



UNIVERSITEIT VAN PRETORIA
UNIVERSITY OF PRETORIA
YUNIBESITHI YA PRETORIA

Make today matter

Optimising vehicle velocity profiles in real-time

By

Riaan Frederick Meeser

A thesis in partial fulfilment for the degree

Philosophiae Doctor (Mechanical Engineering)

in the

Department of Mechanical and Aeronautical Engineering

Faculty of Engineering, the Built Environment and Information Technology

University of Pretoria

South Africa

October 2023

Abstract

Reducing natural resource consumption is essential in moving towards a sustainable society. One such application field is the mining sector, where electric locomotives haul tons of ore for many kilometres underground. Improving the efficiency of these vehicles will result in higher machine utilisation, increased profitability and reduced overall cost of ownership of these vehicles. To optimise the energy consumption of such a vehicle, it is essential to know its energy requirements.

The main contributing factors to energy consumption are the type of vehicle, the route being travelled, the gross combined mass and the velocity profile to complete the route in the logistical constraints imposed. Regenerative braking, various powertrain efficiencies and environmental conditions also significantly affect the overall energy consumed on a route.

The broad goal of this thesis is the development of a method that is capable of providing a vehicle controller with an optimal velocity profile to travel a route by, considering the route, loads and logistics in real-time. This project has three phases: route identification, mass estimation and velocity profile optimisation. The individual aspects are handled as generically as possible to broaden the fields of application.

A novel method of route identification is developed to identify an underground route being travelled in real-time from a database without the use of external communication. The route is identified by searching for patterns in heading and altitude data by comparing these to the routes available in a database, which is a computationally efficient method for real-time applications. If the route cannot be identified, the data is stored as a new route in the database and used for route identification in future traversing. The route topography from the recorded data is now available and used throughout the rest of the study. Extensive above-ground and limited underground tests confirm the usability of the strategy.

A patented real-time mass estimation strategy is developed that uses a simple torsional load cell on the driving axle of the vehicle. This method makes mass estimation simpler as it doesn't require multiple load cells and doesn't need trailers to be instrumented either. The driving torque, route topography and vehicle dynamics equations are applied to Newton's second law of motion for the vehicle, with the equations rearranged to solve for the unknown mass of the vehicle rig. Test results show that the vehicle's gross combined mass can be estimated to be within 5% of the actual value.

The only controllable variable in energy consumption for a specific vehicle on a specific route is the vehicle's velocity profile. A strategy is developed that can robustly optimise the velocity profile for the vehicle on a known route with a known mass such that the overall energy consumed is reduced, while still adhering to the logistical time requirements. The results show an energy reduction from 8.5 MJ to 5.7 MJ (33%) is obtained over a real-world velocity profile for a 10 km test route. The key to the

optimiser's robust operation lies in how the initial guess to the local solver, *fmincon*, is executed using strategic inverting, scaling and shifting of the route's topographic profile in a low fidelity, quickly executable optimisation.

The initial guess strategy proved successful in yielding estimates of the optimal velocity profile close to the complete optimisation solution's values, typically within 10%, but at a fraction of the computational time, typically being around 20 seconds as compared to 3 hours for the full optimisation. The low-fidelity model facilitates an easily implemented lookup table to retrieve the scale and shift parameter values, thus estimating the optimal velocity profile within a fraction of a millisecond, rather than hours. This fast method facilitates real-time optimisation of a vehicle's velocity profile if the route is known.

This thesis set out to develop a system capable of reducing the energy consumption of a vehicle and proved successful in performing this optimisation in real-time due to the advancements made in real-time route identification, real-time mass estimation and real-time optimisation of the velocity profile for the vehicle.

Table of contents

Abstract	i
Table of contents.....	iii
List of Figures	vi
List of Tables	viii
List of Equations.....	ix
List of Symbols/Abbreviations	ix
1. Introduction	1
1.1. Motivation for the study	1
1.2. Problem statement.....	1
1.3. Research Objectives and Questions.....	2
1.3.1. Specific objectives of this work.....	2
1.3.2. Research questions that this work aims to answer	2
1.4. Scope and limitations.....	2
1.5. Summary	5
2. Underground Route Identification	6
2.1. Introduction.....	6
2.2. Materials and Methods.....	7
2.2.1. Hybrid Vehicle Systems	7
2.2.2. Route identification and currently applied methods.....	8
2.2.3. Topography/incline data acquisition.....	8
2.2.4. Heading data acquisition.....	11
2.2.5. Odometer	11
2.2.6. GPS as reference for odometer, velocity and barometric altitude.....	12
2.3. Calculations: modelling of the route identification system	12
2.3.1. Data Acquisition System.....	12
2.3.2. Above-ground tests.....	14
2.3.3. Underground tests	21
2.4. Results and Discussion	22
2.4.1. Above-ground route identification	22
2.4.2. Basement test results.....	24
2.4.3. Gold Reef City Underground Test Results	25

