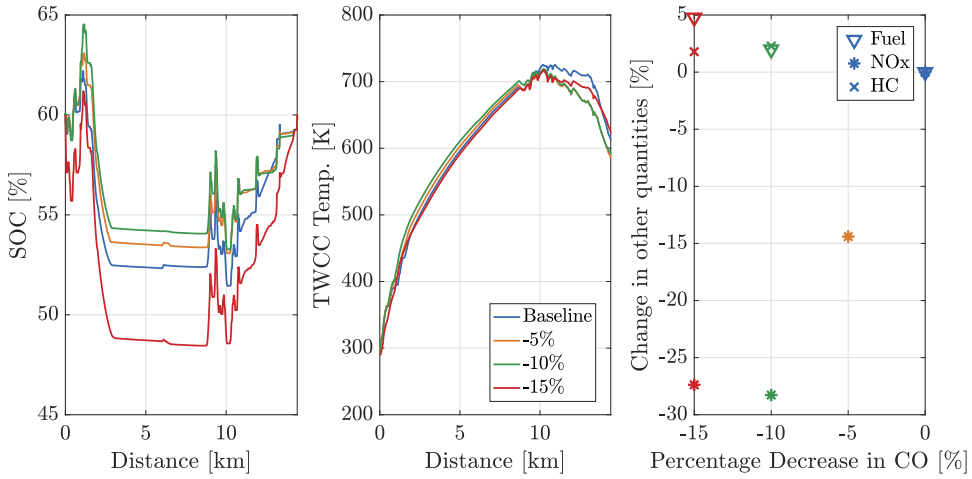


Figure 5.15: CO Constraint Results

be examined shortly, that the engine switching strategy has altered slightly. The limiting case for CO reduction is only 15% below the baseline quantity. The total fuel consumption is increased by 5%, the NO_x is reduced by 30% and the quantity of HC is increased marginally by 3%. The limitation for this means of reduction comes in the form of the ability for the powertrain to recover the energy later in the route if the SOC is discharged early on. While it may be more optimal in terms of emissions to reduce engine power and slowly heat the TWC, followed by burning fuel later to refill the ES, this is not necessarily optimal in terms of fuel economy.

Nitrogen Oxide Constraints

Constraints on the total NO_x give similar results to those for CO (see Figure 5.16). Reductions in total NO_x are accomplished in the same manner as CO. In the NO_x cases, the temperature trajectories are more difficult to discern. SOC trajectories are also much more close to the baseline than either of the previous two emissions species. These results suggest that a large decrease in NO_x can be achieved by altering the engine powering only slightly. This smaller adjustment leads to modest changes in the heating time: within ±5s of the baseline. The limit for reduction of NO_x is approximately 30% for this route and the time to reach the NO_x light-off temperature is increased by 10s to 200s over the baseline.

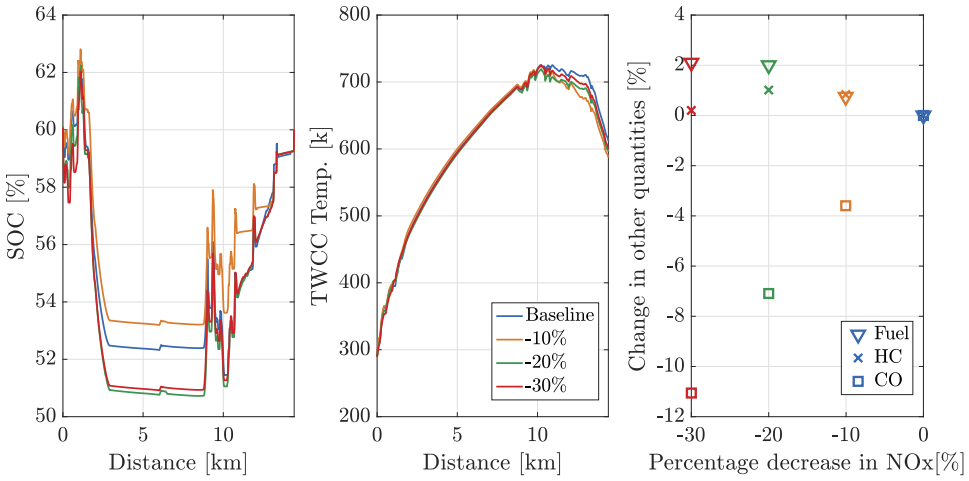


Figure 5.16: NO_x Constraint Results

To aid in the exposition of the solutions, engine powers are compared for the limiting cases of each emissions species in Figure 5.17. The resulting trajectories during warmup and near the end of the journey are compared. The highway driving section is the same for all emissions types and so is omitted. The HC reduction strategy clearly increases the engine power relative to the baseline during the first 2km. The engine is still used for power until it shuts off around 12km; beyond this the car is in electric vehicle mode. The other strategy becomes clear for the reduction of CO and NO_x. This is difficult to discern from the SOC plots alone. Each constraint requires the engine to operate below $\tilde{P}_e$. The NO_x trajectory switches off rapidly but requires a slightly higher power deployment.

During the initial section of the drive, the engine powers above $\tilde{P}_e$ are not used. During the latter stages of the drive the solution remains broadly the same as the baseline. This explains why there is little change in the developed SOC and T_c trajectories. Higher engine powers are avoided, but these are rare in the baseline case. The CO trace does not switch off as frequently and is at a lower power. In both of these cases there is a small increase in the total fuelling.

Driving Changes

For each constrained solution, the maximum arrival time boundary condition is active. This implies that the average velocity for each driving situation is the

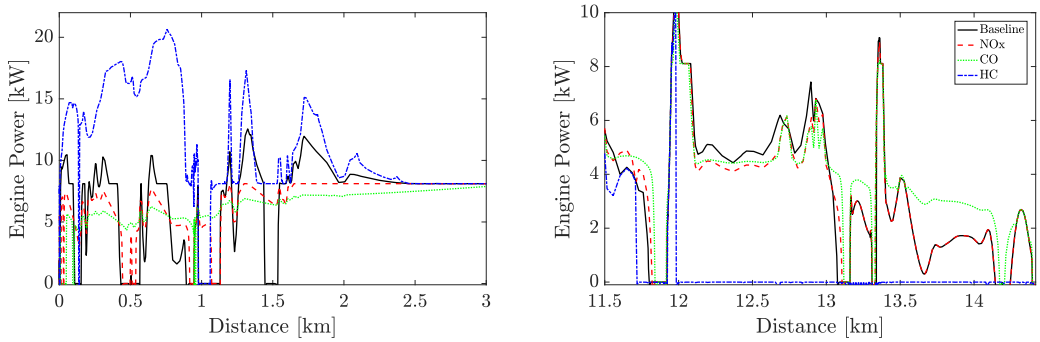


Figure 5.17: Comparison of engine output powers for each of the limiting cases for each emissions type.

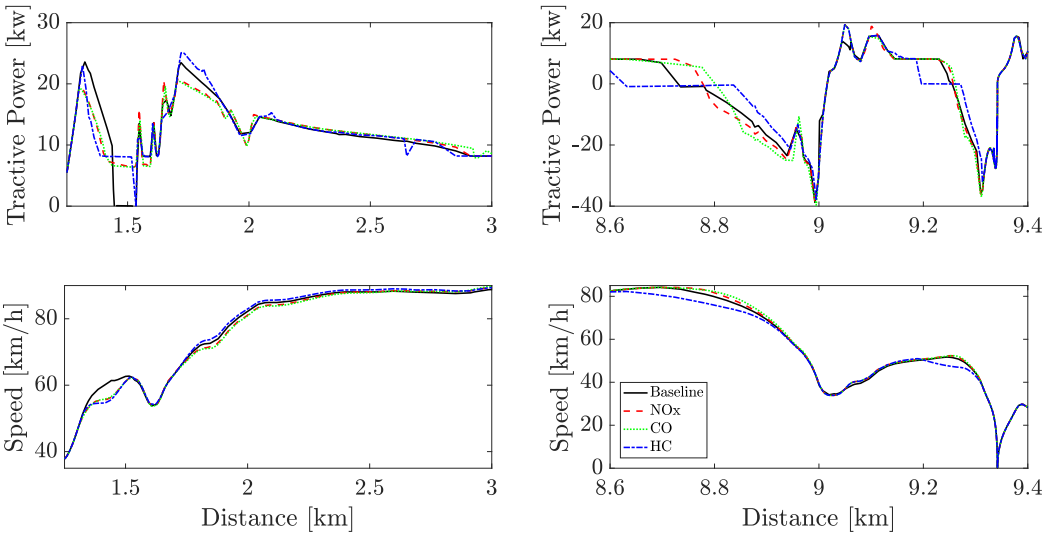


Figure 5.18: Changes to the driving part of the optimisation. The baseline trajectory is in black. Limiting constraints are imposed, NO_x is dashed, CO is dotted and HC is dot-dashed.

same. In spite of this fact, subtle changes in the optimal driving solution can arise as was shown in Chapter 4. These are expected to be less drastic for a hybrid vehicle, as the tractive power and engine powers are effectively decoupled. Any differences are found during acceleration and deceleration events where the vehicle is not speed limited - there is no flexibility for altering the power delivery when the car is at a speed limit. Example segments where changes in the driving are noted are plotted in Figure 5.18. These reflect the differences in engine power, as this is still the predominant power source.

Firstly, consider the subtle changes between 1.3 and 2.5km. The baseline

control drives the vehicle more quickly and shuts the tractive power off entirely at around 1.5km. For all emissions constraints the power is applied more smoothly, leading to a slightly lower speed in this short section. The increased or decreased ICE power appears at the wheels here as slightly higher or lower vehicle speeds during the highway driving.

Secondly, between 8.6 and 9.4 km, the HC constrained control has run the vehicle harder during the earlier stages of the drive. This means a slightly lower speed is realised at the end of the highway. The tractive effort is more readily turned off and the vehicle allowed to coast. The CO/NO_x trajectories require higher velocities at the end of the route to compensate for the marginally lower speeds earlier due to lower tractive power. This is reflected in the more aggressive driving which appears as late braking at 8.8km.

Conclusions

It has been shown that reductions in emissions can be achieved by altering the power-split along the route. In the case of HC, this is achieved by heating the TWC rapidly. An increase of emissions during the initial part of the drive are offset by greater efficiency in reduction later. The excess energy liberated from fuel is stored in the ES for deployment later. For NO_x and CO a slightly different strategy arises which operates the ICE at a lower power than the baseline. These trajectories require less engine switching and lower emissions flow rates. Tractive power demands are met by discharging the ES. Once the TWC has reached its operating temperature, the terminal SOC constraint can be met by increased engine use.

Minor changes in the driving are also achieved, especially during the entry and exit of the highway driving segment. Compared to the vehicle and powertrain in Chapter 4, changes relative to the baseline are less drastic. The ability of the powertrain to alter the engine power separately from the tractive power helps this.

It is unclear whether these results are specific to this route. It may be the case that different results arise depending on the route taken. To investigate this further, the course will be reversed and similar constraints on emissions imposed.

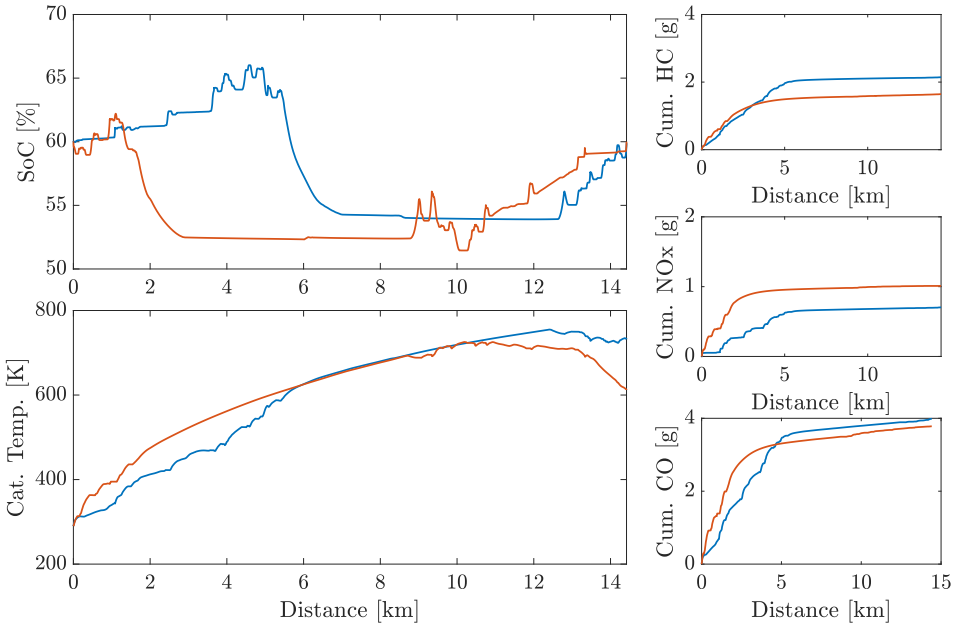


Figure 5.19: Red is the original result of the previous example (Route B1) and blue is the same route in reverse (Route B2).

5.2.3 Case Study 2 - Route B2

The initial route has an immediate period of high tractive power due to the acceleration to highway speeds at around 2 km. This enables high engine powers to be realised during the heating phase in the baseline strategy, which result in rapid heating. The return journey (denoted B2) along the same route offers a good comparison with the original drive. The initial part of the return drive is subject to an increased number of stops and lower speed limits during the first 5 km. The nature of the route during the time for the TWC to heat can determine the maximum allowable reductions in emissions. This is both in terms of how much is generated in the baseline case and how much flexibility there is for reduction. A comparison of SOC, T_c and cumulative emissions are plotted in Figure 5.19 for the fuel optimal solutions for both journeys (with the same regularisation weights). Despite the same journey distance, time, and approximately the same initial and final elevation, the power demands are significantly different. This is reflected in the ES use along the route. The ES is charged considerably more than the original case in the first

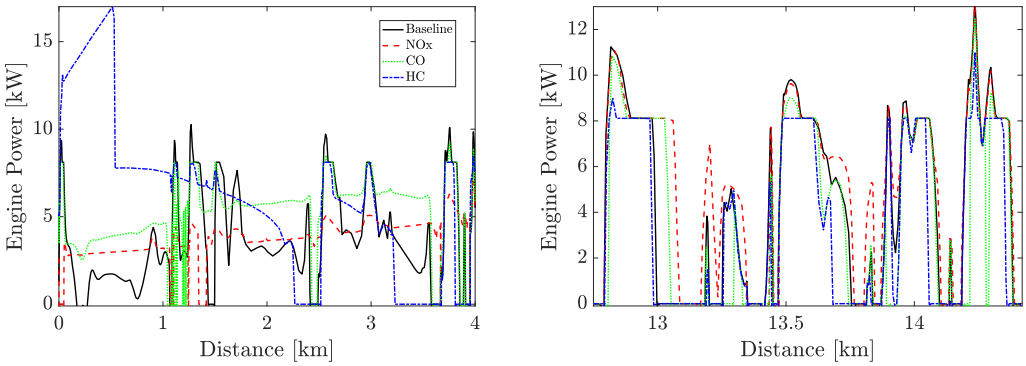


Figure 5.20: Engine power trajectories for the limiting emissions reduction cases for Route B2.

5 km and discharged to help accelerate the vehicle at the start of the highway. The terminal SOC constraint is met during the braking phases of the final 2 km.

In the B2 direction, the TWC reaches a higher temperature than traversing the route in the opposite direction, but it does so more slowly. This slower heating means an increase in the total HC produced. The rate at which HC is produced is similar for both cases during the heating phase. The mass flow rate of CO is lower during the heating period, as would be expected for the reduced power output. However, the longer period before the reduction of the TWC is effective means a greater total production. The lower engine power during heating produces less NO_x and, as the NO_x light-off temperature is the lowest, the total amount produced is lower than for Route B1. The total FC is reduced from 456g to 452g.

In order to appraise the emissions reduction strategies, total integral constraints can be imposed on this return journey relative to the new baseline. Comparison of fuel consumption, engine power trajectories and driving are considered in this case.

Unburnt Hydrocarbon Constraints

For large reductions in HC, the same overpowering heating strategy is obtained for the reverse route. Considering the engine power trajectory (blue in Figure 5.20), it is considerably higher than the baseline. During the first 2km in the baseline case, the vehicle was propelled solely by the ICE. The low speed limit infers a singular

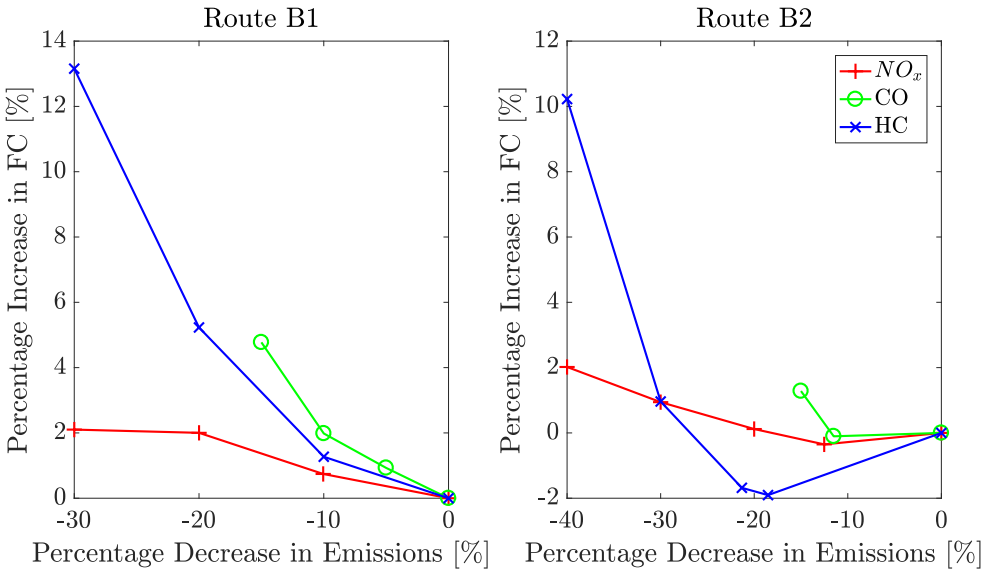


Figure 5.21: Comparison of the percentage changes in fuel consumption for each of three emissions constraints for the traversal of the route in both directions. Red - NO_x, blue - CO, green - HC.

power during this time, leading to more scope for over-powering of the engine. The time to reach the light off temperature is reduced from 400 s to approximately 50 s. A detriment to the fuel economy of 10% is realised in this case (see Figure 5.21). A greater maximum reduction in HC is possible than in the original B1 direction. This is due in part to the greater quantity emitted in the baseline drive and partly a feature of the magnitude of the tractive power demand.

There appears to be a reduction of FC for low reductions in HC due to a drastic alteration in the engine switching strategy and the power application. This can be explained as follows. The baseline case favoured smooth P_e trajectory with a regularisation percentage of less than 0.5%. In fact there is a viable solution whereby HC and FC are reduced by switching the engine an increased number of times during the first 3 km. This has an increased regularisation weighting of 0.9%, but at a lower total cost. Adjustment of the regularisation weighting in the baseline may lead to this solution being computed. Increasing the weighting may prevent this solution from being computed in this case. Solutions of this form were rejected originally due to the excessive regularisation percentage.

Nitrogen Oxide Constraints

When constraints on the total NO_x are enforced on Route B2, a more typical pareto front is realised (red in Figure 5.21). A marginal increase in the FC is seen across all NO_x emissions reductions. This result is consistent with the changes to fuel consumption along the first route. In the limiting reduction case, the engine is powered marginally more highly than the baseline (the red dashed line in Figure 5.20). A faster heating of the TWC results from this engine power (reduced by 70 s to approximately 320 s)

In order to reduce the total NO_x produced for either route, it appears that engine power of between 5 and 7 kW is required. The rate of engine production of NO_x emissions during warm up must be reduced. For B1 this power was lower than the baseline P_e, whereas for B2 this is higher. This in turn can lead to either faster or slower TWC heating. In either case, the maximum reduction in NO_x only results in a 2% increase in FC in either direction.

Carbon Monoxide Constraints

There is little reduction available for the CO species (green in Figure 5.21). In a similar manner to the NO_x constrained cases, the engine power trajectory is similar to the B1 case of constraints. In the extreme reduction case, the time to reach the light off temperature is decreased by about 120 seconds. It is now apparent that further reductions are made more difficult due to the lower conversion efficiency (compared to the other emissions species) at high temperature. In Figure 5.19, the CO emissions do not stop increasing during the drive. For this species, the trade-off for increasing the output emissions during a faster warmup are not offset by the emissions produced later along the route. The only way to reduce the CO is to run the engine between 4-8 kW, which reduces the mass flow rate before the reduction efficiency has increased.