2.5. Conclusion	27
2.6. Future work	27
3. Real-time gross mass estimation.....	28
3.1. Introduction.....	28
3.1.1. Background.....	28
3.1.2. Motivation for this project	28
3.1.3. Project requirements and beneficiaries	29
3.2. Literature study	29
3.3. Theory/calculation (theoretical modelling).....	32
3.3.1. Determining forces for mass estimation	32
3.3.2. Sum of forces	35
3.3.3. Sensors implemented and vehicle used to facilitate mass estimation	36
3.3.3. Octave program structure schematic.....	41
3.3.4. Experimental setup for mass estimation.....	41
3.3.5. Testing Procedure	42
3.4. Results and Discussion	43
3.5. Conclusion	47
3.6. Future work	48
4. Velocity profile optimisation.....	49
4.1. Introduction.....	49
4.2. Relevant literature	50
4.2.1. Vehicle type	50
4.2.2. Vehicle total mass	51
4.2.3. Route	51
4.2.4. Velocity.....	52
4.2.5. Optimisation Techniques.....	53
4.3. Model	57
4.3.1 The Test Vehicle	58
4.3.2 Route Energy Model.....	58
4.3.3 Sensitivity studies on route energy.....	61
4.3.4. The Optimisation Strategy.....	64
4.4. Simulations & Results	66

4.4.1. Road grade sensitivity analysis	67
4.4.2. Constant initial guess of distance/time	71
4.4.3. Random initial guesses for varying fidelity	75
4.4.4. Fine sweep of values to find success rate	78
4.4.5. Inverted topographic as an initial guess.....	79
4.4.6. The effect of regeneration.....	84
4.4.7. The effect of mass changes	87
4.5. Conclusions.....	88
4.6. Future work	90
5. Real-time route velocity profile optimisation	92
5.1. Background.....	92
5.2. Method.....	92
5.3. Real-time optimal velocity profiles.....	95
5.4. The effect of incorrect route identification.....	98
5.4. Conclusion	100
5.5. Recommendations.....	100
6. Conclusion of work	101
6.1. Route Identification	101
6.2. Mass estimation	101
6.3. Velocity Profile Optimisation.....	102
6.4. Real-time velocity optimisation.....	103
6.5. Overall project	103
7. Recommendations and future work.....	105
7.1. Route Identification	105
7.2. Mass estimation	105
7.3. Optimising of velocity profiles.....	105
7.4. Real-time optimisation of the velocity profile.....	105
7.5. Factors not directly relevant to the work performed.....	105
8. Acknowledgements	107
9. Bibliography	108

List of Figures

Chapter 1

Figure 1.1: Overall project scope	3
---	---

Chapter 2

Figure 2.1: Barometric altitude sensor data	11
Figure 2.2: Arduino data acquisition system operational schematic	13
Figure 2.3: Topographic plot. a) Barometric & GPS vs. distance. b) 3 Passes of the same route	14
Figure 2.4: Route data for 3 passes of the same route. a) Altitude gain, b) Heading	14
Figure 2.5: Matlab program operation schematic	16
Figure 2.6: Probability matrices processing steps. Left: Same route. Right: different route	19
Figure 2.7: Above-ground routes travelled.....	21
Figure 2.8: Surveyor's wheel for underground testing.....	22
Figure 2.9: Correlation vs. Distance when travelling Route 1.....	23
Figure 2.10: Basement parking level 1 test. a) Heading-, b) barometric altitude vs. distance	25
Figure 2.11: Underground gold mine test correlation	26
Figure 2.12: Underground RIS results. a) Basement tests, b) Gold mine tests.....	27

Chapter 3

Figure 3.1: Free body diagram of a vehicle on an incline	32
Figure 3.2: Test vehicle on weighing bridge, Mass = 2090kg, full fuel tank, no payload, no driver. Accuracy within 20kg (1%).....	36
Figure 3.3: Estimated barometric altitude by riding the elevator from floor 9 to 15 to 3 to 9	38
Figure 3.4: Live axle with 8 strain gauges mounted (numbering is consistent with figure 3.5b)	39
Figure 3.5: Strain gauges mounted to the rear axle. a) front right mounting point, b) Wheatstone bridge wiring.....	39
Figure 3.6: Axle torque whilst accelerating in gears 1 through 4	39
Figure 3.7: Vehicle mass estimation sensing circuit diagram	40
Figure 3.8: Octave program schematic for mass estimation	41
Figure 3.9: Test vehicle with a small trailer, GCM 2755kg (75x5kg+230kg+2150kg [driver and full tank]). $A_f \approx 2.33m^2$	42
Figure 3.10: Test vehicle with a large trailer, GCM 3270kg. $A_f \approx 3.3m^2$	42
Figure 3.11: Raw mass estimation, sampled at 5Hz. a) The total mass plot, b) zoomed in	43
Figure 3.12: Test data for calculated mass versus: a) Power, b) Velocity, c) Acceleration	44
Figure 3.13: Estimated Mass for denominator value	44
Figure 3.14: Mass estimated with parameter limits imposed. Blue = calculated real-time mass cut data, Red = LPF.....	45

Chapter 4

Figure 4.1: Quality of solution in optimisation	50
Figure 4.2: Factors influencing the route energy	53
Figure 4.3: Optimisation problem classification [63]	56
Figure 4.4: Route energy calculation in Simulink	59
Figure 4.5: Real-world test routes. a) Topographic profiles, b) Original velocity profiles.....	60
Figure 4.6: Route 4.1: Energy consumed vs. distance	61
Figure 4.7: Route energy vs. Mass for real-world test route 4.1	62
Figure 4.8: Route energy sensitivity to velocity. a) Route energy for the entire route with varying velocity, b) final route energy for varying velocity.....	63
Figure 4.9: Final route energy vs. completion time for real-world test route 4.1	63
Figure 4.10: Velocity profile optimisation framework	67
Figure 4.11: Standard topography cases for evaluating optimal velocity results	68
Figure 4.12: Optimal velocity profiles for straight profile routes with a fidelity of 10.....	68
Figure 4.13: Optimal velocity profiles for 5 % Up-Down & 5 % Down-Up with a fidelity of 10.....	69
Figure 4.14: Optimal velocity profiles for 7 % Up-Down-Up-Down & 7 % Down-Up-Down-Up with a fidelity of 10	69
Figure 4.15: Optimal velocity profile for 7 % up-down-up-down and 7 % down-up-down-up with fidelity of 50.....	70
Figure 4.16: Constant initial guess for test routes, fidelity of 50. a) Optimal velocity profiles, b) Energy. 72	72
Figure 4.17: Optimal velocity profile for constant initial guess, 7% routes, fidelity of 50	72
Figure 4.18: Optimal velocity vector for constant initial guess	74
Figure 4.19: Simulation time vs. model fidelity for distance/time initial guess.....	75
Figure 4.20: Variance in optimisation result for random starts.....	76
Figure 4.21: Optimal velocity vector for best of random initial guesses for route 4.1	77
Figure 4.22: Simulation time vs. model fidelity for a random initial guess for route 4.1.....	77
Figure 4.23: Optimal route 4.1 energy for swept initial values for X1 and X2. a) Swept from 1-30 m/s in increments of 1 m/s, b) Swept close to the optimal in increments of 0.025 m/s	78
Figure 4.24: Route 4.1 simulations for varying noise factor and fidelity.....	81
Figure 4.25: a) Optimal solutions for route 4.1 energy vs. <i>Noise_factor</i> , b) Standard deviation vs. <i>Noise_factor</i>	82
Figure 4.26: Inverted topographic vs. complete optimal solution. a) Velocity profiles, b) Energy.....	83
Figure 4.27: Optimal velocity profiles for the test routes with 0% or 40% regeneration efficiency.....	84
Figure 4.28: Optimal velocity profiles for various regeneration efficiencies, model fidelity of 50, for Route 4.1.....	85
Figure 4.29: Varying mass on route 4.1. a) Optimal velocity profiles, b) Route energy.....	88
Figure 4.30: Summary of steps to find the optimal velocity profile of a vehicle on a known route.....	89

Chapter 5

Figure 5.1: Scale (a) and shift (b) parameter surface plot vs. regeneration efficiency & maximum completion time for route 4.1	93
Figure 5.2: Smoothed scale a) and shift b) parameters for the zero regeneration line	93
Figure 5.3: Smoothed a)scale and b)shift parameter surface plots for a sweep of regeneration efficiency and maximum time for route 4.1	94
Figure 5.4: Real-time program schematic.....	94
Figure 5.5: Optimal route velocity profile for route 4.1 using the optimal for that route as input. a) original m and c, b) smoothed m and c parameters	95
Figure 5.6: Real-time updating of the optimal velocity profile of route 4.1.....	96
Figure 5.7: Real-time and final optimal velocity profile. a) Route 4.2, b) Route 4.3 and c) Route 4.4.....	96
Figure 5.8: Route energy profiles.....	97
Figure 5.9: Total energy wasted by not following the optimal velocity profile	98
Figure 5.10: Optimal velocity profiles for the 4 real-world routes. a) complete, b) zoomed in initial	99
Figure 5. 11: Route energy with incorrect route identification. R4.3 travelled with the velocity optimum of R4.4.....	99