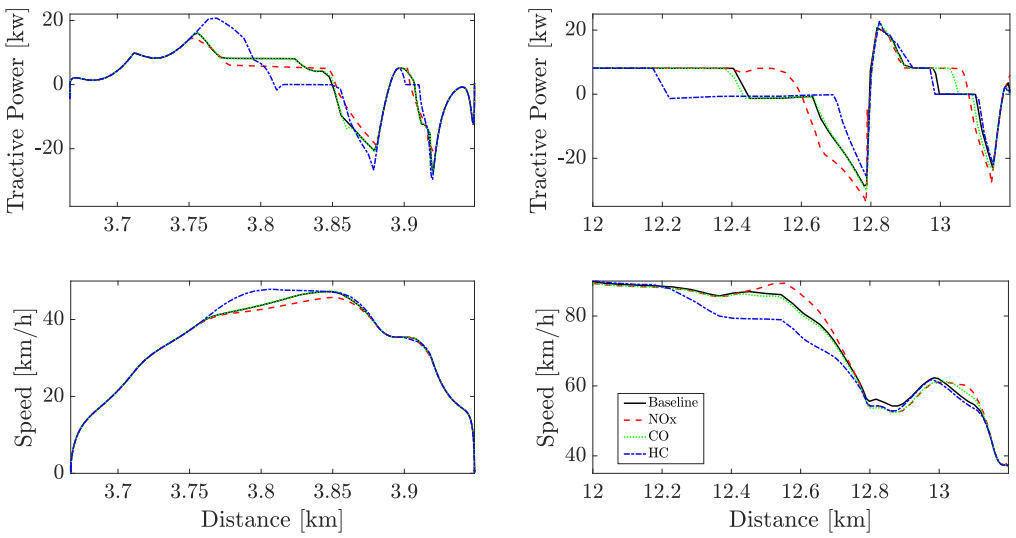


Figure 5.22: Part of the driving changes arising from the limiting constraint cases for Route B2.

Driving Changes

Again changes to the developed drive cycle can be discussed. Example changes for the limiting emissions constrained cases are plotted in Figure 5.22.

Consistent with the traversal of the route in the original direction, the arrival time is always met. This implies the same average velocity. Changes to the drive cycle are only available during periods where the vehicle is not speed limited. Noticeable tractive power differences arise around 3.65 to 3.9km. As this appears during the heating phase, the HC constrained uses higher tractive power than the baseline. In order to avoid violating any speed constraints, the power is reduced to zero at 3.8km. The limiting CO case does not require large changes to the driving, due to the low reduction of the requirement. For the most part, this coincides with the baseline. For the limited NO_x constrained case a lower power delivery is required, which in turn leads to a marginally slower speed through this section.

Further changes can be noted between 12 and 13.5km. The CO constrained case is broadly the same as the baseline again. As the HC case car was driven at high speed during the first part of the drive, the power delivery towards the end of the route must be lower. This ensures that the average speed is maintained.

The opposite is true in the case of the limiting NO_x constraint; the speed was reduced during warmup. The car during this example is propelled for longer and braked more aggressively at 12.8km.

5.3 Conclusions

This chapter discussed simultaneous optimisation of both driving and powersplit control for a parallel hybrid powertrain with thermal TWC characteristics. This was achieved by formulating a minimum fuel optimal control problem in the continuous elapsed distance domain. A detailed powertrain model formed dynamic constraints for the problem. This included powertrain efficiency models and emissions characteristics.

This OCP was transformed in a finite dimension nonlinear program and solved using GPOPS-II. The problem was split into multiple phases that made it easy to include instantaneous distance based changes in speed limit and stopping instances, in the same way as in Chapter 4. The route incorporates both urban and highway driving and, depending on the driving direction, different speed limits are encountered leading to differing power requirements. This allows for consideration of the early route power demands on optimal abatement strategies.

Fuel optimal driving is solved for both routes for charge sustaining scenarios. This set of boundary constraints meant that the energy store is not heavily discharged, leading to prolonged ICE use. The problem formulation includes the ability to adjust engine switching, based on penalising the gradient of the engine speed trajectory. Care is needed when choosing the weighting between this and the fuel consumption. If it is made too large, the engine can no longer switch on and off, if it is too small then the optimiser chooses to rapidly switch the engine. This strategy is impractical and could lead to premature engine fatigue.

Limits on the total emissions production are introduced to evaluate the fuel optimal means to conform to such limits. The sensitivity of the solution trajectories to changes in the total emissions is quantified in terms of powersplit changes and driving changes. It appears that only in the case of HC emissions it is optimal

to decrease the time it takes to heat up the catalytic converter. In the case of NO_x and CO emissions, a low engine power is required which decreases the flow rates of each species during warmup.

Fuel consumption penalties arise for all emissions reductions apart from a spurious case when a 10% reduction in HC was required on B2. The solution for this optimisation was to drastically alter the switching strategy leading to a reduction in the fuel consumption. The relatively high efficiency of the energy storage and electrical machines ensures there are minimal additional losses. In practice, this method could be used to evaluate the requirements of the vehicle based on the initial SOC and the route (in terms of estimated power demands and arrival time) about to be driven in order to select a suitable power-split strategy.

Driving changes can occur which reflect the low or high demanded engine power. In all cases, the average velocity remains the same, which limits the flexibility to alter the drive cycle. By comparison with Chapter 4, the changes to the solution for this route/chapter are considerably smaller.

Despite the fact that the introduction of a catalytic converter has decreased the overall emissions, large reductions in the total emissions produced can be achieved by using a hybrid powertrain.

The key limitation of this method for solving the problem is highlighted when emissions constraints are introduced. The engine switching strategy can change drastically, leading to apparent reductions in fuel consumption over baseline quantities. The car is assumed to not remain stationary for any length of time. This results in no time for the TWC to cool during powertrain shut-off. It also does not account for the ability to charge the energy store while the vehicle is stationary, further improving the heating time. This is infeasible within the chosen distance based framework. Further splitting of the problem into a phases corresponding to the vehicle at rest could offer a means to overcome this. Care would be needed to ensure that the independent variable (time or distance) was consistent with the rest of the problem formulation.

The novelty of this framework, and the subsequent results, suggest that there is capability for this method to be used in the future for longer routes. Further investigations could consider the effects of initial and final SOC constraints which would reflect the discharging capabilities of a plug-in hybrid vehicle.

I build engines and put wheels on them

— Enzo Ferrari, Unknown

6

Optimal Control of Gasoline Engine Combustion Mode Switching

Contents

6.1 Engine Model 160
6.1.1 Cylinder Model 166
6.1.2 Emissions formation 171
6.1.3 Aftertreatment Devices 172
6.2 Optimal Control Formulation 173
6.2.1 Scaling 175
6.3 Optimisation Results 176
6.3.1 Constant power 176
6.3.2 Drive Cycle 181
6.4 Conclusions 184

THUS far, consideration has been given to simultaneous fuel optimal driving and powertrain control. The engine operating trajectories which result from these optimal control problems can be utilised in the creation of optimal engine calibrations. These could augment existing drive cycle information for use in the development of engine and powertrain controllers. One area of interest for improving fuel economy and emissions production is lean combustion engines (for example [8]). Conventional gasoline engines operate at a stoichiometric air-to-fuel ratio (AFR). This ensures complete combustion and good power delivery. Lean combustion

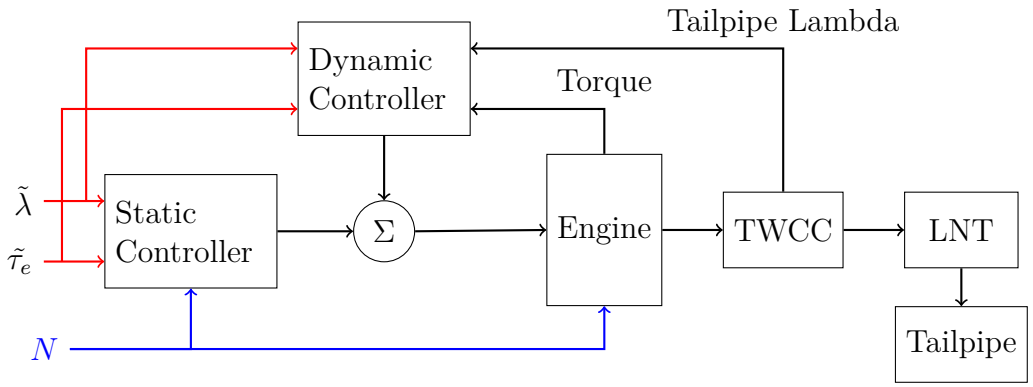


Figure 6.1: Engine controller and signal flow layout. Engine speed N is a measured input to the system and is governed by the vehicle speed. Higher controllers determine the requested torque, $\tilde{\tau}_e$ and tailpipe λ (shown in red).

engines are equipped with systems that allow the operation of the engine at higher air to fuel ratios. This enables more efficient combustion at part loads. However, the three way catalytic converters (TWC) are most efficient at converting the exhaust gases when the engine is operated in a stoichiometric mode (see Figure 6.11). This coupled with increases in NOx due to increased cylinder temperatures can require the use of addition aftertreatment such as a lean NOx trap (LNT). NO_x is trapped in the substrate and released during periods of richer combustion, where the oxygen content of the exhaust stream is reduced. One of the key challenges for control in these engines is in managing the transitions between lean and stoichiometric combustion modes. A two output control problem is formed: ensure the torque output is unaffected, and undetected by the driver, while also ensuring timely transition in tailpipe air quality to avoid emissions spikes.

Figure 6.1 shows a typical engine controller layout [10]. The onboard engine control unit (ECU) consists of stationary lookup maps, with some transient compensation.

In this study, feedback variables are the tracked torque and tailpipe oxygen content (measured by the tailpipe λ sensor).

Generation of the static part of the engine calibration requires the creation and tuning of steady state lookup maps. Historically, when the engine had few control inputs, this was achieved on the test bench by a calibration engineer [10].

System complexity has increased along with the number of controlled actuators, making it difficult to conceive all possible interactions between inputs and outputs. This has also increased the requirement for engine test bed time. Additionally, it is no longer clear how to create an engine calibration in order to maintain fuel economy and conform to emissions limits.

Model-based calibration, together with offline optimisation routines, offers a reasonable solution to this problem. Steady-state engine testing requires a gridded input space to be populated efficiently to reduce the number of tests. Design-of-Experiments (DoE) can reduce the required testing space and the number of tests and, using good modelling techniques, achieve reasonable accuracy for interpolation and extrapolation. High dimension statistical models are a current state of the art for generating static engine calibrations [72]. Neural networks are another means of utilising optimisation for generating engine models suitable for calibration [86]. Feedforward control maps can be created using these static engine models over drive cycle scenarios to minimise fuel consumption and emissions generation.

Dynamic controller creation is more complex. Closed loop controllers, such as PID, are commonly used in combination with the feedforward part of the control. Advanced control methods such as gain-scheduling and model predictive control are constrained by online capabilities. Other advanced techniques exist such as Lyapunov-based controllers [87], and linear-parameter varying (LPV) $\mathcal{H}_\infty$ [88] and LQR [89]. Model predictive control is utilised in order to effectively track the AFR target for a lean combustion engine in [90]. A naturally aspirated engine is the subject of this study, transport delays in measurement of AFR at the sensor in the tailpipe is a key limitation that needs to be overcome.

The design and validation of controller rely on models such as *mean value models*. These models appear extensively in the literature, particularly for turbocharged diesel engines [91–96]. Engine models of this nature incorporate filling/emptying dynamics of air-spaces, turbocharger dynamics and thermal models for aftertreatment devices. Physics based models are a good descriptor for these dynamic quantities.

Dynamic optimal control techniques, such as those discussed already in this thesis, can be applied to models of this type to compute optimal control trajectories for diesel engines [14, 97]. The results of such optimisation provide benchmarks for suboptimal online controller designs. Improvements to existing controller structures can be obtained in terms of fuel efficiency and emissions generation. The optimal control of lean combustion gasoline engines, using dynamic programming was the subject of [98].

This final chapter is focused on utilising dynamic optimal control to examine the available improvements during transient combustion mode switches. A mean value model is combined with physical and phenomenological models for cylinder processes for a turbocharged engine equipped with variable valve timing. Gaussian process models (the same as Chapter 4) are used to predict emissions flow rates and the mass flow rate of air through the engine block. Simple TWC and LNT models are imposed on the tailpipe gas flows in order to predict the tailpipe air quality. The model is parametrised using a representative data set. An optimal control problem is formed which minimises tracking error in output torque and tailpipe λ for the dynamic portion of the engine controller. A feedforward controller is utilised in the loop. Two test cases are optimised over: a constant power output regeneration event and a portion of a test cycle generated using the methods of Chapter 4. The dynamic optimal control problem is solved using the global pseudospectral method in the GPOPS-II optimal control suite. The resulting optimised trajectories for all control inputs determine how a dynamic controller should operate during these transients. This work represents an important step in the creation of closed loop controllers.

The results of this chapter were presented at the Symposium on Combustion Control, Aachen, Germany, June, 2018.

6.1 Engine Model

The model detailed here takes a similar form to those found in the literature. Model data is taken from the Simulink Powertrain blockset for a turbocharged gasoline engine [99]. If the model is to be used for optimisation, it needs to be

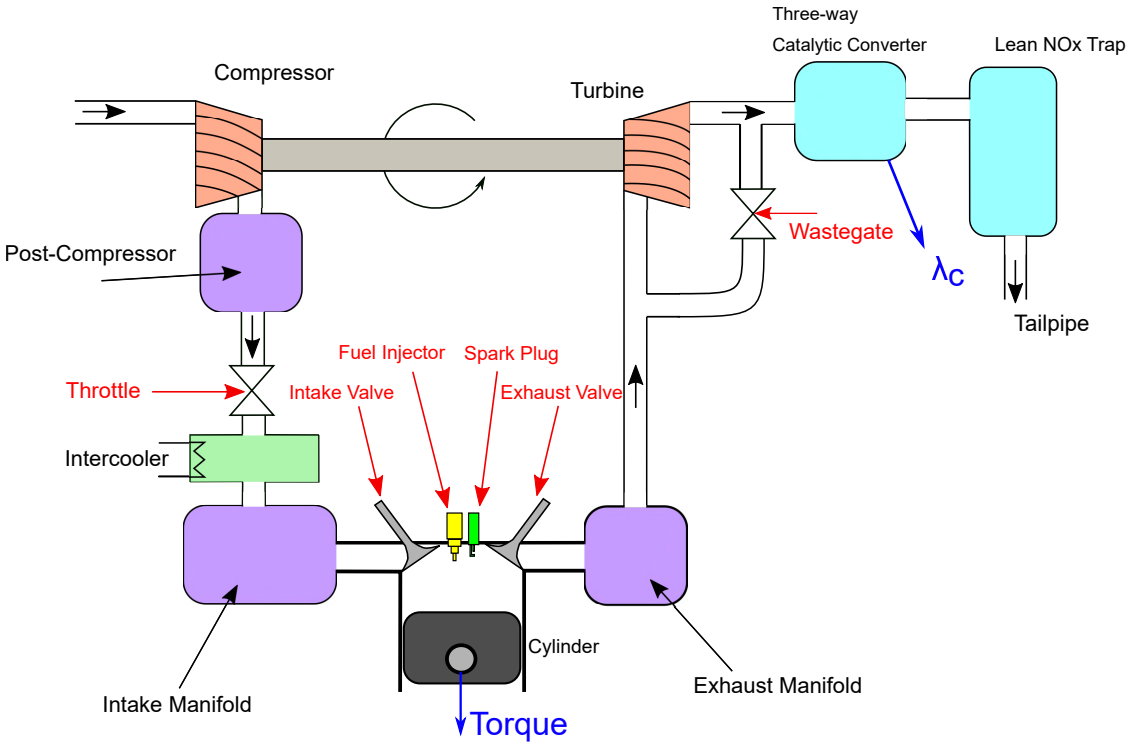


Figure 6.2: The engine model layout. Control input signals are in red, output signals are in blue.

computationally efficient to evaluate and have smooth and computable partial derivatives. Dynamic effects in the air path side of the engine are dominated by the pressure filling and emptying in each of the air spaces and the turbocharger rotational dynamics. The engine layout is shown in Figure 6.2.

Temperature dynamics are notoriously slow and can, for the most part, be ignored [91]. Three gas state equations can be formed for: post compressor pressure p_c; intake manifold pressure p_i; and exhaust manifold pressure p_e (see Figure 6.2). Each of these spaces is large enough to have slow pressure dynamics associated with it. The dynamics for each of these read:

$$\dot{p}_c = \frac{RT_c}{V_c}(\dot{m}_c - \dot{m}_{th}) \tag{6.1}$$

$$\dot{p}_i = \frac{RT_i}{V_i}(\dot{m}_{th} - \dot{m}_a) \tag{6.2}$$

$$\dot{p}_e = \frac{RT_e}{V_e}(\dot{m}_e - \dot{m}_{exh}) \tag{6.3}$$

Post-compressor pressure dynamics (6.1) are determined by the temperature outlet from the compressor T_c, the mass flow rate of air through the compressor $\dot{m}_c$ and the throttle $\dot{m}_{th}$. The intake manifold is separated from the post-compressor plenum by the throttle and intercooler. The pressure dynamics in the intake (6.2) are defined by the temperature out of the intercooler T_i, the mass flow rate of air through the throttle and the aspirated mass of air into the engine $\dot{m}_a$. Finally, the exhaust manifold pressure dynamics (6.3) are dependent on the mass flow rate of exhaust gas out of the engine $\dot{m}_e$ and the flow rate of gas through the exhaust $\dot{m}_t$. Each of the pressure states is present in their corresponding volume spaces for post-compressor V_c, intake manifold V_i and exhaust manifold V_e.

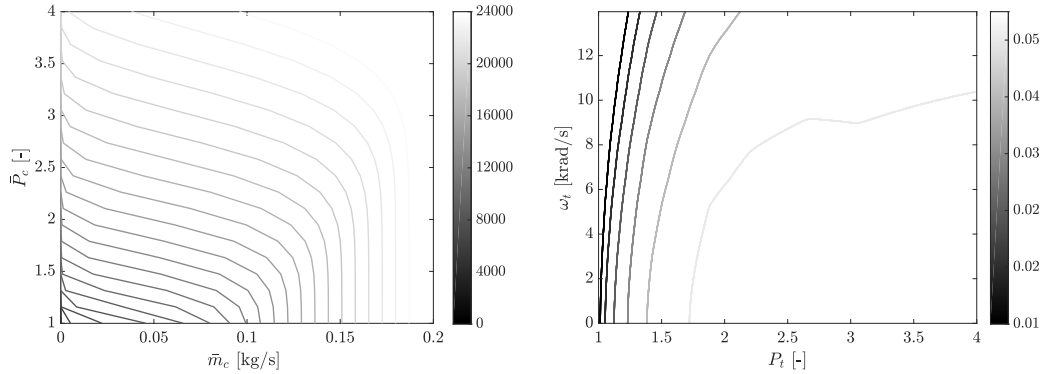
Acting on a similar timescale to the pressure dynamics, the inertial dynamics of the turbocharger are required to complete the air-side model. Taking a rotational dynamic balance along the shaft between the torque created by the turbine τ_t and that required by the compressor τ_c gives:

$$\dot{\omega}_t = \frac{\tau_t - \tau_c}{I_t} = \frac{W_t - W_c}{I_t \omega_t} \tag{6.4}$$

I_t is the rotational inertia of the shaft. The turbine and compressor powers, W_t, W_c can be readily computed based on the models created later in this chapter.

Compressor

The modelling of automotive compressors has been studied extensively in the literature [100–102]. In general, elliptical curves are fitted to both speed and efficiency contours. The model used here is based on the one in [102]. Surge and choke are the two extremes of compressor operation and are difficult to model due to their unstable and undesirable characteristics. For the sake of simplicity, it is assumed that the compressor operates away from either of these extremes of operation. The compressor mass flow rate is determined by the rotational speed of the compressor and the pressure ratio across it. Additionally, the map created is corrected based on a test pressure and temperature which accounts for differing ambient conditions which may affect the characteristic. The turbocharger corrected



(a) The model compressor map. Contours are compressor shaft speed [rad/s].

(b) The model turbine map. Contours are turbine mass flow rate [kg/s].

Figure 6.3: Turbocharger maps.

shaft speed depends on the ambient temperature T_a feeding the compressor, relative to the reference temperature at which the test data is taken:

$$\overline{\omega}_t = \frac{\omega_t}{\sqrt{T_a/T_{ref}}} \tag{6.5}$$

The compressor mass flow rate is similarly corrected based on the inlet pressure p_a and temperature:

$$\overline{\dot{m}}_c = \dot{m}_c \frac{\sqrt{T_a/T_{ref}}}{p_a/p_{ref}} \tag{6.6}$$

The compressor model equations read:

$$P_s = k_{00} + k_{01}\overline{\omega}_t^{k_{02}} \tag{6.7}$$

$$K_1 = k_{10} + k_{11}\overline{\omega}_t + k_{12}\overline{\omega}_t^2 \tag{6.8}$$

$$\overline{\dot{m}}_c = (k_{20} + k_{21}\overline{\omega}_t)\left(1 - \left(\frac{P_c}{P_s}\right)^{K_1}\right)^{1/k_{22}} \tag{6.9}$$

where $P_c = \frac{p_c}{p_a}$ is the compressor pressure ratio. The parameters k_{ij} are optimised using non-linear least squares on the data set. The resulting map is plotted in Figure 6.3a.

Compressor power can be determined based on the pressure ratio across the compressor:

$$W_c = \dot{m}_c c_p T_a \frac{1}{\eta_c}\left(P_c^{\frac{\gamma-1}{\gamma}} - 1\right) \tag{6.10}$$

The compressor isentropic efficiency η_c is assumed constant. The temperature rise across the compressor can also be calculated from the isentropic efficiency η_c:

$$T_c = T_a(1 + \frac{1}{\eta_c}\left(P_c^{\frac{\gamma-1}{\gamma}} - 1\right) \tag{6.11}$$

The ratio of specific heat capacities is γ.