Chapter 6

Figure 6. 1: Schematic of real-time velocity profile optimising system functioning.....	104
---	-----

List of Tables

Chapter 2

Table 2.1: Test route lengths	21
Table 2.2: Summary of route identification results using $\sigma_{alt\ gain} = 3m/100m$, $\sigma_{heading} = 8^\circ$	24

Chapter 3

Table 3.1: Mass estimation parameter calibration values	46
Table 3.2: Mass estimation results.....	47

Chapter 4

Table 4.1: Typical rolling resistance and friction coefficients.....	52
Table 4.2: Final route energy using various initial conditions for a model fidelity of 50	72
Table 4.3: Optimal route final energy for a constant initial guess.....	74
Table 4.4: Random starting points route 4.1 energy data.....	76
Table 4.5: Mean and standard deviation for optimal inverted topographic initial.....	82
Table 4.6: Best route 4.1 energy results of complete optimisation.....	82
Table 4.7: The effect of regeneration on route 4.1 energy consumption	87

Chapter 5

Table 5.1: Energy saving comparison	97
---	----

List of Equations**Chapter 2**

(Eq. 2.1) International barometric formula	9
--	---

Chapter 3

(Eq.3.1) Newton's second law	33
(Eq.3.2) Aerodynamic force	33
(Eq.3.3) Rolling resistance force.....	33
(Eq.3.4) Improved rolling resistance force.....	34
(Eq.3.5) Incline force.....	34
(Eq.3.6) Propulsion force	35
(Eq.3.7) Sum of forces 1.....	35
(Eq.3.8) Sum of forces 2.....	35
(Eq.3.9) Mass equation.....	35
(Eq.3.10) International barometric formula	37

Chapter 4

(Eq.4.1) Aerodynamic force	51
(Eq.4.2) Rolling resistance.....	51
(Eq.4.3) Initial guess with random noise	79

Chapter 5

(Eq.5.1) Route scoring for velocity vector	100
--	-----

List of Symbols/Abbreviations

Latin symbol	Description	Unit (if applicable)
a	Acceleration	m/s^2
A_f	Frontal area	m^2
ADC	Analogue to digital converter	
c	Shift parameter for inverted topographic	
C_d	Drag coefficient	

CMEM	Comprehensive modal emissions model	
CP	Convex programming	
C_{rr}	Coefficient of rolling resistance	
DP	Dynamic programming	
EM	Electric machines	
EMS	Energy management system	
ESS	Energy storage system	
EV	Electric vehicle	
F	Force	N
<i>fmincon</i>	Constrained local solver in Matlab	
<i>fminsearch</i>	Unconstrained local solver in Matlab	
g	Gravitational acceleration	9.81 m/s ²
g	Grams	g
GA	Genetic algorithm	
GCM	Gross combined mass	Kg
GPS	Global positioning system	
GVM	Gross vehicle mass	kg
h	Hour	3600 s
HEV	Hybrid electric vehicle	
hPa	Hectopascal	1 hPa = 100 Pa
Hz	Frequency	Cycles/second
IC	Internal combustion	
ID-EMS	Internet distributed-energy management system	
IR	Infrared	
J	Joule	Nm/s
km	kilometre	1000 m
KE	Kinetic energy	J
kW	kilowatt	1000 W
kts	Knots	1.852 km/h
L	Litre	1e-3 m ³
LIDAR	Light detection and ranging	
LPF	Low-pass filter	
m	metre	m
m	Scale parameter for inverted topographic	
M	Mass	kg
MPC	Model predictive control	
N	Normal force	Newton
n	Number of routes	
Nm	Newton metre	Newton metre
P	Measured pressure	Pa
P	Power	W
P_0	Pressure at sea level	Pa
Pa	Pascal	N/m ²
P1 - P5	Parking level 1 through 5	
PE	Potential energy	J
RIS	Route identification strategy	

SOC	State of charge	%
s	Seconds	seconds
T_{axle}	Torque on the propulsion axle of the vehicle	Nm
V	Velocity	m/s
v	Volt	volt
V2G	Vehicle to grid	
VSP	Vehicle specific power	Kw/ton
$\vec{V}_{initial}$	Initial guess vector	
\vec{V}_{opt}	Optimal inverted topographic vector	
USD	United States Dollar	\$

Greek symbol	Description	Unit (if applicable)
σ	Standard deviation	
μ	Mean	
ρ	Air density	Kg/m ³
θ	Angle	Degrees (°)

1. Introduction

1.1. Motivation for the study

The mining sector in South Africa consumes almost one fifth of all the electricity produced in South Africa annually and contributes significantly to the gross domestic profit of the country [1]. In a typical mine, the electricity costs account for 24% of the mine's total expenditure [2]. The country of South Africa's electricity supplier, ESCOM, is currently under extreme pressure to produce enough electricity for all of its industries. Any potential improvement in electricity consumption will be invaluable to not only the company operating the equipment but also to the whole country.