Throttle valves

The intake throttle and wastegate are modelled as restrictions in the piping, with variable areas which form part of the control actuators for the engine. Consider first the intake side throttle; the mass flow rate through which is determined by:

$$\dot{m}_{th} = A_{th}(u_{th})\frac{p_c}{\sqrt{RT_c}}\Phi(P_{th}) \tag{6.12}$$

The mass flow rate is determined by the conditions upstream, p_c and T_c, together with the throttle expansion ratio, $P_{th} = \frac{p_i}{p_c}$, and opening area $A_{th}(u_{th})$. The mapping of the control signal u_{th} to the orifice area is linear in this study. Under normal (sub-sonic) flow conditions the flow rate correction function takes the form

$$\Phi(P_{th}) = \sqrt{\frac{2\gamma}{\gamma-1}\left(P_{th}^{\frac{2}{\gamma}} - P_{th}^{\frac{\gamma+1}{\gamma}}\right)}. \tag{6.13}$$

The function Φ saturates when the flow becomes choked. Below this pressure ratio, the mass flow rate cannot increase further. Choked flow occurs at the critical expansion ratio

$$P_{th}^{crit} = \sqrt{\left(\frac{2}{\gamma+1}\right)^{\frac{\gamma}{\gamma-1}}} = 0.5283, \tag{6.14}$$

corresponding to

$$\Phi(P_{th}^{crit}) = \sqrt{\gamma\left(\frac{2}{\gamma+1}\right)^{\frac{\gamma+1}{\gamma-1}}} = 0.6847. \tag{6.15}$$

A further simplification can be used [103] which increases model simplicity but maintains accuracy:

$$\tilde{\Phi}(P_{th}) \approx \sqrt{A_0 P_{th}(1 - P_{th})} \tag{6.16}$$

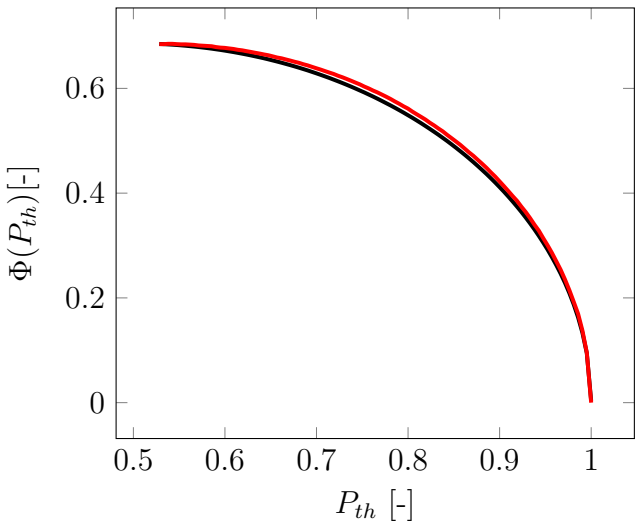


Figure 6.4: A comparison of the true flow rate correction function in red (Equation (6.13)), against the approximation function in black (Equation (6.16)).

The parameter A_0 is chosen such that the choked expansion ratio gives the correct saturation value. A comparison of both pressure ratio functions is plotted in Figure 6.4.

In order to increase the density of the intake manifold air, the charge air is cooled in the intercooler with effectiveness ε. This is simply a heat exchanger transferring heat from the intake air stream to incoming ambient air. The intake manifold temperature is

$$T_i = \varepsilon T_a + (1 - \varepsilon) T_c \tag{6.17}$$

It is assumed that there is no pressure drop across the intercooler.

Turbine

The mass flow through the turbine is a function of both the expansion ratio $P_t = \frac{p_a}{p_t}$ and the normalised rotational speed.

$$\dot{m}_t = a_0 \overline{\omega}_t + a_1 (1 - a_2 e^{-a_3 P_t}) \tag{6.18}$$

Again the model parameters are estimated using nonlinear least squares. The resulting map is plotted in Figure 6.3b. The turbine power is a function of the exhaust manifold temperature T_e, the flow rate through the turbine m_t and the

expansion ratio across it P_t

$$W_t = \dot{m}_t c_p T_e \eta_m \eta_t \left(1 - P_t^{\frac{\gamma-1}{\gamma}}\right). \tag{6.19}$$

Similarly to the compressor, the turbine isentropic efficiency η_t is assumed constant. The other control variable in the pressure charging side of the airpath regulates the turbocharger speed by diverting some of the exhaust gas past the turbine. As the wastegate is a throttle, the same flow rate governing equation can be used. Inlet conditions are replaced with those in the exhaust manifold - p_e, T_e. Again the control signal to orifice area mapping is linear $A_{wg}(u_{wg})$. The resulting mass flow rate is:

$$\dot{m}_{wg} = A_{wg}(u_{wg}) \frac{p_e}{\sqrt{RT_e}} \Phi(P_t) \tag{6.20}$$

The total flow down the exhaust pipe is then the sum of the flows through the turbine and the wastegate:

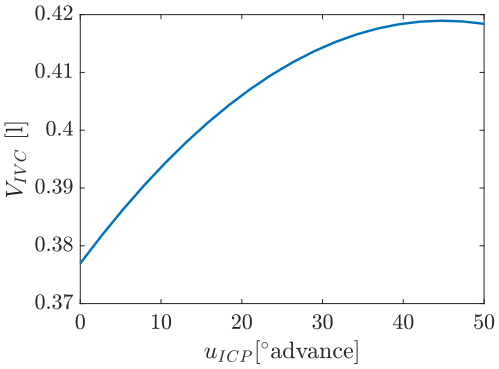
$$\dot{m}_{exh} = \dot{m}_{wg} + \dot{m}_t. \tag{6.21}$$

6.1.1 Cylinder Model

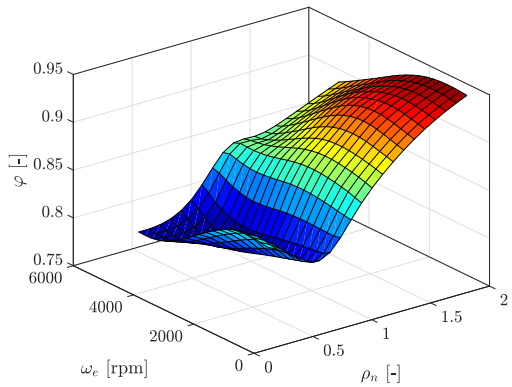
The engine considered in this study is equipped with independent variable exhaust and intake valve timings. Model parameters and structure are taken from [99]. Airflow effects in cylinder are difficult to model based on physics. For example, complex flow interactions are possible when valves overlap. This part of the model relies on empirical data and uses Gaussian processes for correction functions. The mass of air is estimated by considering the ideal gas law for the air trapped when the exhaust valve closes

$$m_{trap} = \frac{p_i V_{IVC}}{RT_i} \tag{6.22}$$

The volume of the cylinder at intake valve closing (IVC), V_{IVC}, is a function of the intake valve timing, measured as degrees advance from top-dead-centre (TDC) (u_{ICP}). The relationship between the two is described by a quadratic polynomial, plotted in Figure 6.5a.



(a) The volume of the cylinder as an adjusted function of the intake cam phasing.



(b) Corrected Mass

Figure 6.5: Intake side empirical functions taken from [99].

A normalised density is defined

$$\rho_n = \frac{p_i T_a}{T_i p_a}. \tag{6.23}$$

This is the ratio of gas density in the intake manifold to that of the aspirated air before the compressor. This is similar to the idea of volumetric efficiency [104].

The trapped mass correction factor φ, is an empirical correction based on the normalised density and the engine speed. This goes some way to account for engine speed-based cylinder pulsations.

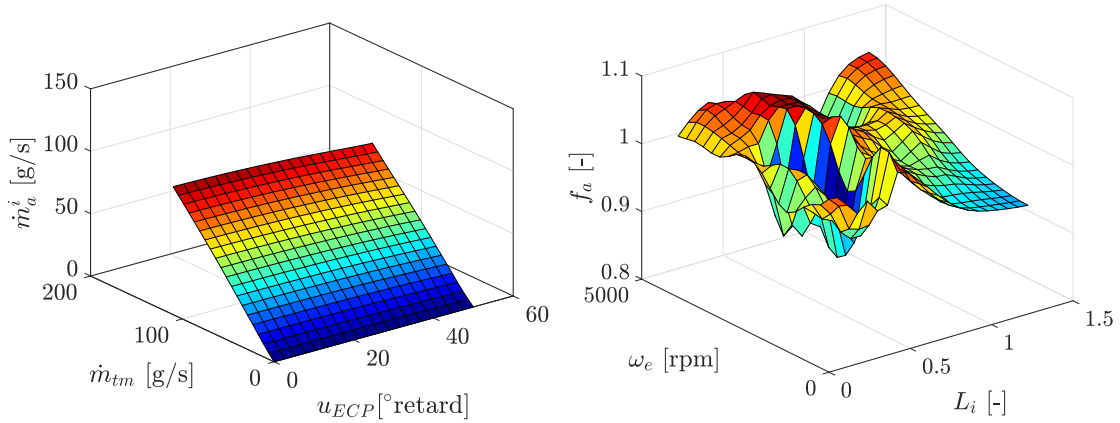
$$\dot{m}_{tm} = \frac{\varphi(\rho_n, N_e) m_{trap} N_e}{30} \tag{6.24}$$

The exhaust cam phasing determines the ideal flow of air through the engine block

$$\dot{m}_a^i = \dot{m}_a^i(u_{ECP}, \dot{m}_{tm}). \tag{6.25}$$

At steady state exhaust cam phase (ECP) angles, the air mass flow is corrected once more, based on the engine speed and the ideal engine load. Ideal load, L_i, refers to the normalised cylinder air mass - the mass of air aspirated compared to that at the ambient conditions, filling the displacement volume V_d,

$$L_i = \frac{0.12 \dot{m}_a^i T_a R}{N_e V_d p_a} \tag{6.26}$$



(a) Dependence of the trapped mass on the exhaust cam phasing (ECP).

(b) The volume of the cylinder as an adjusted function of the intake cam phasing.

Figure 6.6: Maps of engine mass flow rate corrections taken from [99].

Finally, the air mass flow is corrected with the map in Figure 6.6b which is modelled with a Gaussian process model.

$$\dot{m}_a = \dot{m}_a^i f_a(L_i, N) \tag{6.27}$$

The engine load is also corrected based on this function:

$$L = L_i f_a(L_i, N) \tag{6.28}$$

Fuel Flow

The control input for fuelling is the pulse width of injection, u_{inj}, which together with the engine speed, determine the mass flow rate of fuel $\dot{m}_f$,

$$\dot{m}_f = BN_e u_{Inj} \tag{6.29}$$

The in-cylinder AFR is ratio of mass flow rate of air to fuel flow rate

$$AFR = \frac{\dot{m}_a}{\dot{m}_f} \tag{6.30}$$

resulting in the air-fuel equivalence ratio

$$\lambda_e = \frac{AFR}{AFR_s}. \tag{6.31}$$

Where AFR_s is the stoichiometric AFR (14.7 for a gasoline engine). $\lambda_e > 1$ implies lean combustion, conversely $\lambda_e < 1$ implies rich combustion.

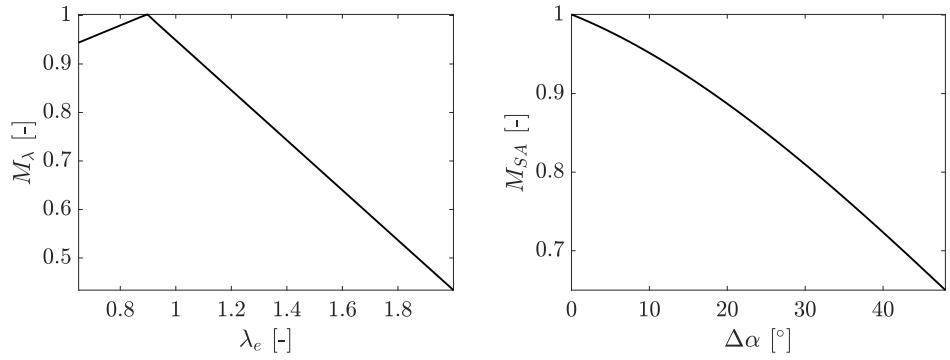


Figure 6.7: Torque lambda efficiency.

Torque Model

Torque is predicted based on a model similar to that presented in [105]. Inner torque is predicted based on the engine speed and load. This is degraded by two efficiency factors: one compensating for changes in the air/fuel ratio and the other compensating for adjustments to the spark timing.

The brake torque τ_b, is therefore given by

$$\tau_b = M_\lambda M_{SA} \tau_{ind}(N_e, L) - \tau_d(N_e, L). \tag{6.32}$$

The indicated torque is an affine function of the load, with a small dependency on the engine speed, plotted in Figure 6.9a,

$$\tau_{ind} = g_0 L(1 - e^{-g_1 N_e}). \tag{6.33}$$

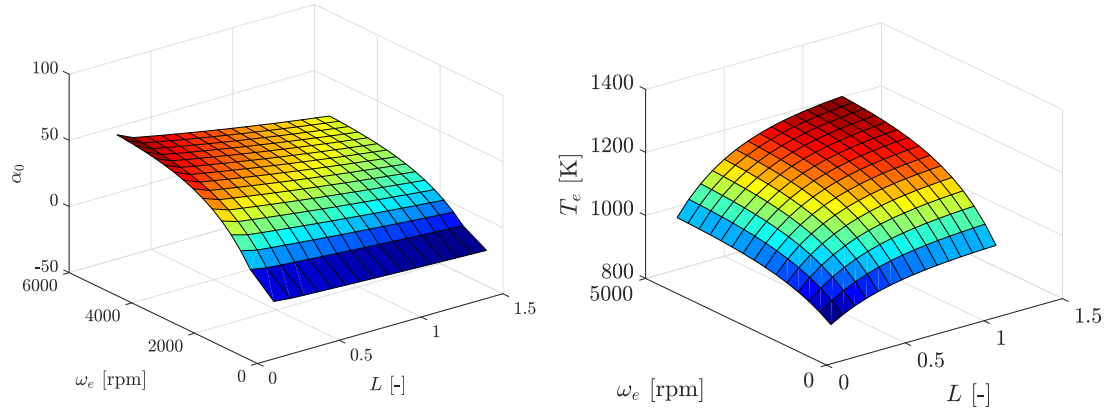
Optimal spark timing is given by

$$\alpha_0(N_e, L) = c_0 + c_1(1 - e^{-c_2 L})(-e^{-c_3 N_e}). \tag{6.34}$$

The timing of the ignition is made earlier (advanced) than optimum in order to reduce the torque output of the engine. The spark advance is

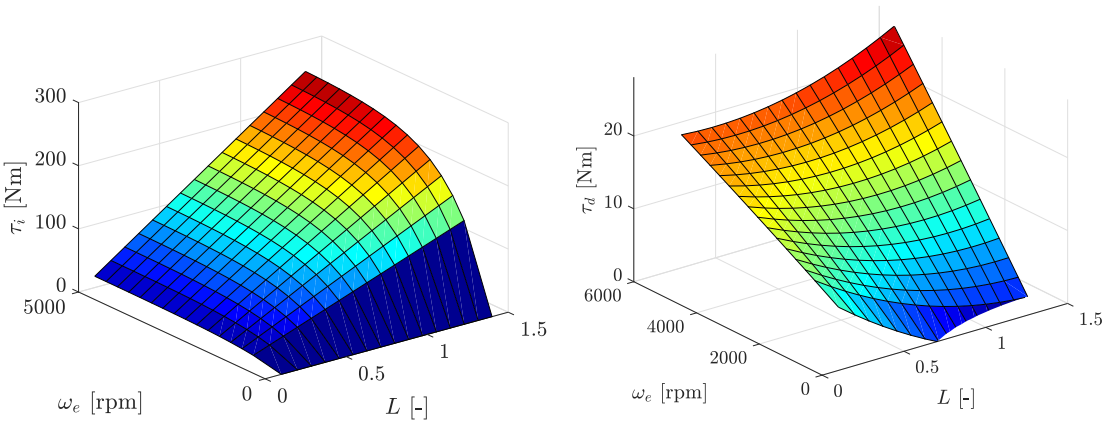
$$\Delta\alpha = \alpha_0(N_e, L) - u_{SA} \tag{6.35}$$

The spark-advance efficiency, M_{SA} is a third order polynomial of $\Delta\alpha_{SA}$ and is plotted in the RHS of Figure 6.7. The efficiency due to in-cylinder λ, denoted M_λ takes a piecewise linear definition (LHS of Figure 6.7), with maximum torque created



(a) Optimum spark timing map. **(b)** Exhaust temperature.

Figure 6.8: Engine maps for temperature and spark timing.



(a) Indicated Torque map for engine **(b)** Friction torque map for the engine.

Figure 6.9: Engine torque maps

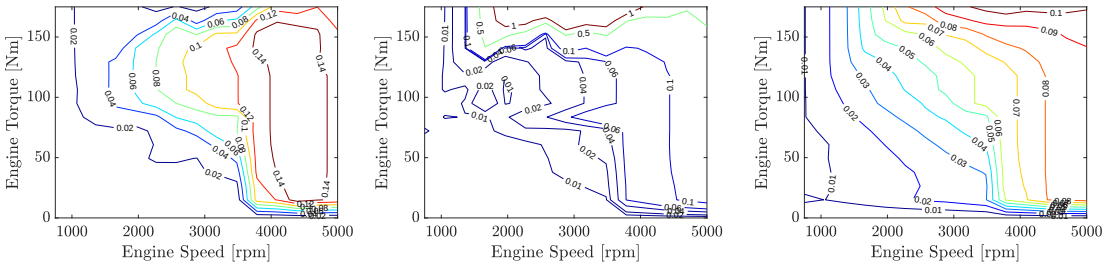
slightly rich of stoichiometric combustion (around $\lambda_e = 0.9$). The drag torque is a quadratic mapping of both engine load and speed and is also plotted in Figure 6.9b.

$$\tau_d = \tau_{fric}(L, N_e). \tag{6.36}$$

Finally the exhaust temperature model takes the form:

$$T_e = n_0 + n_1(1 - e^{-n_2 N_e})(1 - e^{-n_3 L}) \tag{6.37}$$

All of the parameters of this section are fitted to the data set (in [99]) using nonlinear least squares.



(a) Mass flow rate of NO_x. **(b)** Mass flow rate of CO. **(c)** Mass flow rate of HC.

Figure 6.10: Engine emissions maps, contours are mass flow rates [g/s].

6.1.2 Emissions formation

Three major types of emissions are considered in this work: the oxides of nitrogen (NO_x), carbon monoxide (CO) and unburnt hydrocarbons (HC). Contrasting conditions about the stoichiometric AFR result in trade-off production for each of these species. CO formation is most favourable under rich combustion conditions. Incomplete combustion prevents full oxidation to CO_2, due to the lack of oxygen. HC formation is reduced with the presence of more oxygen (increasing λ). This effect increases until the reduced flammability of the mixture causes incomplete combustion of the fuel. In general, higher flame temperatures are synonymous with slightly lean mixtures and therefore lead to the highest NO_x production.

The production of each of the species is most heavily correlated to the engine speed and brake torque, which includes information regarding the airpath variables through the load maps. Two dimensional response surfaces are created using Gaussian process models (GPM) the same as in Chapter 4. The output data (from [99]) is taken at stoichiometric combustion and the response surfaces are plotted in Figures 6.10a, 6.10b and 6.10c. Changes in emissions formation due to deviations away from stoichiometric combustion take the form

$$\eta_{\lambda CO} = 8e7e^{-20\lambda_e} + 0.85, \tag{6.38}$$

$$\eta_{\lambda HC} = 7.81(\lambda_e - 1.08)^2 + 0.95, \tag{6.39}$$

$$\eta_{\lambda NOx} = \frac{5\lambda_e}{8} e^{-(\lambda_e - 1.05)^2/32} + 0.39. \tag{6.40}$$

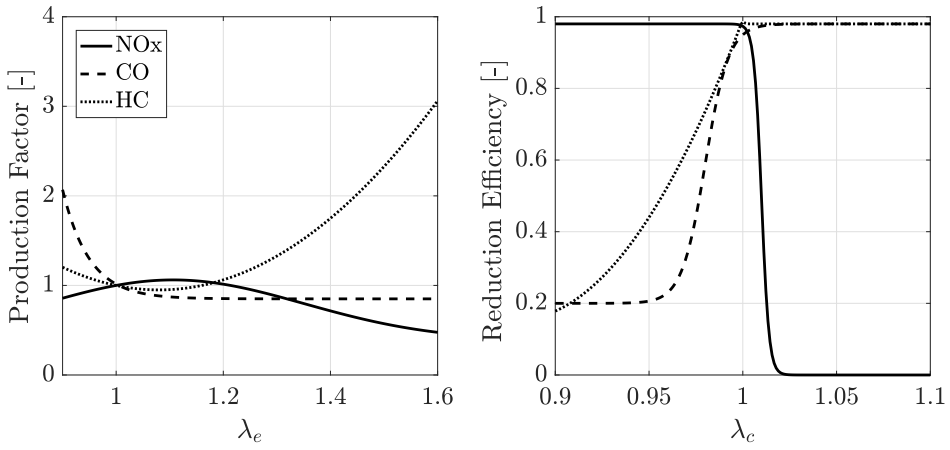


Figure 6.11: Left: The dependence of engine out emissions on in-cylinder λ_e. Right: The dependence of TWC reduction efficiency on the local air quality λ_c.

These are plotted in the left hand side of Figure 6.11. The total engine out emissions are the product of these functions and the output of the GPM predictions.

$$\dot{m}^e_{\mathrm{CO}} = \eta_{\lambda\mathrm{CO}}\mathrm{GPM}(N_e, \tau_b), \tag{6.41}$$

$$\dot{m}^e_{\mathrm{NO_x}} = \eta_{\lambda\mathrm{NO_x}}\mathrm{GPM}_{\mathrm{NO_x}}(N_e, \tau_b), \tag{6.42}$$

$$\dot{m}^e_{\mathrm{HC}} = \eta_{\lambda\mathrm{HC}}\mathrm{GPM}_{\mathrm{HC}}(N_e, \tau_b). \tag{6.43}$$

6.1.3 Aftertreatment Devices

The three-way catalytic converter (TWC) reduces the tailpipe quantities of HC, CO and NO_x. CO and HC are oxidised into CO_2 and H_2O while the NO_x is reduced to N_2 . Each of the reactions is most favourable at high temperatures in the TWC. This relies on the substrate reaching the *light-off* temperature where above 50% reduction can occur. In this study, it is assumed that the engine has sufficiently warmed the catalytic converter. In addition to the temperature characteristic, there is also a dependence for emissions reduction based on the quantity of oxygen in the exhaust stream. This is related to the local air/fuel ratio λ_c according to the curves in the right hand side of Figure 6.11 [76].

There is a first order dynamic describing the transport of gas through the exhaust system to the measurement in the TWC:

$$\dot{\lambda}_c = C(\lambda_e - \lambda_c). \tag{6.44}$$

This local AFR, λ_c, feeds the conversion efficiency equations:

$$\eta_{NOx} = 0.49(1 - \tanh(250(\lambda_c - 1.01))), \tag{6.45}$$

$$\eta_{CO} = 0.98 + 0.39(\tanh(120(\lambda_c - 0.985)) - 1), \tag{6.46}$$

$$\eta_{HC} = 116\lambda_c^2 - 205\lambda_c + 90 - (58\lambda_c^2 + 102\lambda_c - 44.5))(1 + \tanh(10^3(\lambda_c - 1))). \tag{6.47}$$

Finally, the LNT has some trapping efficiency dynamics, which are switched depending on λ_c.

$$\dot{\eta}_{LNT} = -\frac{\eta_{LNT}}{200}\left(1 + \tanh\left(100\left(\lambda_c - 1.05\right)\right)\right) + \frac{1 - \eta_{LNT}}{2}\left(1 - \tanh\left(100\left(\lambda_c - 1.05\right)\right)\right). \tag{6.48}$$

Tailpipe emissions flow rates can therefore be calculated from

$$\dot{m}^t_{\mathrm{CO}} = \dot{m}^e_{CO}(1 - \eta_{CO}), \tag{6.49}$$

$$\dot{m}^t_{\mathrm{NO_x}} = \dot{m}^e_{NO_\mathrm{x}}(1 - \eta_{LNT})(1 - \eta_{NO_\mathrm{x}}), \tag{6.50}$$

$$\dot{m}^t_{\mathrm{HC}} = \dot{m}^e_{HC}(1 - \eta_{HC}). \tag{6.51}$$

6.2 Optimal Control Formulation

The optimal control problem can now be posed. The model is populated with some of the parameters in Table 6.1. This aims to determine optimal interaction of dynamic feedback control paths in order to facilitate better tracking of λ targets while ensuring good torque output. The optimisation has full knowledge of all state and output variables so is expected to outperform a nominal dual loop P+I controller. The steady state part of the controller is optimised beforehand and ensures that the optimisation can be computed quickly.

Table 6.1: Parameter values and descriptions for the gasoline airpath model.

Parameter	Value	Description [Units]
V_c	0.0026	Post-compressor volume [m^3]
V_i	0.0029	Intake manifold volume [m^3]
V_e	0.0016	Exhaust manifold volume [m^3]
V_d	0.0015	Engine displacement volume [m^3]
I_t	1.55×10^{-5}	Turbocharger Inertia [kgm^{-2}]
T_a	298	Ambient temperature [K]
p_a	1	Ambient pressure [bar]
η_c, η_t	0.7, 0.7	Compressor/turbine Isentropic Efficiency [-]
η_m	0.9	Turbo Mechanical Efficiency [-]
c_p	1003	Heat capacity of air [J/K/kg]
γ	1.4	Ratio of Specific Heats [-]
A_{wg}	3.1×10^{-4}	Wastegate opening area parameter [m^2]
A_{th}	1.4×10^{-5}	Throttle opening area parameter [m^2]
ε	0.7	Intercooler Effectiveness [K]
B	2.2×10^{-7}	Fuelling Parameter [kg/s]
C	10	Transport dynamic timescale [s^{-1}]

The cost function is made up of three terms; the first penalising the deviation of each of the outputs (λ_c and τ_e) from their targets, the second penalising the deviations away from the steady-state controller and the final term penalises the derivative of the control action. This final term behaves as a regularisation for the system, helping to avoid chattering controls. The cost function reads:

$$J = \int_{t_0}^{t_f} (\lambda_c(\mathbf{x}, \mathbf{u}) - \bar{\lambda}_c)^2 + r_1(\tau_b(\mathbf{x}, \mathbf{u}) - \bar{\tau})^2 + r_2(\mathbf{u} - \bar{\mathbf{u}})^2 + r_3 \dot{\mathbf{u}}^2 dt \tag{6.52}$$

The dynamic state vector is made up of the post compressor, intake and exhaust pressures, the turbocharger speed, the tailpipe λ_c, the LNT trapping effectiveness, and the three tailpipe emissions species.

$$\mathbf{x} = [p_c, p_i, p_e, \omega_t, \lambda_c, \eta_{LNT}, m^t_{\mathrm{CO}}, m^t_{\mathrm{NO_x}}, m^t_{\mathrm{HC}}]^\top \tag{6.53}$$

The corresponding dynamic equations are (6.1), (6.2) (6.3), (6.4), (6.44), (6.48), (6.49), (6.50) and (6.51).

The output equations required for the cost function λ_c and τ_b are given by (6.44) and (6.32). The model inputs consist of the throttle valve position, the

injection quantity, the wastegate position, the spark timing, and cam phasing for both intake and exhaust valves:

$$\mathbf{u} = [u_{thr}, u_{inj}, u_{WG}, u_{SA}, u_{ICP}, u_{ECP}]^\top, \tag{6.54}$$

An inequality constraint is required which ensures proper combustion throughout the problem horizon

$$0.9 \leq \lambda_e \leq 1.7. \tag{6.55}$$

Further inequality constraints ensure flow is not reversed through the various flow restrictions:

$$p_c > p_i, \tag{6.56}$$

$$p_e > p_a. \tag{6.57}$$

6.2.1 Scaling

In order to make the optimisation problem computationally efficient, it is necessary to remove the notion of large and small decision variables. This could involve creating a set of length, time and mass scales, and use these to normalise the problem data (see Chapter 2). This does not lend itself to good use in the case of the turbocharger rotational speed. To ensure that the scaled version of the state was adequately small, this would force the reciprocal time scale to be large. Consequently, the time scale would be small, resulting in a simulation horizon length much greater than reality. This could prove problematic for optimal control software.

A different approach is proposed whereby four new state variables related to the airpath are defined by scaling the rotational speed state by the turbocharger inertia and pressure states by their volumes. The new scaled portion of the state vector reads:

$$\mathbf{x}' = [p_c', p_i', p_e', \omega_t']^\top = [\frac{p_c V_c}{R}, \frac{p_i V_i}{R}, \frac{p_e V_e}{R}, \omega_t I_t]^\top \tag{6.58}$$

The corresponding scaled dynamic equations read:

$$\dot{p}'_c = T_c(\dot{m}_c - \dot{m}_{th}), \tag{6.59}$$

$$\dot{p}'_i = T_i(\dot{m}_{th} - \dot{m}_a), \tag{6.60}$$

$$\dot{p}'_e = T_e(\dot{m}_{eo} - \dot{m}_{exh}), \tag{6.61}$$

$$\dot{\omega}'_t = I\frac{W_t - W_c}{\omega'_t}. \tag{6.62}$$

6.3 Optimisation Results

The baseline static controller is optimised by setting constant output targets across the operating space of $N_e - \tau_b - \lambda_c$. The static controller then interpolates based on these operating points to determine the steady state control strategy for the engine. The dynamic portion of the baseline controller is tuned for reasonable performance under step changes in reference values. This consists one PI loop measuring λ_c and adjusting the fuelling quantity; the other measures the error in torque output and controls the intake throttle. This is a standard means to dynamically control an engine [10].

6.3.1 Constant power

The first operating condition is a constant power output from the engine, with a five second LNT regeneration period. The total optimisation horizon is 15 seconds. The switch from lean to rich combustion occurs at 5 seconds, the second switch at 10 seconds is the return to lean running. The engine is set to operate at 2000rpm and required to produce 80Nm of torque. Switching modes from lean to rich combustion requires a combination of both increasing the mass flow rate of fuel and decreasing the mass flow rate of air through the engine. This is achieved by decreasing the intake manifold pressure, which will in turn require a decrease in the compressor and exhaust manifold pressures. In order to see how this is achieved, it is necessary to consider both the optimised control trajectories and the optimised states (Figures 6.12 and 6.13).

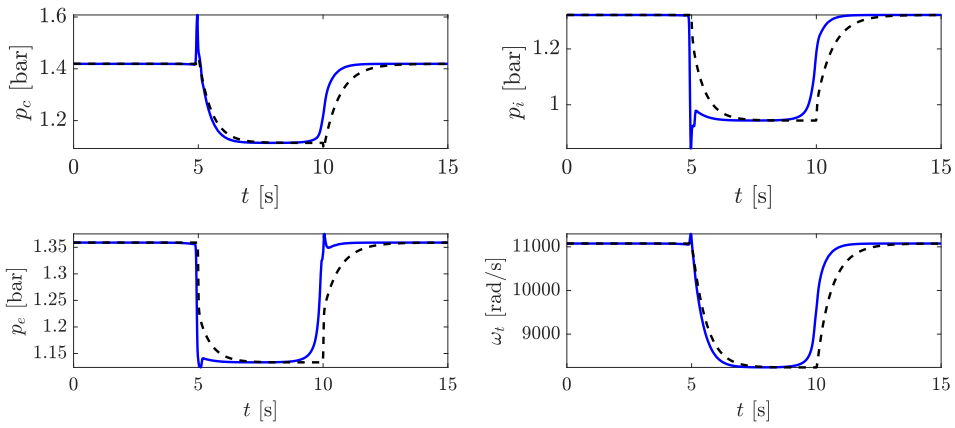


Figure 6.12: Airpath state trajectory comparison for the optimised controller action (solid blue), compared to the baseline (P+I) (dashed black).

This decrease is achieved by lowering the boost that is created by the compressor. As the turbocharger inertia is a slow dynamic state, this cannot be achieved instantaneously, so the intake pressure is decreased by closing the throttle. A side-effect of this is the sudden increase in post-compressor pressure. The wastegate is opened, which slows the turbine and in turn slowly decreases the turbocharger speed state. This only occurs marginally quicker than for the baseline control. During this deceleration, the throttle is slowly re-opened and air flow into the engine is maintained. The intake and exhaust cam phases are not changed from the baseline during the lean to rich step event. As the engine load is maintained correctly as well, the spark-advance channel is not used to compensate for extra torque. The injection quantity signal controls the tailpipe lambda very closely. Due to the delay between measure and actuation, the mass flow rate of air is actually altered before the step request occurs. The injection quantity is adjusted in order to maintain the output tracking.

Larger control actions are required when the engine is requested to transition back to lean running. The turbocharger is required to boost again, increasing the flow rate into the engine. Compared to the lean to rich switch, increasing the intake manifold pressure is more difficult. The turbocharger takes a considerable length of time (relative to the rest of the airpath) to speed up fully. Therefore the wastegate is actuated approximately 1.5 seconds before the step transition.

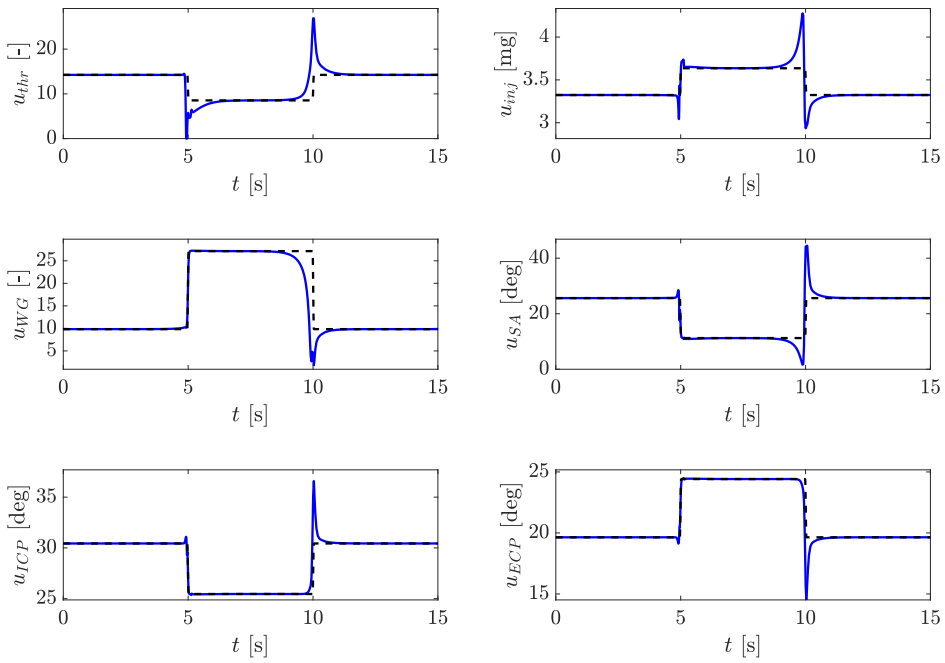


Figure 6.13: Control actions for baseline (black dashed) and optimised (solid blue).

The pressure in the exhaust manifold is also increased by allowing more air to flow through the engine. During the period of rich combustion, this would increase the lambda as the air flow increases. This is compensated for by increasing the injection quantity. At the transition itself, the throttle is instantaneously closed and the fuel injection dramatically decreased. In combination with changes to the timing of both valves, this increases the mass flow of air aspirated by the engine. The spark timing is altered during this transient which suggests the torque tracking is made more difficult.

Indeed, there is a large spike in torque for the baseline controller case, which is clearly reduced with dynamic control of all the actuators Figure 6.14. A large increase in the air-aspirated will cause the spike in torque production. This is effectively reduced by altering the spark timing to degrade torque production.

The behaviour of the air-fuel ratios at various points in the engine are also considered. The baseline case (in black dash) has a lethargic transition from both rich to lean and vice-versa. The in-cylinder λ is driven quickly to force a fast transition

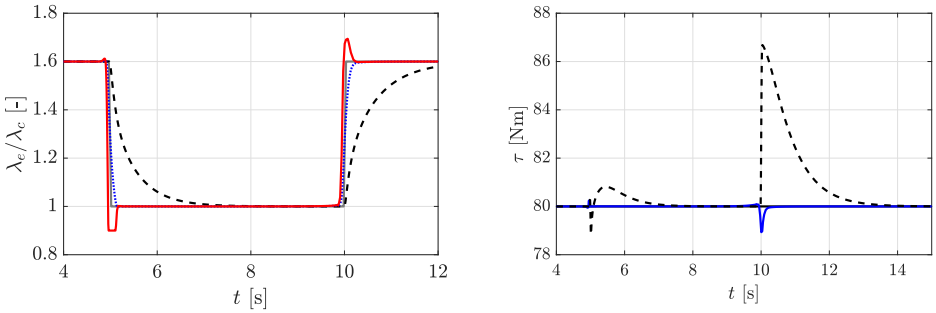


Figure 6.14: Output tracking trajectories for the constant power output case. The left hand figure compares the optimised trajectories for λ_e (red), λ_c blue dotted compared to the baseline (black dashed). The right hand figure compares the torque tracking for the baseline controller (black) and optimised trajectory (blue).

in the tailpipe air (red trace). Importantly, the combustion quality constraint (6.55) becomes active during the transition. This limits how fast the switch-over can occur.

Consider now the emissions flow rates out of the engine and in the tailpipe plotted in Figure 6.15. Poor λ_e switching during the rich to lean transient leads to unnecessary spikes in the NO_x flow rates out of the engine. This is due to the increased production for slightly richer than stoichiometric combustion (λ slightly greater than 1) - see Figure 6.11.

In the optimised case, this phenomenon is removed entirely. At the tailpipe, the transient NO_x is abated considerably. The slow transition between lean and rich during the initial transient causes a wide spike of emissions in the tailpipe NO_x. The LNT is full during this transient; meaning it is necessary to increase the TWC effectiveness by quickly transitioning to stoichiometric (or close to) combustion as quickly as possible. At the switch back to lean combustion at 10seconds, the LNT is now empty and can effectively remove all NO_x emissions from the tailpipe. In terms of NO_x emissions, the total tailpipe emissions are best reduced by quick transitioning to rich combustion. For the test case here, a 38% reduction is achieved.

The HC produced is largely the same between the baseline and optimised trajectory case. In this case, the production of HC is lower for the baseline during the slow switch of the engine lambda back to lean combustion after 10 seconds. This is reflected in the tailpipe flow rates, which are, on the whole, already very low anyway. The TWC operates well under lean conditions for the reduction of HC.

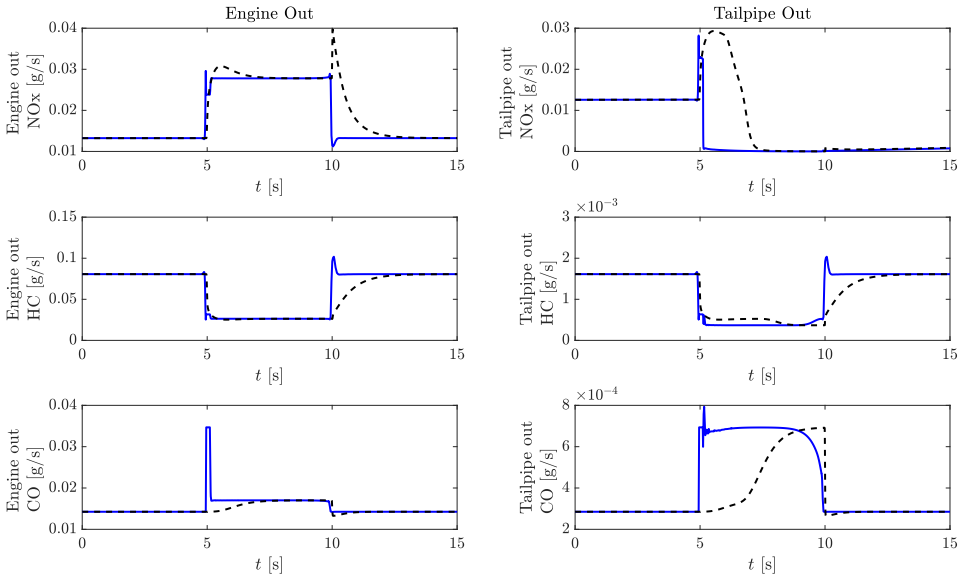


Figure 6.15: Mass flow rates of NO_x, HC and CO (descending); engine out on the left and out of the tailpipe on the right. Response under baseline control is black dashed. Dynamic optimised response is solid blue.

The decrease in effectiveness under rich conditions does not effect the total tailpipe production. Similar results arise for CO emissions in the optimised case. A large spike in CO emissions during the initial transient arises due to the steeper rolloff of TWC effectiveness for λ_c below 1. However, the TWC efficiency at reducing CO prevents this new emissions spike from appearing at the tailpipe. HC and CO total production at the tailpipe increase by 2.4% and 16.8% respectively. While this appears to be detrimental, it should be noted that the actual totals of these quantities are much lower than the NO_x, and as such are more sensitive to slight changes in the flows. The total fuel consumption is increased by less than 0.5%.

This simple test case shows the potential for improvements in the reduction of the oxides of nitrogen at the tailpipe during the regeneration period of the LNT. This comes at the detriment of other emissions species and fuel consumption. However, the absolute value of these quantities was already low due to the lean combustion regime used for most of the operation. The optimisation was solved on a desktop PC with an 8 core 3.9GHz CPU. The problem was initialised with the baseline control action to speed the computation up resulting in a computation time of 386 seconds.

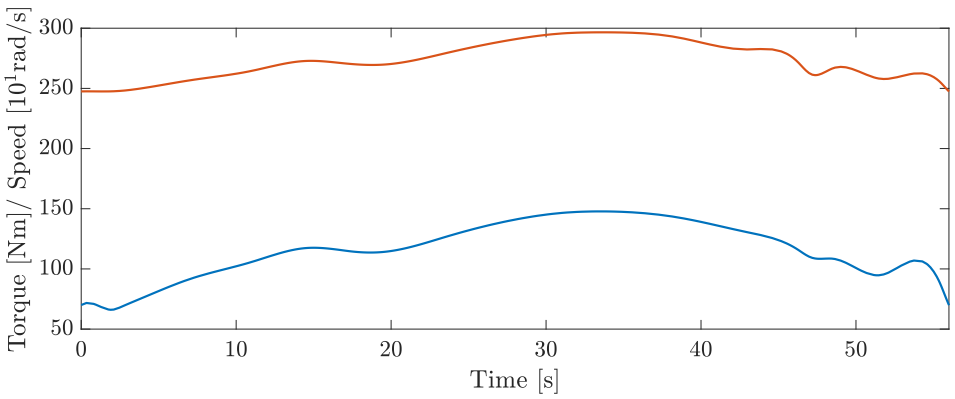


Figure 6.16: The drive cycle engine torque (blue) and speed (red) trajectories for the transient optimisation case.

6.3.2 Drive Cycle

A further test case is considered which represents a more realistic engine operating trajectory. This utilises the methods of Chapter 4 to create the required torque and engine speed characteristics. A fuel optimal solution to the traversal of Route A is computed and a small segment of the resulting trajectories are optimised over.