One field of application where improvements can be made is on underground electric locomotives, used extensively in the mining industry. A typical mine has in excess of sixty of these locomotives in constant operation. These vehicles transport people and tons of ore for many kilometres underground daily. Improving the efficiency of these vehicles will result in increased machine utilisation, profitability and reduced overall cost of ownership of these vehicles as well as a reduced impact on the environment. The vehicles have been produced in various powertrain options, initially running on internal combustion (IC) engines, more recently applications use a hybrid of IC and electric systems, with the most recent being fully electric vehicles utilising large onboard battery energy storage systems (plug-in electric vehicle - PEV). Due to the constant expanding of the mines it is impractical to make use of conventional overhead trolley assisted electrical power systems like those used in normal railway applications [1]. This forces the use of onboard energy supply and storage systems. There are various application fields to improving the resource consumption of vehicles. It can be attempted to improve the mechanical and electrical hardware systems onboard to increase efficiency. The energy converters can be operated in more efficient ways [3], [4]. The loads and operational shifts can be optimised [1]. The driving style/velocity profile can be optimised [5].

1.2. Problem statement

Optimising a vehicle's energy consumption is only possible if the past and future energy requirements of the vehicle's drivetrain are known. The energy requirements depend on many factors, such as the vehicle type, the surface on which the vehicle is travelling, the route that is being travelled along with its corresponding topographic profile and overall length, the gross combined mass of the vehicle as well as mass changes during the completion of a route, logistical constraints, the velocity profile and legislative restrictions on velocity. Regenerative braking, various powertrain efficiencies and environmental conditions also contribute significantly to the energy consumed while traversing a route.

The goal of this project is to develop a robust method of finding the optimal velocity profile that a vehicle should travel a route at to reduce the energy consumption of the vehicle whilst still adhering to logistical constraints. The method is required to operate in real-time, as this allows for graceful accommodation of unforeseen changes to the completion time, facilitating real-world application.

1.3. Research Objectives and Questions

The main goal of the research study is to reduce the overall energy consumption of a vehicle. As mentioned in the problem statement, the energy consumption is a function of several main parameters, such as the route being travelled, the vehicle type, the payload and logistical constraints. The vehicle type will be considered fixed for this study, which leads to the specific objectives of this work, as stated below. Each of these are studied in detail, first in isolation and then brought together later on in the work.

1.3.1. Specific objectives of this work

- Identify the main role-players in the consumption of energy of a vehicle
- Look into the effect of route characteristics on energy consumption
- Look into the effect of mass on energy consumption
- Look into velocity profiles and their effect on route energy

For this study, the focus will be on parameters external to the vehicle itself, which allows vehicle-specific characteristics to be easily updated as required.

1.3.2. Research questions that this work aims to answer

- What are the key role-players in the consumption of energy for a vehicle?
- Can a route be identified easily and in real-time without requiring external infrastructure?
- Can the total mass of the entire vehicle rig be estimated in real-time?
- Do velocity profiles exist that reduce energy consumption but still meet the logistical constraints?
- Can optimal velocity profiles, or estimates to them, be found in real-time?
- What are the trade-offs encountered when performing optimisation in real-time?

1.4. Scope and limitations

In the realm of PhD research, contributions can be broadly categorised into those that offer depth and those that extend breadth in a field of study. Contributions of depth delve deeply into a specific area, aiming to provide profound, detailed insights or solve intricate problems within a narrow scope. This type of contribution is often characterised by intensive analysis, a focus on developing or refining theories, and a commitment to advancing specialised knowledge. It's akin to digging a well, going deeper into a confined space to unearth hidden truths or solutions. On the other hand, contributions of breadth seek to expand the horizons of a field, covering a wider range of topics or interdisciplinary connections. These contributions are more about spanning across different areas, akin to building bridges between various aspects of a discipline or between multiple disciplines. Such work is valuable for opening new avenues of inquiry and fostering a broader understanding and application of concepts. Both types of contributions are crucial: while depth allows for expert mastery and groundbreaking discoveries in a specific area, breadth fosters innovation, interdisciplinary dialogue, and a comprehensive field view. The contribution of this work is that of the latter, where novel products and strategies are developed and incorporated into an optimisation framework to deliver a strategy that allows simple technologies to be used to reduce energy wastage.