A high power segment of the resulting trajectories is optimised over. A regeneration event is also imposed during increasing power delivery.

This will show how the previous results of the simple step change relate directly over to the dynamic test case. In the context of engine calibration, this study represents part of the design and iteration loop for designing controllers for engine use.

Optimised control trajectories are compared to the baseline in Figure 6.17.

The transition from lean combustion to rich at 17 seconds is predominantly driven by dynamic throttle use. There is little dynamic control for any of the other control channels, apart from the injection quantity to ensure the λ target is tracked. Similarly, during the transition after regeneration is complete, the wastegate is opened fully to promote an increase in the intake manifold pressure. The throttle is again opened and the valve timings changed to promote higher mass flows of air. The fuelling is adjusted for this high air flow condition to ensure the combustion remains in the rich region.

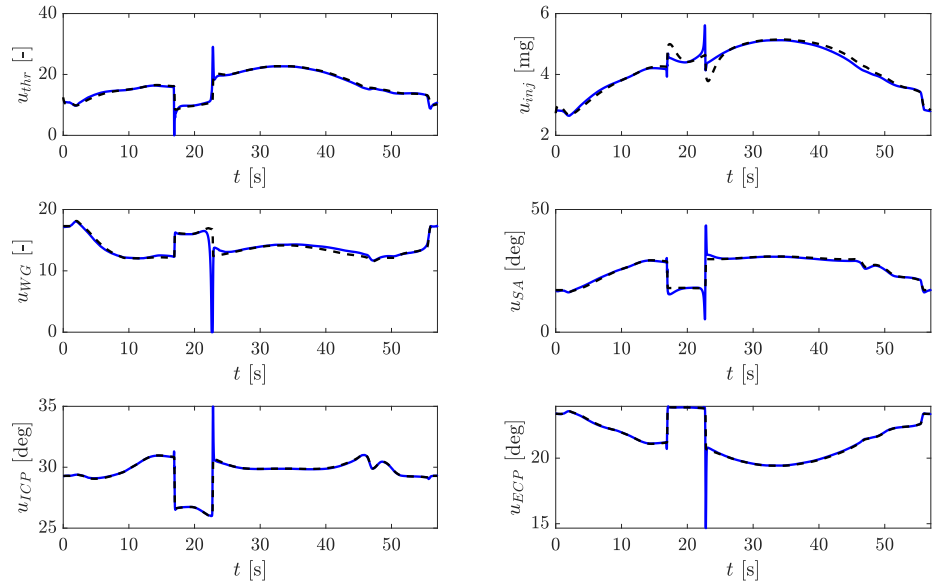


Figure 6.17: Optimal control trajectories (blue) for the drive cycle case compared to the nominal baseline controller (black dashed).

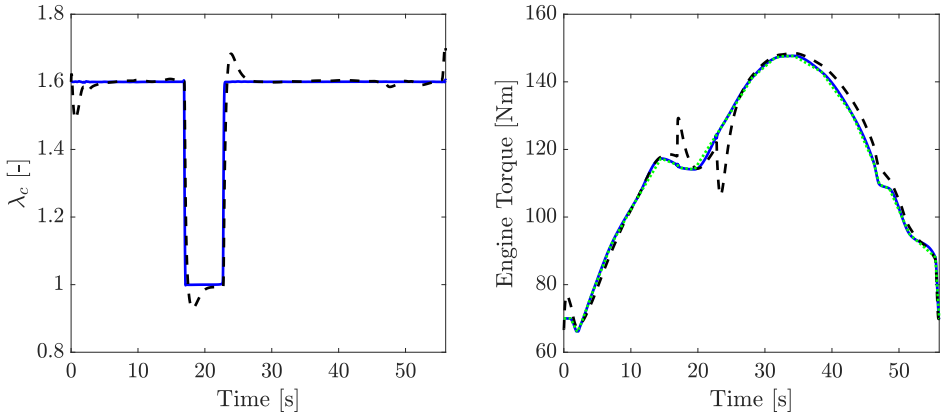


Figure 6.18: Output tracking trajectories - left is the tailpipe λ, blue is optimised, black dashed is baseline. Right is output torque, blue is optimised, black dashed is baseline control. The torque target is plotted as dotted green.

The tracking of outputs is improved in Figure 6.18. Over the baseline control, the following of references is particularly improved during the transient events. Over/under-shoot was seen in both λ_c and engine torque during the transitions between combustion modes, this is all but eliminated with full dynamic control.

The effect of optimising the control on the emissions is compared in Figure 6.19. Despite most of the NO_x being created during the transient, there is still scope

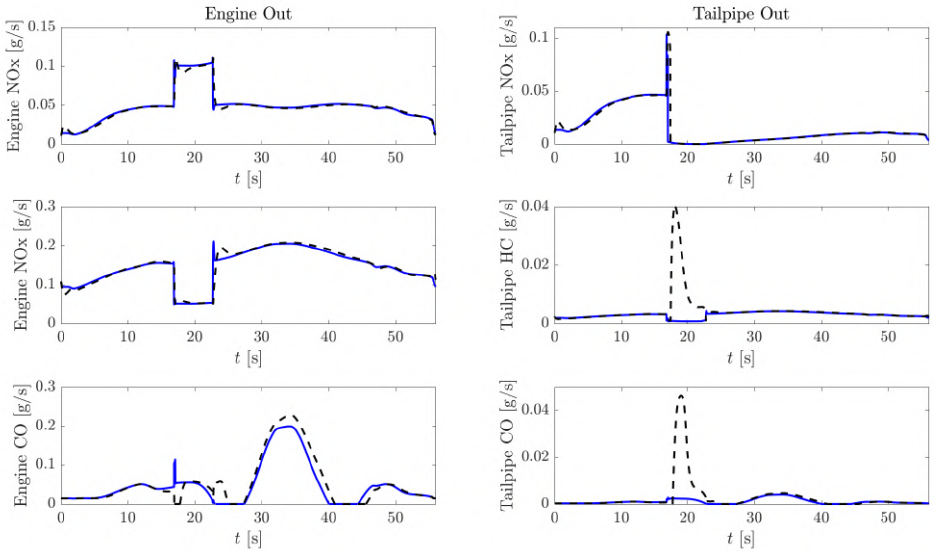


Figure 6.19: Emissions trajectories for the drive cycle case. The optimised trajectories are blue, the baseline controller are black dashed.

for reduction at the tailpipe. Engine out NO_x, is broadly unchanged compared to the baseline. Again, compared to the step switching case, the width of the tailpipe NO_x spike is reduced. This, in total, leads to a 3.7% reduction in the total produced over the test cycle period.

In the case of HC emissions, the engine out emissions are unchanged again. Due to the poor λ_c tracking in the baseline, a spike in tailpipe emissions is seen. The overly rich period of combustion at the start of the regeneration period degrades the effectiveness of the TWC to reduce HC. This results in a reduction of total HC by 33.6%. Finally, CO emissions are reduced, due in part to better torque tracking - there is a strong correlation between the torque produced and the CO emissions. In the same manner as HC emissions, the improvement to tailpipe λ following reduces the tailpipe CO emissions. The total reduction is 62.6%. The same as in the steady state output case, there is a very small increase in the fuel consumption across the horizon. The problem was again initialised with the baseline controller and the total time to solve the problem was 1185 seconds.

6.4 Conclusions

This final chapter was focused on utilising dynamic optimal control to investigate the means of reducing cumulative emissions for lean combustion gasoline engines equipped with lean-NO_x traps. An existing model structure is modified to provide one which is suitable for optimal control formulation with a suitably smooth and accurate model. An eight state model describes the evolution of the dynamic air components in the engine and the aftertreatment and cumulative emissions. The model utilises physics based models for air-spaces and some airpath components such as the intercooler. The compressor model is an elliptical form model found in the literature. Cylinder processes related to torque generation and aspirated air mass are empirically modelled. These models are fitted to representative data from an existing engine model (that found in [99]). Emission production and certain parts of the engine breathing sub-model are modelled with Gaussian process models. The parameters of which are chosen in the same method as Chapter 4.

A typical engine control architecture incorporates a static feedforward controller with a dynamic feedback controller. The optimal control problem in this chapter creates the output of the dynamic portion of the controller.

This chapter utilises offline numerical optimal control to determine optimal operating trajectories for some example test cases. A baseline controller is created which performs reasonably well with heuristic tuning. Improvements to the tracking of output variables enable effective reductions in emissions without harming torque output or increasing fuel consumption dramatically.

The results of this optimisation also highlight one method to improve transient operation. Moving forward, new feedback controllers would implement some of the ideas presented here. As the turbocharger rotational speed dynamic is the slowest in the system, pre-emptive control of this can aid in faster switching in combustion modes. If the air-flow is changed in the engine, then this must be compensated for by the fuelling control. When switching from rich to lean combustion, the most effective method utilises all available actuators to achieve this. To avoid large errors

in tracking, the pressure boosting system in this case needs an increased amount of pre-emptive control compared to switching from lean to rich.

These results are also reflected in the longer driving case with a moving torque target. Better tracking of both control outputs, together with improvements in the transient emissions imply that there are gains to be made over a dual loop, P+I controller.

A story has no beginning or end: arbitrarily one chooses that moment of experience from which to look back or from which to look ahead.

— Graham Greene, *The End of the Affair*

7 Conclusions

THE aim of this thesis was to investigate uses for dynamic optimisation to solve fuel-related and emissions constrained optimisation problems which arise in automotive powertrain control. Specifically, the control of powertrains for optimal driving and transient engine operation were considered. A number of results were found across Chapters 3-6 which are reflected on here.

Chapter 2 presented some of the key concepts required for the later parts of this thesis. A broad discussion of continuous time optimal control, resulting in Pontryagin's Maximum Principle. As the optimal control problems posed and solved in this thesis are highly non-linear and complex, numerical methods for solving these problems were required.

Before tackling the problem of optimal driving on a complex route, a smaller, more easily dissected problem was tackled in Chapter 3. Initially, a car was driven up a slope of constant incline, in the absence and presence of a constant speed constraint. An analytical control solution was derived for this both cases, along with minimum arrival times and minimum energy arrival times. In the case that a speed constraint is present, a singular arc occurs. This solution was validated using computational techniques required for later problems in this thesis. Next, a similar course was ascended which included a lateral, curvature based acceleration constraint. As the complexity of the OCP was increased, numerical methods were required to compute

a solution. Finally, a vertically undulating course was driven. As the route now required braking, the effects of regenerative versus dissipative braking is investigated in terms of the resulting velocity and force profiles. All of these examples have relevance in the construction of the optimal driving problem of the next two chapters.

One of the key issues facing continued internal combustion engine use is the requirement for effective engine calibration and testing. Both procedures rely on fixed drive cycles which reflect poorly real world driving. Chapter 4 investigated a means to generate test cycle data by completing a fuel-optimal drive of a representative driving route. The vehicle model incorporated statistical models for the prediction of emissions. By splitting the route into phases of differing legislative speed limit, the route in question was split into twelve separate segments. An OCP was posed within each, with a different state bound. State continuity was enforced between mesh intervals and total integrals were summed. The sensitivity of the resulting solution to limits on the total route emissions and gearbox type were discussed. The effects of random stops on the solution were also considered. It was shown that the methods presented offer an effective way to compute simultaneous optimal driving and powertrain control trajectories. This is one of the main results of this thesis. The resulting speed traces are compared to real driving data logged on the route in Appendix C.

A novel extension to the optimal driving problem was considered in Chapter 5. This chapter investigated the possibility for emissions reduction by altering the heating strategy during the warmup phase for a electric-hybrid vehicle. The powertrain in question was capable of running the engine at a power independent of that at the wheels. A novel powertrain model for a parallel series hybrid vehicle is utilised in order to avoid imposing either a fixed ICE switching law or require discrete decision variables. An engine friction model was incorporated with a constraint enforcing power production above the idle speed. The optimisation chose the duration and number of engine switches during the drive. A regularisation term was appended to the cost function to control the engine switching; the results of the optimisation show that this requires careful selection. An unrealistic frequency

of switching can arise in the case that this weighting is too low; too high and the engine cannot be switched at all. The results of emissions constraints were examined. It was found that the differing species require different strategies for abatement. The case of traversing the route in the opposite direction was also examined; different results occurred here as well. This example driving scenario also found that large reductions in emissions, up to 40% of unconstrained quantities, can be achieved for a minor increase in fuel consumption. Not only does this further push the usefulness for hybrid powertrain technology for fuel consumption reduction, great emissions reductions can be achieved.

Finally, Chapter 6 considered the application of dynamic optimal control to lean ICE combustion mode switching. A model was created based on a standard data set and structure for the turbocharged, direct injection engine. Several smooth approximations are completed for cylinder processes and for the turbocharger.

The engine in this chapter was designed to operate in a lean combustion regime, so this required the periodic regeneration of a lean-NO_x trap equipped in the tailpipe. This poses problems for a typical dual-loop engine controller. Dynamic optimal control was used to investigate the means in which unwanted transient emissions can be reduced. It was found that combined control of the airpath actuators to pre-empt the switch in combustion mode was beneficial. The slowest system response, in this case the turbocharger inertia, has to be adjusted before the switch-over occurs. This result was found for both a simple step case, and for a longer drive cycle segment.

7.1 Future work

The results presented in this thesis are by no means an exhaustive consideration of all of these problems. There are a number of key extensions that can be explored based on the methods utilised.

Future areas of interest are listed below.

1. **Drive Cycle Set Generation**

   A set of drive cycles could be generated based on the traversal of the route in

Chapter 4. When creating engine calibrations it is useful to optimise controller maps and gains based on a single drive cycle. The validation of the controller can be completed for the same route, but a different arrival time. A set of drive cycles can be made for a single route with differing arrival times and number of stops. Further, additional validation would be required using real driving data to ensure the acceptability of the resulting calibration.

2. **Powertrain Control Maps**

   In the formulation presented in Chapter 4, the gearbox shifting law is kept simple. It was optimised simultaneously with the engine torque and vehicle speed. In reality, the gearbox shifting law is more complex than (4.13). Not only does the gear selected depend on the forward speed of the vehicle, but also the accelerator pedal position (an example heuristic map is used in [19]). A smooth representation of this shifting law could, be solved for simultaneously along with the optimal driving trajectories. The accelerator pedal would also require mapping to engine torque demand. Again, this depends on the vehicle speed. Finally, a brake pedal control signal will be required. A constraint would also need to be imposed which ensured both pedals were not actuated simultaneously.

   This formulation provides a more realistic view of actual driver control inputs and optimises the input of both pedals, together with the control maps. The only drawback to this is the requirement for the control maps to be easily parametrised for the sake of tractability. It could then be investigated whether or not total emissions constraints affect the resulting control laws which would be implemented in the engine control unit.

3. **Further case studies for power-split optimisation**

   Chapter 5 considered the optimisation of powersplit along a single route (traversed in either direction) under a series of scenarios. Additionally to this, the SOC state was constrained to begin and end at 60% of the maximum. A number of further cases can be considered. A large number of vehicles are

used for commuting - i.e the traversal of the same route in opposite directions on a given day. This opens the possibility to have a daily terminal SOC constraint, but relax the constraint at the end of the initial commute (and the beginning of the return journey). The case where the energy store is filled on the first journey, while the TWC is hot, could enable emissions to be reduced on the return journey by extensive electric running. This will be dependent on the route in question and limitations on the total emissions allowed.

Plug-in hybrid electric vehicles (PHEV) have the capability to be recharged external to the vehicle. In this case, the terminal SOC constraint can be removed. This opens up the flexibility for the vehicle to use differing strategies for fuel-optimal driving. It is likely that the whole reserve of electrical energy will be used along the route as the optimisation will favour avoiding ICE usage. In the case of emission-constrained operation, the optimisation can now choose to discharge the energy store fully or to complete the route with a partially filled energy store.

It was found that the nature of the route had a profound effect on the total emissions produced as well as the means by which they can be abated. It would be prudent to see the effect of lengthening the journey, and the chosen route on the the total emissions. This could be investigated for both the charge-sustaining scenario of Chapter 5 and for the aforementioned PHEV case. The effects of this increased complexity may have a detrimental effect on the computation time.

4. **Increase model complexity for hybrid powertrain**

   A number of modelling simplifications were made in the computation of the results in Chapter 5. One clear area for expansion of the model used is to incorporate temperature dynamics in the engine itself. Steady state temperature of the exiting gases were assumed in this thesis. While there will be some degradation from heating the cold engine block, the main effect will be on the friction characteristic. Additional models for coolant circulation and

the heat generated in the battery could also be included. The effects of rapid heating of the engine on reduced friction losses represents another means for the reduction in fuel consumption. Engine block temperature effects on the creation of emissions is another way in which the model complexity can be increased. An increase in the number of dynamic states is required which will further slow the optimisation routine. It is necessary to appraise the increase in computation time and the increased information that the results have gained.

The powertrain model could be further increased in the same means as for the extensions to the diesel car. Control maps, mapping the accelerator and brake pedals to torques for each machine and the generator speed could be parsed in the optimal control formulation. Again, these would have to be parametrised so that the OCP could still be solved efficiently.

5. **Controller law design**

   The major motivator for completing the chapter on supervisory heating control was to discuss the effects of route knowledge on optimal/sub-optimal heating. Emissions abatement is an increasingly important facet of online control.

   The method for optimising the control effort in Chapter 5 is designed for offline control. However, the results of which can prove useful in creating pre-determined rulesets for simple controller architectures to follow. Using a more simple model for the vehicle, the demanded tractive power effort can be estimated before the trip begins (or during the early stages of which). This can then be used, along with knowledge of the journey, to decide whether to heat the TWC rapidly or not. In order to populate such a parameter space, the optimisation would have to be computed for a large number of scenarios, and the solution of each checked. Differing initial SOC, route length, predicted power demands and terminal SOC could form a parameter space from which parameters for a simple online control law can be found.

The alternative to this would be a reduced order model that will incorporate the important features of the system, and utilising an online, receding horizon control algorithm. This would calculate future control inputs that were optimal at the current time. The horizon of such an algorithm would have to be sufficiently long as to include the entire period over which the heating was required.

6. **Increase complexity of gasoline airpath models.**

   The tractability of the airpath optimisation in Chapter 6 relied on a number of simplifications in the model. Most notably is the steady state data set used for the prediction of emissions out of the engine. A model of greater complexity should be utilised instead of the two dimensional ones used. Higher order models can better capture the effects of intake gas composition and valve timing. Further to this, dynamic compensation can be incorporated based on the rate of change of torque. Another simplification that could be addressed in the future is the constant isentropic efficiencies assumed for the compressor and turbine. The addition of this information into the optimal control problem could cause issue in terms of computational tractability.

7. **Real-time airpath controller**

   Based on the results of Chapter 6, it was seen that pre-emptive control of the airpath actuators enabled better transient response when tracking dual targets of tailpipe λ and the output torque.

   Real-time capable controllers which can actuate all of the available control inputs can be designed based on the results of the optimal control problem. A regeneration command signal could switch over tracking targets to the values required for lean combustion, rather than requiring an instantaneous change-over in the output targets.

   Another viable means to incorporate previewing of the regeneration command is to utilise receding horizon model predictive control. For gasoline airpaths, linear-time varying approximations to the dynamics is a feasible means to

implement such a controller. The optimisation horizon would have to be long enough to allow the controller time to actuate the system far enough in advance of the switch.

Appendices

A

Distance Domain Reformulation

In this thesis, the minimum energy (or fuel) optimal control problem is discussed and solved in both the time and distance domains. It was noted that there is a direct comparison to draw between the two. Chapter 3 saw conclusions drawn about the problem structure in the time domain which were assumed to carry over into the distance domain. The short analysis here provides some comfort in knowing that the same optimal control structure can be obtained if the minimum energy OCP was completely analysed in the distance domain.

A.1 Time domain

The original optimal control problem is repeated here for comparison. The objective is to minimise the total input energy E_t, defined as the following integral

$$E_{ft} = \int_0^{t_f} F(t)v(t)\,dt. \tag{A.1}$$

This cost is minimised under the presence of the dynamic constraints:

$$\dot{v}(t) = \frac{F(t)}{M} - K \tag{A.2}$$

$$\dot{s}(t) = v(t) \tag{A.3}$$

The standard constraints are placed on the speed $v(t)$ and the controlled force $F(t)$.

$$0 \leq v(t) \leq v_m \tag{A.4}$$

$$0 \leq F(t) \leq F_m \tag{A.5}$$

The boundary conditions are required as follows

$$s(0) = 0,\ s(t_f) = s_f \tag{A.6}$$

$$v(0) = 0,\ v(t_f) = \text{free}. \tag{A.7}$$

A.1.1 Distance domain

This minimum energy OCP can now be reformulated into the distance domain. The variable transformation between t- and s- domains is one-to-one, provided that the vehicle speed is strictly positive. This takes the form:

$$\frac{d}{ds} = \frac{1}{v}\frac{d}{dt}. \tag{A.8}$$

The transformed OCP reads: minimise the integral cost function

$$E_{fs} = \int_0^{s_f} F(s)\, ds. \tag{A.9}$$

subject to the dynamic constraints:

$$\frac{dv(s)}{ds} = \frac{F(s)}{v(s)M} - \frac{K}{v(s)} \tag{A.10}$$

$$\frac{dt(s)}{ds} = \frac{1}{v(s)} \tag{A.11}$$

Path constraints on the speed $v(s)$ and the controlled force $F(s)$ are

$$\epsilon \leq v(s) \leq v_m \tag{A.12}$$

$$0 \leq F(s) \leq F_m \tag{A.13}$$

and the required boundary conditions are

$$v(0) = \epsilon,\ v(s_f) = \text{free} \tag{A.14}$$

$$t(0) = 0,\ t(s_f) = t_f. \tag{A.15}$$

In this formulation, it is important to note that the velocity is a divisor in both of the state equations. To ensure tractability, the velocity is lower bounded by a small quantity ϵ. Furthermore, the total number of states has not changed; transforming the OCP into the distance domain does not reduce the problem size. This does not allow for periods of stopped operation (by way of a zero velocity limit). Indeed, as the independent variable is always increasing, this is not feasible within the problem framework.