The work performed throughout the completion of this study is summarised in Figure 1.1. In Phase 1, a real-time route identification strategy is presented. In Phase 2, a real-time mass estimation strategy is presented. In Phase 3, the outputs of phases 1 and 2 are used to determine an optimal velocity profile for a vehicle to travel that results in the lowest overall energy consumed, given some logistical constraint.

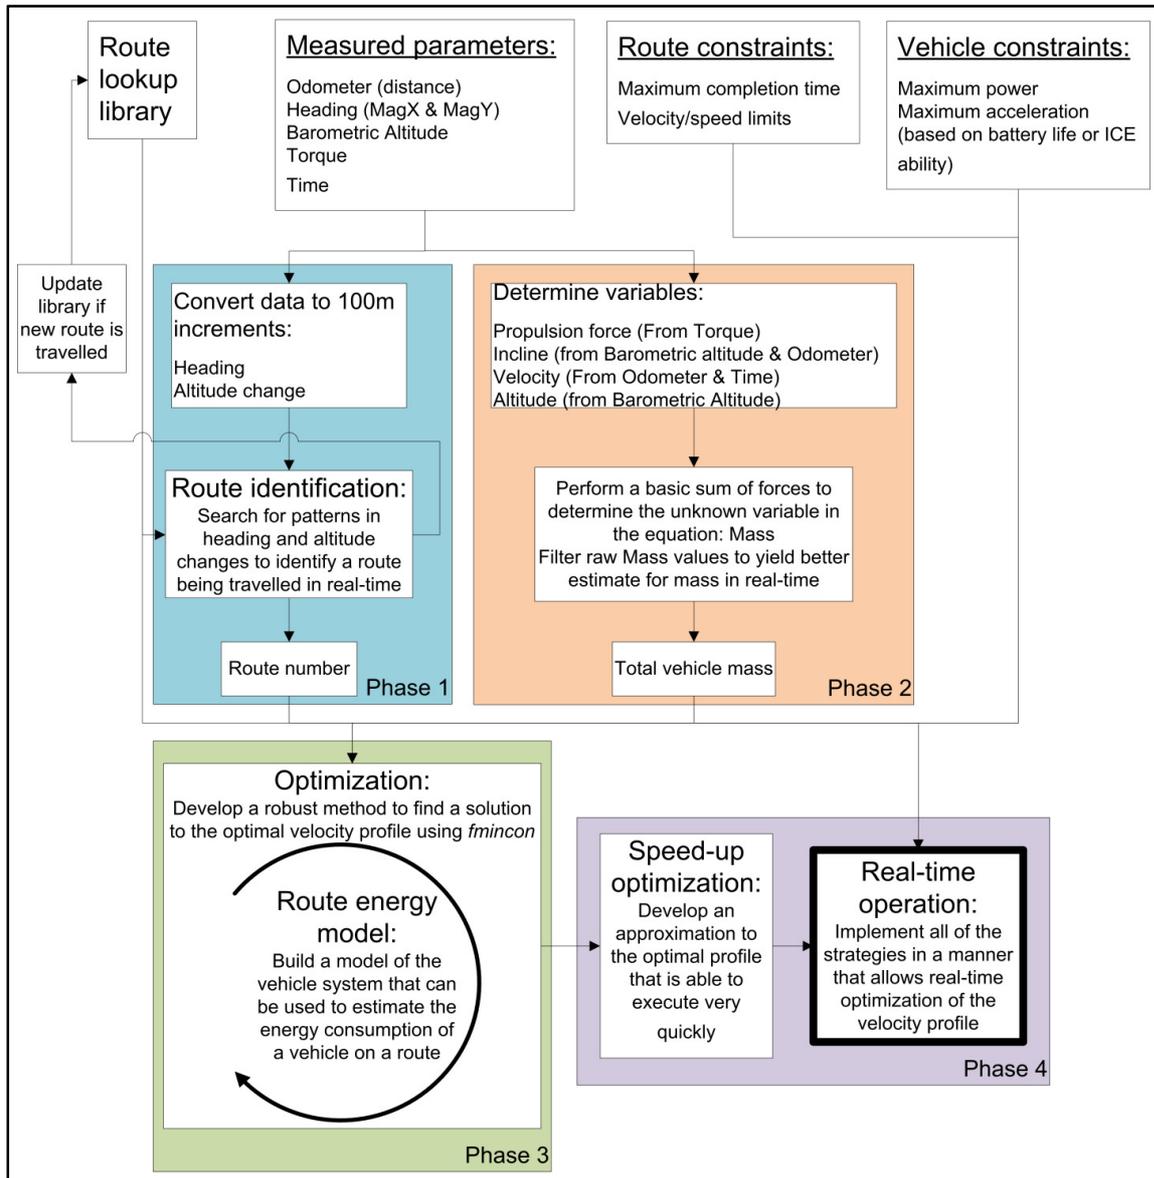


Figure 1.1: Overall project scope

The work is performed in three phases, each of which is a self-contained advancement towards developing a method of finding an optimal velocity profile for a vehicle over its route. It is based on a published paper (Phase 1), a filed patent (Phase 2) and two papers due for submission in relevant journals (Phase 2 and 3). Each chapter is written to be self-contained, allowing the reader to read each

chapter in isolation. Consequently, there is some duplication in the literature review and motivations. However, these are framed within the context of each phase.

Chapter 2 presents a published novel method of route identification that can identify an underground route being travelled in real-time using a database of known routes without external communication [6]. This was the first phase of the project, labelled as Phase 1 in Figure 1.1. The route is identified by searching for patterns in heading and altitude data, recorded as a function of distance and evaluated in real-time by simple, low-cost sensors and comparing these values to an available database in a computationally simple method. If the route cannot be identified, the data is stored as a new route in the database and is used in future traversing of the routes. An accurate route topography dataset is available from the recorded data, which is used throughout the rest of the study phases. Extensive above-ground and limited underground tests confirm the usability of the strategy. The route identification strategy can identify a route being travelled with adequate accuracy within a short distance of usually less than 40 data points. This system finds direct practical application in optimising the energy management of an underground locomotive used by the mining industry in South Africa. Still, it can also benefit applications where route identification is required and using a GPS is not feasible.

In Chapter 3, a patented real-time mass estimation strategy that uses a simple torsional load cell on the driving axle of the vehicle is presented [7]. Submission of an article on this work will take place in 2024. This was the second phase of the project, labelled as Phase 2 in Figure 1.1. The presented method of measuring the vehicle mass can estimate not only the mass of a vehicle but also include the mass of trailers, if added, without using additional wheel load sensors for each axle. The system works by measuring the torque on the driving axle of the vehicle using strain gauges, and together with a barometric pressure sensor for road slope estimation and an optical encoder for distance it can solve the force balance equation for a vehicle in motion, with the only unknown for this study being the vehicle mass. The proposed mass estimation strategy has an error with a standard deviation over 27 tests of 5.2% in estimating the total vehicle mass for a road vehicle with and without a trailer. The mass estimation strategy proposed in this study will benefit hybrid and electric vehicle optimisation strategies, allowing better estimation of the total energy required to perform a route. It can assist commercial freight applications where total vehicle payload is governed by safety and legislation, with heavy fines potentially applicable for infringements on the maximum weight. In future applications, it may even be used in automatic vehicle tolling systems, reducing the need for weighbridges and easing congestion.

If the vehicle type, road conditions and overall mass are known it means the only controllable variable in energy consumption is the vehicle velocity profile. Chapter 4 presents a strategy that can robustly optimise the velocity profile for a vehicle on a known route with a known mass characteristic such that the overall energy consumed is reduced while still adhering to time constraints. This is the third phase of the project, labelled as Phase 3 in Figure 1.1. The results show that an energy reduction from 8.5 MJ to 5.7 MJ (33%) is obtained over a normal driver velocity profile for a specific 10 km test route having the same completion times.

In Chapter 5, the method of obtaining a good initial guess to the optimisation of Chapter 4 is used to construct a lookup table that can find an approximate solution to the optimisation problem for a vehicle on a known route in real-time. This real-time optimisation can read the data stream from the sensors onboard the vehicle, work out its location (Chapter 2), estimate its mass (Chapter 3) and provide a target optimal velocity profile for the rest of the route that will reduce the overall energy consumption. The methodology can gracefully accommodate unexpected changes in the velocity profile. Such as over-speeding or traffic in certain parts of the route. This is Phase 4 of Figure 1.1. The real-time optimal velocity profile can reduce energy consumption by more than 20% for a normal vehicle without regeneration ability.

1.5. Summary

This project aims to build bridges between the fields of mining, sensing and optimisation by providing a vehicle controller with real-time target velocities to follow to reduce the overall amount of energy required to complete a route. The proposed method makes use of low complexity calculations that allows for graceful adapting to unknown disturbances in the operating environment, like traffic, in real-time.