The new Hamiltonian in this case is

$$\mathcal{H}(s) = F(s) + \lambda_v(s)\left(\frac{F(s) - MK}{Mv(s)}\right) + \frac{\lambda_t(s)}{v(s)} + \mu(s)(v(s) - v_m). \tag{A.16}$$

The costate dynamics are as follows

$$\frac{d\lambda_v(s)}{ds} = \frac{\lambda_v(s)(F(s) - MK) + M\lambda_t(s)}{Mv(s)^2} - \mu(s), \tag{A.17}$$

$$\frac{d\lambda_t(s)}{ds} = 0. \tag{A.18}$$

In this case $\lambda_t(s) = \lambda_t$ is a constant. The control again takes the form:

$$F^*(s) = \begin{cases} 0, & \text{if} \quad \Phi_s(s) > 0 \\ F_s(s), & \text{if} \quad \Phi_s(s) \equiv 0 \\ F_m, & \text{if} \quad \Phi_s(s) < 0 \end{cases} \tag{A.19}$$

The new switching function is

$$\Phi_s = 1 + \frac{\lambda_v(s)}{Mv(s)}. \tag{A.20}$$

Taking the derivative once again with respect to distance will enable the solution structure to be uncovered.

$$\frac{d\Phi_s}{ds} = \frac{1}{Mv(s)^2}\left[v(s)\frac{d\lambda_v(s)}{ds} - \lambda_v(s)\frac{dv(s)}{ds}\right] \tag{A.21}$$

$$= \frac{1}{Mv(s)^2}\left[v(s)\left(\frac{\lambda_v(s)(F(s) - MK) + M\lambda_t(s) - \mu(s)Mv(s)^2}{Mv(s)^2}\right)\right. \\ \left. - \lambda_v(s)\left(\frac{F(s) - MK}{Mv(s)}\right)\right] \tag{A.22}$$

$$= \frac{\lambda_t}{Mv(s)^3} - \frac{\mu(s)}{Mv(s)} \tag{A.23}$$

Compared to the time domain, the derivative of the switching function is not constant along the trajectory in either speed limited or unconstrained cases. Clearly, the only time that this function can be zero for any finite period is when $\mu(s) = \frac{\lambda_t}{v(s)^2}$. Consider now the second derivative of $\Phi_s(s)$:

$$\frac{d^2\Phi_s}{ds^2} = \frac{-3\lambda_t(s)}{Mv(s)^4}\frac{dv(s)}{ds} - \frac{1}{v(s)^2}\left(v(s)\frac{d\mu(s)}{ds} - \mu(s)\frac{dv(s)}{ds}\right). \tag{A.24}$$

Now the only case for a singular sub-arc is

$$\frac{dv(s)}{ds} = \frac{d\mu(s)}{ds} = 0. \tag{A.25}$$

This implies that the period at which $F(s) = MK$, also has constant, non-zero path multiplier $\mu = \lambda_t/v_m^2$. Therefore, the case that a singular arc occurs is only possible in the case that the velocity constraint is met[1]. The control is piecewise defined again as:

$$F^*(s) = \begin{cases} F_m, & s \in [0, \sigma_1] \\ MK, & s \in (\sigma_1, \sigma_2] \\ 0 & s \in (\sigma_2, s_f] \end{cases} \tag{A.26}$$

The switching distances σ_1 and σ_2 are unknown. The corresponding state trajectories can be found by integrating the dynamic equation (A.10)

$$v_1(s) = \sqrt{2(F_m/M - K)s + \epsilon^2} \tag{A.27}$$

$$v_2(s) = v_m \tag{A.28}$$

$$v_3(s) = \sqrt{2K(\sigma_2 - s) + v_m^2} \tag{A.29}$$

similarly the three sub arc state trajectories for t can be written:

$$t_1(s) = \frac{\sqrt{2(F_m/M - K)s}}{F_m/M - K} \tag{A.30}$$

$$t_2(s) = \frac{\sqrt{2(F_m/M - K)\sigma_1}}{F_m/M - K} + \frac{s - \sigma_1}{v_m} \tag{A.31}$$

$$t_3(s) = \frac{\sqrt{2(F_m/M - K)\sigma_1}}{F_m/M - K} + \frac{\sigma_2 - \sigma_1}{v_m} - \frac{\sqrt{K(\sigma_2 - s)}}{K} \tag{A.32}$$

[1]In the case that the speed limit is sufficiently large, the constraint is never active. $\frac{d\Phi_s}{ds}$ cannot remain at zero for any finite period of time. The control takes a bang-bang form as before.

The switching distances σ_1 and σ_2 can be computed by enforcing continuity between state sub-arcs. In the limit as $\epsilon \to 0$, the trajectories become equivalent to those found in Chapter 3. The expressions describing the switching distances are

$$\sigma_1 = \frac{v_m^2 - \epsilon^2}{2(F_m/M - K)} \tag{A.33}$$

and

$$\sigma_2 = s_f - \frac{v_m^2 - v_f^2}{2K}. \tag{A.34}$$

These are directly analogous to τ_1 and τ_2 in Chapter 3.

Again, the argument for fulfilment of an upper bounding arrival time can be made. In the case of the OCP posed in the distance domain, this requires a terminal state constraint. The solution of the simple example of Chapter 3 is made more difficult in the distance domain, and is perhaps less intuitive as the resulting dynamics are non-linear. This analysis further reinforces the idea of transforming the minimum fuel OCP between the time and distance domain. Of course, the proviso is that the velocity is always positive which does affect the solution accuracy. This can be mitigated provided the velocity boundary condition ϵ is small.

B

Estimation of Route Data

A key requirement for the work presented in Chapters 4 and 5 is the estimation of route information. A brief overview of the method used to transform raw GPS data in a form for optimisation is discussed here.

The two routes chosen both incorporate urban driving and highway driving. Route A is chosen to be closed and slightly longer - giving rise to drive cycle-like length journeys. The second route, Route B, is a commuter type route which aims to reflect a typical urban-highway-urban commute. In Chapter 5 this route is traversed in both directions. In one direction, this is a short 2km section of urban driving followed by the highway section of about 8km. The final section of urban driving is about 5km. Total elevation changes for this route are small.

The requirements for the vehicle models (and optimal control problems) of the previous chapters have neglected the effects of road camber and width. The road curvature forms an integral part of the allowable accelerations, and thus vehicle speeds. It is important therefore, that the data is not too heavily smoothed as the optimal control will choose speeds to negotiate corners which are higher than would be acceptable in real life.

It is also necessary to obtain a suitably smooth approximation of the road inclination along the route. This is required for proper appraisal of negotiating

hills. This appendix considers the estimation and smoothing of both of these data sets obtained from route mapping software.

B.1 Road Curvature

It is convenient to describe the route as a long curve with lateral curvature $\mathcal{C}$ and inclination θ. As the vehicle model has no requirement for road camber, these two descriptors are independent [56]. Progress along the route is always increasing, as positive speeds are considered. Therefore a one-to-one mapping between the distance travelled along the curve s, and ψ and $\mathcal{C}$ exists. The route data set is taken from GPS and is given as a latitude, longitude and elevation triple (λ, ϕ, z). In this thesis, this is obtained from the online route mapping software available at *www.mapmyride.com*. This raw data follows the centreline of the road lane at a random sampling distance. The raw GPS data is transformed into a local Cartesian coordinate frame (x, y, z) based on the UTM (WGS84 Geodetic) system using a mean Earth radius of 6317km. The starting point of the route is deemed the origin of the coordinate frame. Once the data is in the this domain, it can be effectively transformed the curve information along the routes length. The geometry of the transformation between the Cartesian space to the curve quantities is shown in Figure B.1. An incremental change in distance ds along the spine, in the direction of the tangent vector $\mathbf{t}$, is related to incremental changes in the Cartesian coordinates x and y via:

$$dx = ds\,\cos\psi, \tag{B.1}$$

$$dy = ds\,\sin\psi. \tag{B.2}$$

The point A in Figure B.1 is defined by the radius of the curve at that instance $1/\mathcal{C}(s)$ and ψ, the orientation angle of the tangent vector $\mathbf{t}$. The curvature is related to the orientation angle by the following derivative

$$\mathcal{C}(s) = \frac{d\psi(s)}{ds}, \tag{B.3}$$

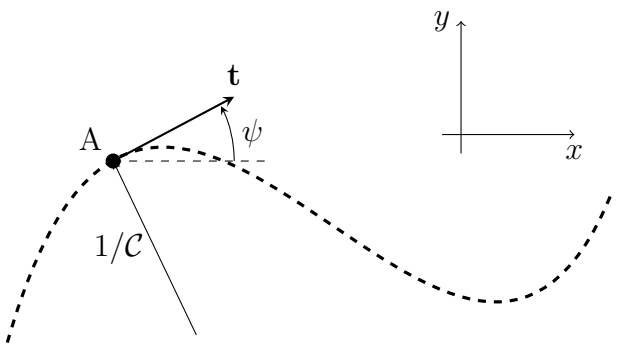


Figure B.1: The geometry of the differential transformation required to estimate the road curvature. The road is drawn as the dashed line. The tangent to the road at point A is $\mathbf{t}$, and it makes an angle ψ with the global x-axis.

and the orientation angle is simply the gradient of the tangent vector in the Cartesian space:

$$\psi(s) = \arctan \frac{dy(s)}{dx}. \tag{B.4}$$

In the absence of noise, numerical differentiation of the data set should reveal the curvature. However, as the route is long, and contains several tight corners, a smoother approximation of the curvature is required. Additionally, GPS data is not guaranteed noise-free - a problem which is amplified by successive numerical differentiation. It is also expected that the data set is not equally spaced along the route. In order to alleviate these issues, an appropriate OCP can be posed which can determine a suitably smooth spine. If the rate of change of curvature is controlled as such

$$u(s) = \frac{d\mathcal{C}(s)}{ds} \tag{B.5}$$

then this expression along with (B.1),(B.2) and (B.3) forms a complete set of dynamics for the problem. The elapsed distance along the GPS data can be estimated, and this forms a quadratic error cost:

$$J = \int_0^{s_f} (u(s)^2 + w_c((x_c(s) - x(s))^2 + (y_c(s) - y(s))^2))\, ds. \tag{B.6}$$

x_c and y_c are the raw GPS data along the centreline of the road, and their elapsed distance is estimated via cumulative summation of the distance between them. Intermediate points are computed via spline interpolation. The optimal control problem is to, therefore, select the rate of change of curvature along the route line to minimise the reconstruction error weighted against smoothing the derivative of the curvature. Estimation noise and rapidly changing curvature can be avoided by tuning the integrating weight w_c. In the case that the route is closed, the following boundary constraints are enforced

$$x(0) = x(s_f) \tag{B.7}$$

$$y(0) = y(s_f) \tag{B.8}$$

$$\mathcal{C}(0) = \mathcal{C}(s_f). \tag{B.9}$$

Else, the initial and final x and y co-ordinates must be the same as that from the GPS data.

B.1.1 Results

The non-linear OCP for curvature estimation is solved using GPOPS-II for Route B (which appears in Chapter 5).

A comparison of the original $x - y$ data with the fitted spine is plotted in Figure B.2. The insets on the right correspond to features of high curvature in the route. Slow moving features, such as in B, are well captured regardless of the weighting. However, the effect of filtering on tight corners such as A or D shows how the weighting selection is important. Over-fitting of the data is feasible in the case when the weighting is too high (the blue trace in A). This rapidly changing curvature is not well suited for the OCP formulations required in this thesis. Again in C, over-fitting of the data is seen, whereas decreasing the weighting provides a reasonable estimate of the curvature and remains accurate. As the spine length is fixed to be the same as that for the raw data, avoiding high curvature can only feasibly be achieved by deviating away from the original data. Poor estimation of the curvature is seen in inset D for the $w_c = 1 \times 10^{-8}$ weighting.

The sign reversal could prove problematic, but this must be traded off against the curvature which too rapidly changes direction.

The curvature estimates in Figure B.3 further highlight the oscillatory nature of an overfitted estimate. The rapidly switching solution for an overly accurate reconstruction is demonstrated at 6km. The curvature estimates for each of the insets in Figure B.2 is also considered in Figure B.3. This data is that which appears in the OCPs in this thesis and as such is required to be smoothed. The poor fitting of tight corners is once again demonstrated in insets A and D. For the OCPs solved in Chapter 5, the weighting of $w_c = 1 \times 10^{-8}$ is used.

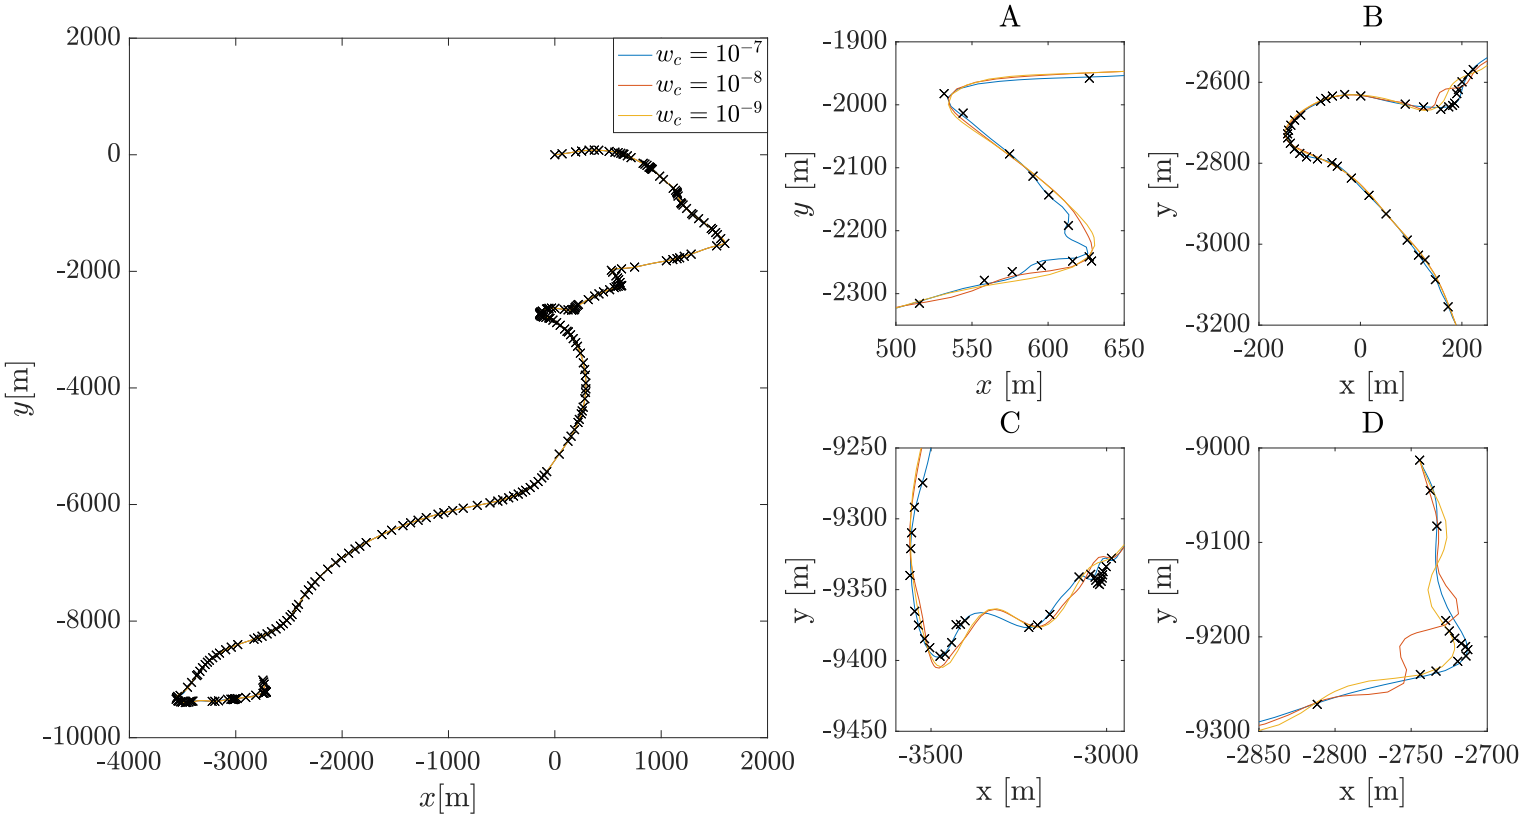


Figure B.2: Route reconstruction with the original GPS data as black crosses. The insets demonstrate over-fitting of the data if w_c is chosen too large.

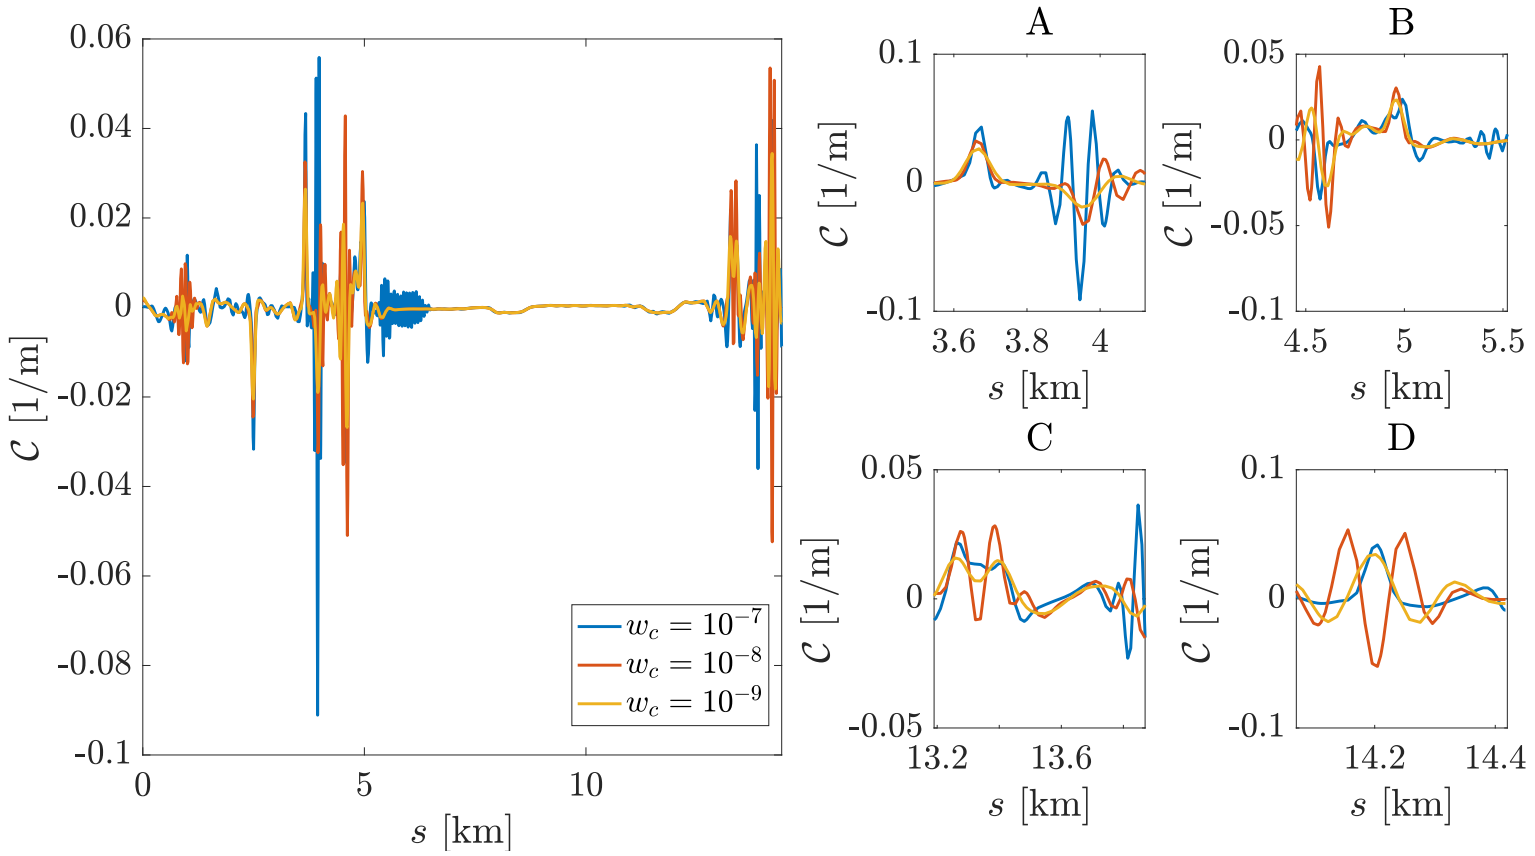


Figure B.3: Curvature estimates plotted as a function of distance along the spine.

B.2 Road Grade

Similarly, the inclination angle $\theta(s)$[1] can be estimated by considering an optimal filtering problem.

This ensures that the data is sufficiently smooth for the purposes of optimal control. The elevation data is calculated based on the height of the road surface above sea-level. The dynamics of the filtering problem can be posed as such:

$$\begin{aligned} \frac{dz}{ds}(s) &= \theta(s) \\ \frac{d\theta}{ds}(s) &= u(s). \end{aligned} \tag{B.10}$$

Continuity can again be enforced at the beginning and end of the loop in the case that the route is closed:

$$\begin{aligned} z(0) &= z(s_f) \\ \theta(0) &= \theta(s_f). \end{aligned} \tag{B.11}$$

A performance index which incorporates both the magnitude of the error between the GPS data and a regularisation term can be formed

$$J = \int_0^{s_f} (z(s) - z_{GPS}(s))^2 + w_e u(s)^2 ds. \tag{B.12}$$

The weighting factor w_e allows the exact reconstruction of the elevation to be traded off against the smoothing of the estimated inclination. The OCP is solved using the same method as for the curvature reconstruction. Considering the data points in Figure B.4, the exact reconstruction of the elevation data is difficult. The GPS data is concentrated in areas where the curvature changes rapidly. This is unlikely to be in the same place where the elevation changes are large; cresting a hill is not likely to be during a sharp turn.

When the weighting w_e is large, hilly features are not captured properly. This is clear in the insets of Figure B.4.

The variation of w_e on the road inclination estimate is examined in Figure B.5. The sensitivity of the inclination estimate to changes in the reconstruction weighting is shown.

[1]It should be noted that $\theta(s)$ should really be defined as $\theta(s) = \tan^{-1}(\frac{\mathrm{d}z}{\mathrm{d}s})$. A small angle approximation can be used to simplify the OCP, i.e $\theta(s) \approx \frac{\mathrm{d}z}{\mathrm{d}s}$.

The results of both optimisations for Route A (found in Chapter 4) are plotted in Figure B.6 for completeness.

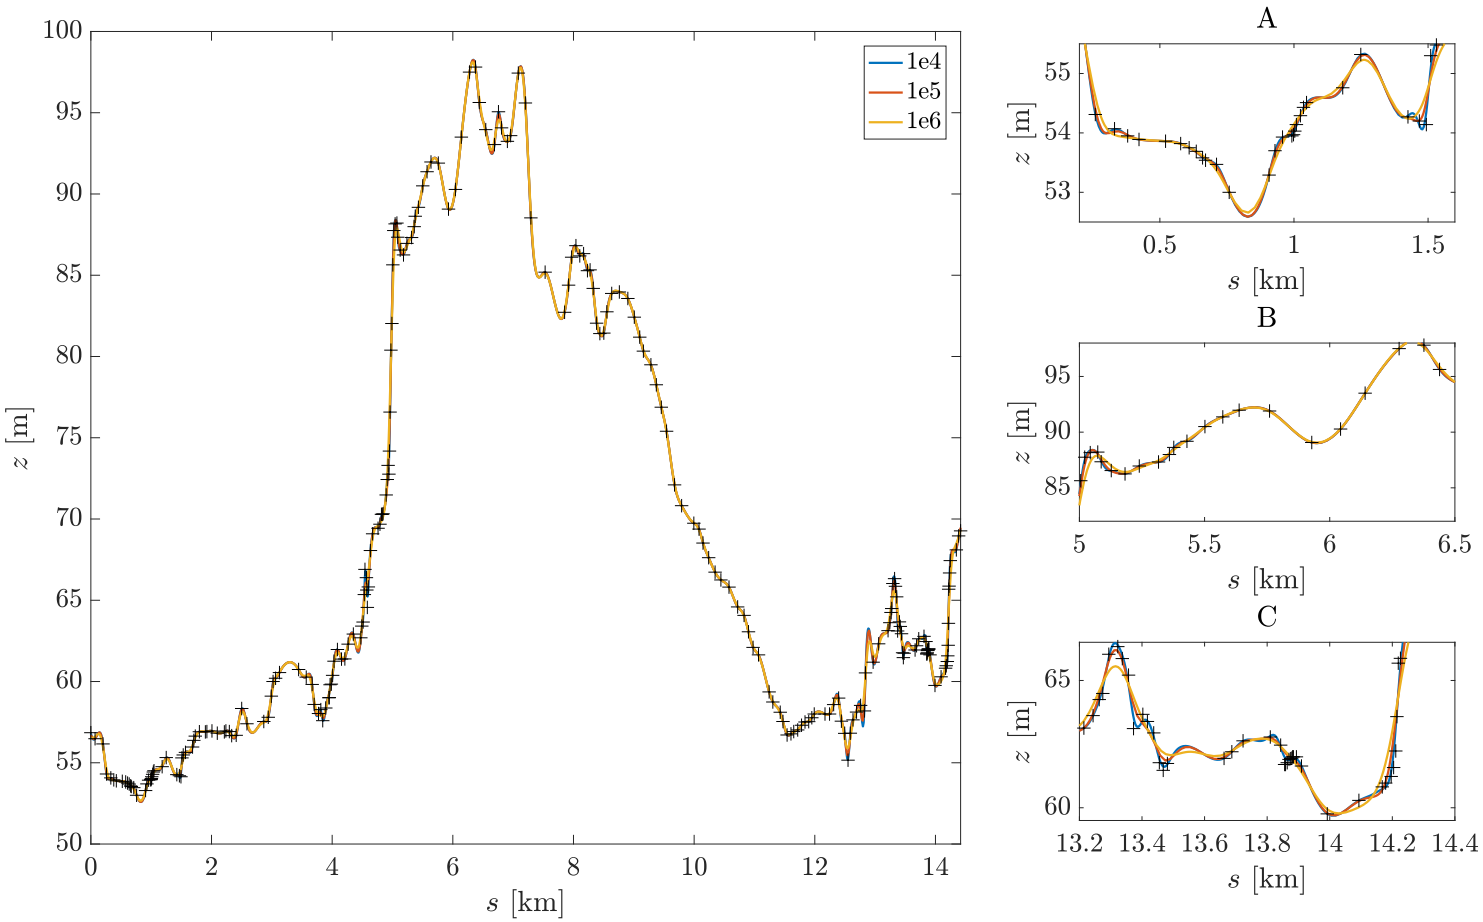


Figure B.4: Route B elevation reconstruction based on the solution to the optimal filtering problem. Highlights on the right-hand side show the effects of over or under-reconstruction.

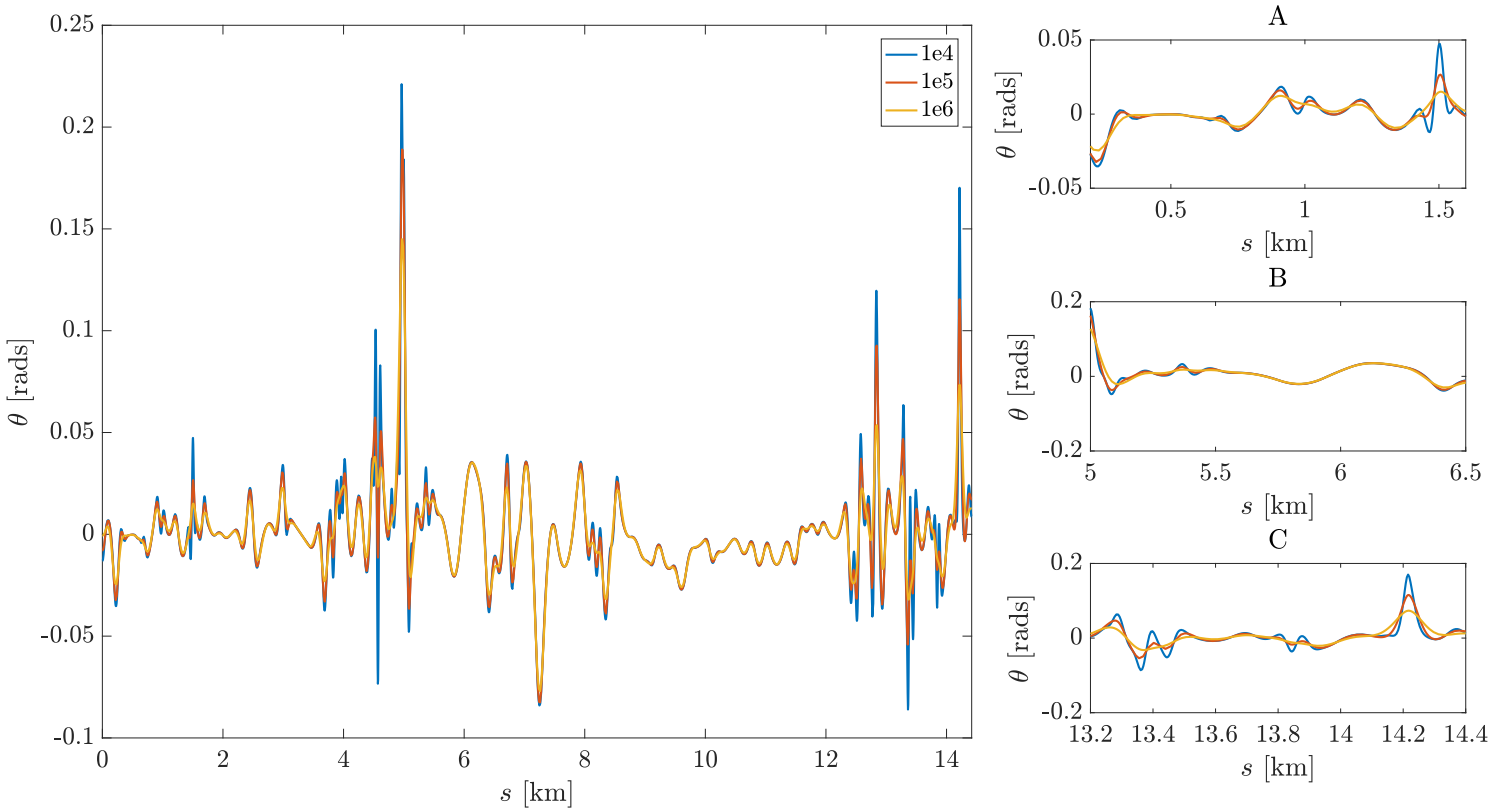


Figure B.5: Comparison of the rapidly changing inclination angles for the elevation reconstruction. Insets are at the same points as Figure B.4.

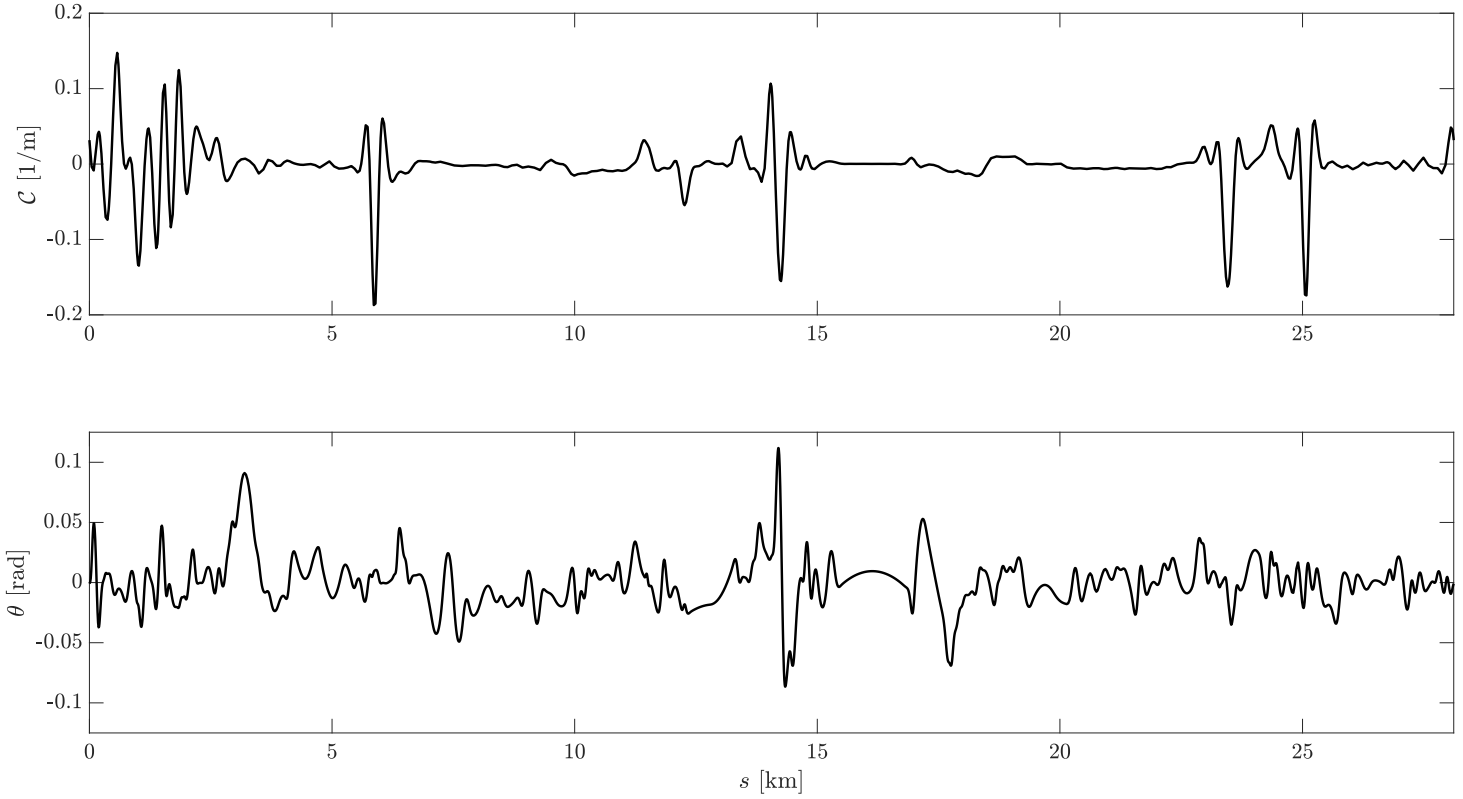


Figure B.6: Route A estimated curvature and inclination.

C
Velocity Profile Validation

In order to validate the velocity profiles generated in Chapter 4, Route A is driven with a real car. This appendix serves to check that the results of the optimisation are reasonable and to also improve the information that the optimisation could use moving into the future.

The route is driven with a Chevrolet Matiz 1.0l petrol engined car, during a time when the roads are quietest to correlate to the low/no traffic assumption. The vehicle used for the computation of driving profiles in Chapter 4 was of a different type; however this validation effort is concerned with the speed trace generated rather than correlation the power profile as well.

Vehicle speed data is obtained over the Controller Area Network (CAN) bus through the on-board diagnostics (OBD) port (Figure C.1b). Another reasonable means to log speed data relies on the use of GPS. At a set frequency, the position of the vehicle is logged; this then requires numerical differentiation to obtain speed data. This is likely to be too noisy compared to the data obtained over the CAN.

Internal measurements of vehicle speed are obtained via an encoder mounted on the driveline. Inside the ECU, this is converted, using knowledge of the tyre radius, to the car's speed. Also, as this speed appears on the dashboard, it is likely that this quantity is larger than the actual speed the vehicle is travelling. Vehicle instrumentation panels are designed to err on the side of caution when



(a) The vehicle used for validation. **(b)** A picture of the OBD port

Figure C.1: Equipment used for the validation in this appendix.

displaying the speed. The tyre pressure and tread wear will effect the radius of the tyre and add further error to the measurement.

In order to quantify the error in the velocity measurements, the route is first traversed with both GPS and OBD logging enabled. Elapsed distance is calculated using a trapezoidal approximation of the integral of the velocity data. The total can be compared to that found with GPS data and the ratio of the two forms a correction factor. This scaling can be applied to all of the logged ECU data to ensure the speed error is compensated for in the remaining. The trace of the route data logged for both ECU and GPS (unscaled) is plotted in Figure C.2. The solution of the run with stops at each instance from Chapter 4 is plotted for comparison with speed data from both GPS and OBD. Without filtering the GPS data, the speed trace is particularly noisy. Despite the OBD logging at a similar sample rate, much greater clarity of the trace is seen.

A series of six validation runs are plotted in Figure C.3 and are compared to the original speed limit set and one of the speed traces from Chapter 4.

The original speed limit information in the first 2.5km, up to the point (A), is too high. Additional stops are apparent; a couple appear at 6.3km and 8.3km. A further stop which is met during every drive is found at (C). The vehicle was required to drive more quickly that the optimised speed profile at the point (B). This is due to the limitations posed by other vehicles on sections of the drive which require higher speeds. Finally at (D), there is another error in the speed limit

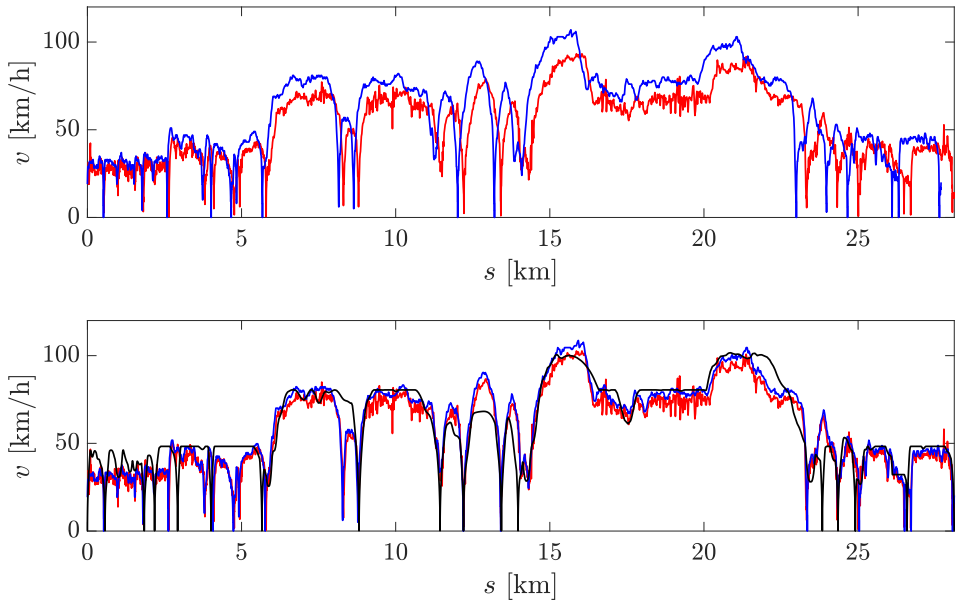


Figure C.2: The original GPS (red) and OBD (blue) data is plotted in the top graph. Bottom: Scaled GPS (red) and OBD (blue) data, is compared to one simulation run for all stops (black).

information. The change from highway to urban driving occurs much later along the route. Part of the reason for the low speeds obtained by the optimiser at (B) is due to the incorrect speed limit at (A). It is also due in part to the selected arrival time constraint enforcding a fixed average velocity which is close enough to the actual logged journey times. Durations for the journey are logged in Table C.1 along with the total time that the vehicle is in motion. Despite the original problem containing some small errors in speed limit information, the obtained driving data is reasonably well correlated to the results of the optimisation. The selected arrival time of 2000 seconds for the all-stopping optimisation case is a little too long. Furthermore, the total time spent at rest is approximately 120 seconds which is 6% of the average total journey time. If the vehicle cannot stop the engine during stationary operation, then the fuel consumption prediction using this offline method will under-estimate the total fuel usage.

This confirms that the results obtained in Chapter 4 are reasonable. Despite the fact that the driver will drive sub-optimally, there is a strong correlation between the logged data and the optimisation results. Greater correlation can be achieved

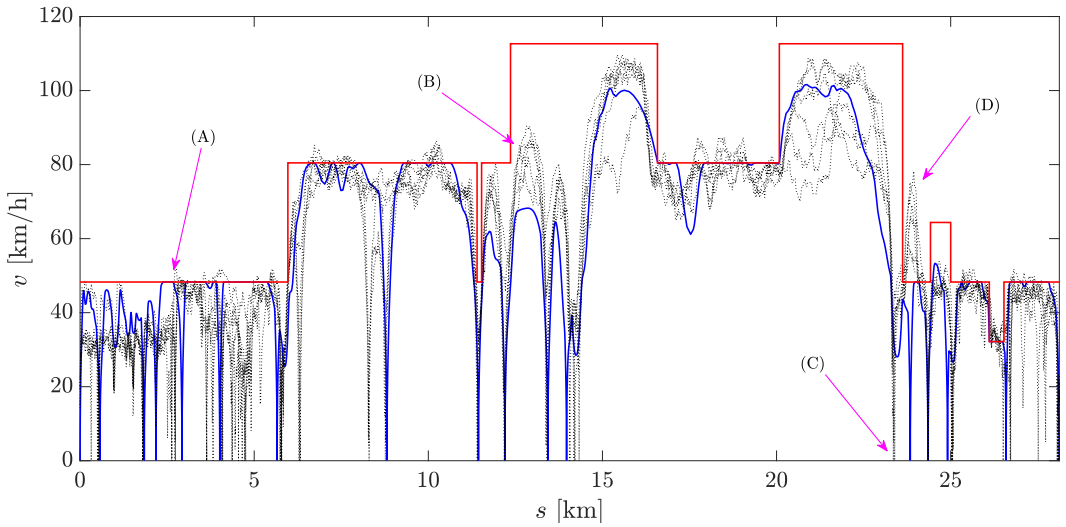


Figure C.3: Adjusted logged data for Route A (dotted). Original speed limit information is red and an optimisation required to stop at all stops ($T = 2000$s) is plotted in solid blue.

Table C.1: Times to complete each validation run with the duration moving as well.

Run	Total Time (s)	Time in Transit (s)
1	2195	2043
2	1978	1866
3	2121	2027
4	2010	1944
5	2125	1974
6	2123	1982
Mean	2092	1973

by re-running the optimisation with this new information regarding the frequency and placement of stops and speed limits. It is evident that, as in other engine calibration exercises, it is an iterative process that proves the best method to obtain the desired results. Other signals that can be used in other validation efforts (which vary from vehicle to vehicle).

... if I have seen further, it is by standing on the shoulders of giants

— Sir Isaac Newton letter to Robert Hooke, 1675

References

[1] *Does air pollution kill 40,000 people each year in the UK?* 2017. URL: https://wintoncentre.maths.cam.ac.uk/news/does-air-pollution-kill-40000-people-each-year-uk.

[2] *Oxford Zero Emission Zone update briefing note*. 2019. URL: https://www.oxford.gov.uk/downloads/file/5377/oxford_zero_emission_zone_update_briefing_note_-_january_2019.

[3] Felix Leach et al. "Road bumps for electric cars". In: *The Economist* 413.9053 (Sept. 2017).

[4] *CO_2 Emissions from Cars: the facts*. Apr. 2018. URL: https://www.transportenvironment.org/sites/te/files/publications/2018_04_CO2_emissions_cars_The_facts_report_final_0_0.pdf.

[5] *Driving into 2025: The Future of Electric Vehicles*. Oct. 10, 2018. URL: https://www.jpmorgan.com/global/research/electric-vehicles.

[6] Erich Brandt, Yanying Wang, and Jessy Grizzle. "Dynamic modeling of a three-way catalyst for SI engine exhaust emission control". In: *IEEE Trans. Control Syst. Technol.* 8.5 (2000), pp. 767–776.

[7] Chinmay Patil, Sanjyot Varade, and Swapnil Wadkar. "A review of engine downsizing and its effects". In: *International Journal of Current Engineering and Technology* Special Issue-7 (2017), pp. 319–324.

[8] *The New Four-Cylinder Gasoline Engines from Mercedes-Benz*. 2013. URL: http://schmidt-cnc.de/media/files/MTZ-Beitrag-Die-neuen-Ottomotoren-en-EN.pdf.

[9] Lei Liu, Zhijun Li, and Boxi Shen. *The Application and Optimization of EGR-LNT Synergetic Control System on Lean-burn Gasoline Engine*. Tech. rep. SAE Technical Paper, 2015.

[10] Rolf Isermann. *Engine Modeling and Control: Modeling and Electronic Management of Internal Combustion Engines*. Berlin, Germany: Springer-Verlag, 2014.

[11] Gokul S Sankar et al. *Fast Calibration of a Robust Model Predictive Controller for Diesel Engine Airpath*. 2018. URL: arXiv:1804.06161v1.

[12] Alberto Bemporad et al. "Model Predictive Control of Turbocharged Gasoline Engines for Mass Production". In: *WCX World Congress Experience*. SAE International, Apr. 2018.

[13] Qinghua Lin, Pingen Chen, and Vitaly Y. Prikhodko. "Model Predictive Control of a Lean-Burn Gasoline Engine Coupled With a Passive Selective Catalytic Reduction System". In: *ASME 2017 Dynamic Systems and Control Conference*. 1. ASME, Oct. 2017, pp. 1–9.

[14] Jonas Asprion, Oscar Chinellato, and Lino Guzzella. "Efficient solution of the diesel-engine optimal control problem by time-domain decomposition". In: *Control Engineering Practice* 30 (2014), pp. 34–44.

[15] Giorgio Mancini et al. "Dynamic Feedforward Control of a Diesel Engine Based on Optimal Transient Compensation Maps". In: *Energies* 7 (2014), pp. 5400–5424.

[16] Tomaž Katrašnik. "Analytical framework for analyzing the energy conversion efficiency of different hybrid electric vehicle topologies". In: *Energy Conversion and Management* 50.8 (2009), pp. 1924–1938.

[17] Jinming Liu and Huei Peng. "Modeling and control of a power-split hybrid vehicle". In: *IEEE Transactions on Control Systems Technology* 16.6 (2008), pp. 1242–1251.

[18] Antonio Sciarretta, Michael Back, and Lino Guzzella. "Optimal control of parallel hybrid electric vehicles". In: *IEEE Transactions on Control Systems Technology* 12.3 (2004), pp. 352–363.

[19] Tobias Nüesch et al. "Equivalent consumption minimization strategy for the control of real driving NOx emissions of a diesel hybrid electric vehicle". In: *Energies* 7.5 (2014), pp. 3148–3178.

[20] Hoseinali Borhan et al. "MPC-Based Energy Management of a Power-Split Hybrid Electric Vehicle". In: *IEEE Transactions on Control Systems Technology* 20.3 (2011), pp. 593–603.

[21] Sebastian East and Mark Cannon. "Fast Optimal Energy Management With Engine On/Off Decisions for Plug-in Hybrid Electric Vehicles". In: *IEEE Control Systems Letters* 3.4 (2019), pp. 1074–1079.

[22] Delphine Sinoquet, Gregory Rousseau, and Yohan Milhau. "Design optimization and optimal control for hybrid vehicles". In: *Optimization and Engineering* 12.1 (Mar. 2011), pp. 199–213.

[23] Stephan Zentner et al. "An equivalent emission minimization strategy for causal optimal control of diesel engines". In: *Energies* 7.3 (2014), pp. 1230–1250.

[24] *Don't Breathe Here: Tackling air pollution from vehicles*. Oct. 14, 2015. URL: `https://www.transportenvironment.org/publications/dont-breathe-here-tackling-air-pollution-vehicles`.

[25] Delphi. *Delphi Worldwide Emissions Standards Passenger Cars and Light Duty Vehicles 2017-2018*. 2017. URL: `https://www.delphi.com/sites/default/files/inline-files/delphi-worldwide-emissions-standards-passenger-cars-light-duty-2016-7.pdf`.

[26] Emilia Silvas et al. "Synthesis of realistic driving cycles with high accuracy and computational speed, including slope information". In: *IEEE Transactions on Vehicular Technology* 65.6 (2016), pp. 4118–4128.

[27] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge, United Kingdom: Cambridge University Press, 2010.

[28] William H Press et al. *Numerical recipes 3rd edition: The art of scientific computing*. Cambridge, United Kingdom: Cambridge University Press, 2007.

[29] Arthur Bryson and Yu-Chi Ho. *Applied Optimal Control: Optimization, Estimation and Control*. London, United Kingdom: John Wiley & Sons, 1975.

[30] Daniel Liberzon. *Calculus of Variations and Optimal Control Theory: A Concise Introduction*. Princeton, United States of America: Princeton University Press, 2012.

[31] Donald E Kirk. *Optimal control theory: An Introduction*. New York, United States of America: Dover Publications, 2004.

[32] John T Betts. *Practical methods for optimal control and estimation using nonlinear programming, Second Edition*. Philadelphia, United States of America: SIAM, 2010.

[33] Jorge Nocedal and Stephen Wright. *Numerical optimization*. New York, United States of America: Springer-Verlag, 2006.

[34] Philip E Gill, Walter Murray, and Michael A Saunders. "SNOPT: An SQP algorithm for large-scale constrained optimization". In: *SIAM review* 47.1 (2005), pp. 99–131.

[35] Andreas Wächter and Lorenz Biegler. "On the Implementation of a Primal-Dual Interior Point Filter Line a Search Algorithm for Large-Scale Nonlinear Programming". In: *Mathematical Programming* 106.1 (2006), pp. 25–57.

[36] David JN Limebeer and Matteo Massaro. *Dynamics and optimal control of road vehicles*. Oxford, United Kingdom: Oxford University Press, 2018.

[37] Lev Semenovich Pontryagin. *Mathematical theory of optimal processes*. London: Routledge, 1987.

[38] Francesco Biral, Enrico Bertolazzi, and Paolo Bosetti. "Notes on numerical methods for solving optimal control problems". In: *IEEJ Journal of Industry Applications* 5.2 (2016), pp. 154–166.

[39] Anil V. Rao et al. "GPOPS- II: A MATLAB Software for Solving Multiple-Phase Optimal Control Problems Using hp-Adaptive Gaussian Quadrature Collocation Methods and Sparse Nonlinear Programming". In: *ACM Transactions on Mathematical Software* 37.2 (2010), pp. 1–39.

[40] Divya Garg et al. "An Overview of Three Pseudospectral Methods for the Numerical Solution of Optimal Control Problems". In: *Advances in the Astronautical Sciences* (2009), pp. 1–17.

[41] Anil V. Rao. "A survey of numerical methods for optimal control". In: *Advances in the Astronautical Sciences* 135.1 (2009), pp. 497–528.

[42] Per E Rutquist and Marcus M Edvall. "Propt-matlab optimal control software". In: *Tomlab Optimization Inc* 260.1 (2010), p. 12.

[43] Oskar von Stryk. *User's Guide for DIRCOL*. 1999.

[44] Roberto Lot et al. "Optimal control of road vehicles: Direct and indirect approaches". In: *15th European Automotive Congress (05/10/17)*. Oct. 2017.

[45] William W Hager. "Runge-Kutta methods in optimal control and the transformed adjoint system". In: *Numerische Mathematik* 87.2 (2000), pp. 247–282.

[46] Sina Ober-Blöbaum, Oliver Junge, and Jerrold E. Marsden. “Discrete mechanics and optimal control: An analysis*”. In: *ESAIM: COCV* 17.2 (2011), pp. 322–352.

[47] J Frédéric Bonnans and Julien Laurent-Varin. “Computation of order conditions for symplectic partitioned Runge-Kutta schemes with application to optimal control”. In: *Numerische Mathematik* 103.1 (2006), pp. 1–10.

[48] James N Lyness and Cleve B Moler. “Numerical differentiation of analytic functions”. In: *SIAM Journal on Numerical Analysis* 4.2 (1967), pp. 202–210.

[49] Andreas Griewank and Andrea Walther. *Evaluating derivatives: principles and techniques of algorithmic differentiation. Second Edition.* Philadelphia, United States of America: SIAM, 2008.

[50] Matthew J. Weinstein and Anil V. Rao. “A source transformation via operator overloading method for the automatic differentiation of mathematical functions in MATLAB”. In: *ACM Transactions on Mathematical Software* 44.2 (2017), p. 25.

[51] Fengjin Liu, William W. Hager, and Anil V. Rao. “An hp Mesh Refinement Method for Optimal Control Using Discontinuity Detection and Mesh Size Reduction”. In: *53rd IEEE Conference on Decision and Control, Los Angeles* (2014), pp. 5868–5873.

[52] Arthur J. Krener. “The High Order Maximal Principle and Its Application to Singular Extremals”. In: *SIAM Journal on Control and Optimization* 15.2 (1977), pp. 256–293.

[53] David J N Limebeer and Anil V. Rao. “Faster, Higher, and Greener: Vehicular Optimal Control”. In: *IEEE Control Systems* 35.2 (2015), pp. 36–56.

[54] David H. Jacobson, Stanley B. Gershwin, and Milind M. Lele. “Computation of Optimal Singular Controls”. In: *IEEE Transactions on Automatic Control* 15.1 (1970), pp. 67–73.

[55] Adam Lowell Schwartz. “Theory and implementation of numerical methods based on Runge-Kutta integration for solving optimal control problems”. PhD thesis. University of California at Berkeley, 1996.

[56] Giacomo Perantoni and David J.N. Limebeer. “Optimal control for a Formula One car with variable parameters”. In: *Vehicle System Dynamics* 52.5 (2014), pp. 653–678.

[57] Giacomo Perantoni. “Optimal Control of Vehicle Systems”. PhD thesis. University of Oxford, 2013.

[58] Antonio Sciarretta, Giovanni De Nunzio, and Luis Leon Ojeda. “Optimal Ecodriving Control: Energy-Efficient Driving of Road Vehicles as an Optimal Control Problem”. In: *IEEE Control Systems* 35.5 (2015), pp. 71–90.

[59] Gareth Pease, David J N Limebeer, and Peter Fussey. “Fuel Minimisation for Diesel Power Cars with Emissions Constraints”. In: *IEEE Transactions on Control Systems Technology* 19.3 (2019), pp. 545–555.

[60] Eric W Weisstein. *Lemniscate*. MathWorld - A Wolfram Web Resource. URL: `http://mathworld.wolfram.com/Lemniscate.html`.

[61] Thijs Van Keulen, Bram De Jager, and Maarten Steinbuch. “Optimal trajectories for vehicles with energy recovery options”. In: *IFAC Proceedings Volumes*. Vol. 44. 1. Milan, Italy: IFAC, 2011, pp. 3831–3836.

[62] Gareth Pease, David Limebeer, and Peter Fussey. “Diesel Engine Optimisation, with Emissions Constraints, on a Prescribed Driving Route”. In: *2017 IEEE AFRICON*. Cape Town, South Africa, 2017, pp. 410–415.

[63] Philipp Polterauer and Harald Waschl. “An Implementable Multi-Calibration Strategy to Address Real Driving Emissions of CI Engines”. In: *2017 American Control Conference (ACC)*. July 3, 2017, pp. 297–302.

[64] Hippolyte Bouvier, Guillaume Colin, and Yann Chamaillard. “Determination and comparison of optimal eco-driving cycles for hybrid electric vehicles”. In: *2015 European Control Conference, ECC 2015*. 1. 2015, pp. 142–147.

[65] Hong Tu Luu, Lydie Nouvelière, and Said Mammar. “Dynamic programming for fuel consumption optimization on light vehicle”. In: *IFAC Proceedings Volumes (IFAC-PapersOnline)* (2010), pp. 372–377.

[66] Roberto Lot and Simos A. Evangelou. “Green driving optimization of a series hybrid electric vehicle”. In: *Proceedings of the IEEE Conference on Decision and Control* (2013), pp. 2200–2207.

[67] Bart Saerens. “Optimal Control Based Eco-Driving”. PhD thesis. Katholieke Universiteit Leuven, 2012.

[68] Wei Zhou et al. “A Pseudospectral Strategy for Optimal Power Management in Series Hybrid Electric Vehicles”. In: *IEEE Transactions on Vehicular Technology* 65.6 (June 2016), pp. 4813–4825.

[69] Shouyang Wei et al. “A pseudospectral method for solving optimal control problem of a hybrid tracked vehicle”. In: *Applied Energy* 194 (2017), pp. 588–595.

[70] Mehdi Imani Masouleh and David J N Limebeer. “Fuel Minimisation for a Vehicle Equipped with a Flywheel and Battery on a Three-Dimensional Track”. In: *IEEE Transactions on Intelligent Vehicles* 2.3 (2017), pp. 161–174.

[71] Benjamin Berger, Florian Rauscher, and Boris Lohmann. “Analysing Gaussian processes for stationary black-box combustion engine modelling”. In: *IFAC Proceedings Volumes (IFAC-PapersOnline)* 18.PART 1 (2011), pp. 10633–10640.

[72] Benjamin Berger. “Modeling and Optimization for Stationary Base Engine Calibration”. PhD thesis. Technische Universitat München, 2012.

[73] Carl E Rasmussen and Christopher K I Williams. *Gaussian Processes for Machine Learning*. 2006.

[74] Daniel P. Kelly. “Lap Time Simulation with Transient Vehicle and Tyre Dynamics”. PhD thesis. Cranfield University, 2008.

[75] F. Biral, M. Da Lio, and E. Bertolazzi. “Combining safety margins and user preferences into a driving criterion for optimal control-based computation of reference maneuvers for an ADAS of the next generation”. In: *IEEE Intelligent Vehicles Symposium, Proceedings*. Vol. 2005. 2005, pp. 36–41.

[76] Richard Stone. *Introduction to Internal Combustion Engines*. 4th Edition. United Kingdom: Palgrave Macmillan, 2012.

[77] Jiahui Tang et al. “Energy Management of a Parallel Hybrid Electric Vehicle with CVT Using Model Predictive Control”. In: *Proceedings of the 35th Chinese Control Conference*. Chengdu, China, 2016, pp. 4396–4401.

[78] Dongsuk Kum, Huei Peng, and Norman K. Bucknor. “Optimal energy and catalyst temperature management of plug-in hybrid electric vehicles for minimum fuel consumption and tail-pipe emissions”. In: *IEEE Transactions on Control Systems Technology* 21.1 (2013), pp. 14–26.

[79] Markus Schori, Thomas Boehme, and Matthias Schultalbers. “Optimal Catalytic Converter Heating in Hybrid Vehicles”. In: *SAE Technical Paper*. Apr. 1, 2014.

[80] Pierre Michel et al. “Optimizing fuel consumption and pollutant emissions of gasoline-HEV with catalytic converter”. In: *Control Engineering Practice* 61 (2017), pp. 198–205.

[81] Vital van Reeven et al. “Optimal Control of Engine Warmup in Hybrid Vehicles”. In: *Oil & Gas Science and Technology - Revue d’IFP Energies nouvelles* 71.1 (2016), p. 14.

[82] Xuefang Li, Simos A. Evangelou, and Roberto Lot. “Integrated Management of Powertrain and Engine Cooling System for Parallel Hybrid Electric Vehicles”. In: *2018 IEEE Vehicle Power and Propulsion Conference (VPPC)*. Aug. 2018.

[83] Hamilton H. Mabie and Charles F. Reinholtz. *Mechanisms and Dynamics of Machinery*. London, United Kingdom: John Wiley & Sons, 1987.

[84] Tony Markel et al. “ADVISOR: A systems analysis tool for advanced vehicle modeling”. In: *Journal of Power Sources* 110.2 (2002), pp. 255–266.

[85] M.C.F. Donkers et al. “Optimal control for integrated emission management in diesel engines”. In: *Control Engineering Practice* 61 (2017), pp. 206–216.

[86] Ke Fang et al. “Optimal-Behavior-Based Dynamic Calibration of the Automotive Diesel Engine”. In: *IEEE Transactions on Control Systems Technology* 24.3 (2016), pp. 979–991.

[87] Mrdjan Jankovic, Miroslava Jankovic, and Ilya Kolmanovsky. “Constructive Lyapunov control design for turbocharged diesel engines”. In: *IEEE Transactions on Control Systems Technology* 8.2 (2000), pp. 288–299.

[88] Mertens Jung and Keith Glover. “Control-oriented linear parameter-varying modelling of a turbocharged diesel engine”. In: *Proceedings of 2003 IEEE Conference on Control Applications*. Vol. 1. 2003, pp. 155–160.

[89] Felipe Castillo Buenaventura et al. “Air Fraction and EGR Proportion Control for Dual Loop EGR Diesel Engines”. In: *Ingenieria y Universidad* 19.1 (2015), p. 115.

[90] Ping Wang, Chaojie Zhu, and Jinwu Gao. “Feedforward Model Predictive Control of Fuel-Air Ratio for Lean-Burn Spark-Ignition Gasoline Engines of Passenger Cars”. In: *IEEE Access* (2019), pp. 73961–73969.

[91] Timothy Broomhead et al. “Model Reduction of Diesel Mean Value Engine Models”. In: *SAE Technical Paper*. Vol. 2015-01-12. 2015.

[92] Daniel Alberer and Luigi del Re. “Optimization of the transient diesel engine operation”. In: *SAE Technical Paper Series*. 4970. SAE International, Sept. 2009.

[93] Alex Darlington, Keith Glover, and Nick Collings. “A Simple Diesel Engine Air-Path Model to Predict the Cylinder Charge During Transients : Strategies for Reducing Transient Emissions Spikes”. In: *SAE Technical Paper Series*. 724. SAE International, Oct. 2006.

[94] Jonas Asprion, Oscar Chinellato, and Lino Guzzella. "Optimal Control of Diesel Engines: Numerical Methods, Applications, and Experimental Validation". In: *Mathematical Problems in Engineering* 2014 (2013), pp. 1–21.

[95] Frank Willems et al. "Integrated emission management for cost optimal EGR-SCR balancing in diesels". In: *IFAC Proceedings Volumes* 46.PART 1 (2013), pp. 711–716.

[96] L. Guzzella and A. Amstutz. "Control of diesel engines". In: *IEEE Control Systems Magazine* 18.5 (1998), pp. 53–71.

[97] Dariusz Cieslar. "Control for transient response of turbocharged engines". PhD thesis. University of Cambridge, 2013.

[98] Jun-Mo Kang, Ilya Kolmanovsky, and JW Grizzle. "Dynamic optimization of lean burn engine aftertreatment". In: *Journal of Dynamic Systems, Measurement, and Control* 123.2 (2001), pp. 153–160.

[99] Mathworks. *Simulink Powertrain Blockset*. 2018. URL: `https://uk.mathworks.com/products/powertrain.html`.

[100] Oskar Leufvén and Lars Eriksson. "Surge and choke capable compressor model". In: *IFAC Proceedings Volumes*. Vol. 18. Elsevier BV, Jan. 2011, pp. 10653–10658.

[101] Oskar Leufvén and Lars Eriksson. "A surge and choke capable compressor flow model-Validation and extrapolation capability". In: *Control Engineering Practice* 21.12 (2013), pp. 1871–1883.

[102] Jonas Asprion. "Optimal control of diesel engines modelling, numerical methods, and applications". PhD thesis. ETH Zürich, 2013.

[103] Viktor Leek, Kristoffer Ekberg, and Lars Eriksson. "Development and Usage of a Continuously Differentiable Heavy Duty Diesel Engine Model Equipped with VGT and EGR". In: *SAE Technical Papers* 2017 (Apr. 2017).

[104] John B Heywood. *Internal Combustion Engine Fundementals*. McGraw-Hill, 1988.

[105] J. Gerhardt, H. Hönninger, and H. Bischof. "A New Approach to Functional and Software Structure for Engine Management Systems - BOSCH ME7". In: *SAE Technical Paper*. SAE International, Feb. 1998.