2. Underground Route Identification

2.1. Introduction

Recently an increasing emphasis has been placed on improving vehicle efficiency and reducing waste in all forms [8]. One such improvement is the use of hybrid and electric drivetrains in vehicles. Many Hybrid Electric Vehicles (HEV) comprise a combination of a fossil fuel Internal Combustion Engine (ICE), Electric Motor (EM), electric motor driver (power electronics that control the power to -and -from the electric motor), Energy Storage Systems (ESS), inverters (to facilitate the energy flow) and a vehicle control system to manage all of these. Hybrid vehicles are equipped with sophisticated control strategies that are applied to achieve higher efficiencies than normal ICE-driven vehicles by allowing the use of the electrical systems to store or recover energy that would otherwise be wasted as well as facilitating that the ICE can be operated in a more efficient region of its operating range [9]. When correctly applying hybrid technology to a vehicle fuel consumption improvements can be up to 40 % [10]. Fully electric drivetrains have all the added benefits of hybrids, except that instead of having the IC engine to convert the chemical energy of a fossil fuel to propulsion they use electric motors only, with batteries acting as energy storage systems. The application field of this study is underground electric locomotives used in South Africa [11].

Successful optimisation of a vehicle's energy system with electric motors added relies on knowing not only the current state of the onboard energy systems, including fuel level, battery state of charge etc., but also the future demands that the vehicle will encounter on its route. These demands depend on variables such as the vehicle, drivetrain, route speed profile, route topography and vehicle mass, all as a function of the route's distance to completion. The topography of a route directly affects the instantaneous energy requirements, as a flat route can have fuel consumption 15 to 20 % less than that of hilly terrain [12]. The advantage of having a hybrid system is also more emphasised by hilly terrain as energy when going downhill can be regenerated using electric braking systems [13], [14], [15]. It can thus be of great benefit if the route being travelled can be identified in real-time and its topography is known to allow the vehicle control system to make better management decisions for the vehicle's energy usage. The energy usage can be altered regarding the power source (in the case of hybrids or multiple electric machine (EM) vehicles), power levels, speed profiles and efficiency profiles.

Vehicle location above-ground can be easily and accurately pinpointed using a global positioning system (GPS), and route identification can then also easily be performed, but, the vehicle used as the focal point for this study mainly operates underground, in the hostile environment of a mine. The objective of this study is to present a method by which routes may be easily and cheaply identified underground with sufficient accuracy, using easily obtainable environmental data, and not requiring any transmitters or beacons to be placed on the routes in the mine.

The novelty of this method lies in the use of magnetic heading and barometric altitude data, which are easily measurable environmental attributes that do not change as a function of driving style or the vehicle used. A 3-axis magnetic heading sensor and a digital barometric pressure sensor were used together with a wheel rotation encoder to record patterns in the heading and altitude-change data, both as a function of distance. A simplistic statistical comparison model was implemented that

compares these values measured in real-time over standard distance increments to those of saved routes. Then, sequential patterns in the results are evaluated to perform route identification.

The Route Identification Strategy (RIS) proposed was first evaluated above-ground using a normal road vehicle and six routes of comparable length. The results showed that the RIS could correctly identify a route, usually within less than 20 % (4km) of the route's total length, with no prior knowledge of the location of the vehicle. The algorithm was set up such that the route's starting point is also not crucial, with a practical limitation that a route won't be identified if more than half of the route is already passed, as little optimisation benefit will be achievable past that point.

A compact handheld test instrument was constructed and used to perform equivalent tests underground to verify whether the RIS can function correctly when traversing underground. The first underground tests were performed in a basement parking lot. After it performed satisfactorily more tests were also conducted inside a small tourist attraction gold mine, where the strategy proved successful in identifying patterns in a route's data sampled there as well.

In summary, the project set out to develop a system that can identify a route that an underground locomotive is travelling on as well as knowing where the locomotive finds itself on that route. Using simple and low-cost sensing equipment, route attributes rather than vehicle and driver attributes were stored and patterns in this data were evaluated and compared to previously recorded data and routes were successfully identified, proving the usability of this system.

2.2. Materials and Methods

2.2.1. Hybrid Vehicle Systems

With the global energy crises, all forms of wasted energy need to be reduced and achieve the highest possible efficiencies for energy converters. One application attempting to achieve this decrease in resource consumption is applying hybrid drive systems to vehicles. Some hybrid vehicles can also be plugged in and charged off of mains grid power systems, potentially yielding even lower costs in vehicle operation. These vehicles are called Plug-in Hybrid Electric Vehicles (PHEV).

Hybrid systems can run at power levels higher than what the IC engine can deliver on its own. It is also capable of allowing the IC engine to run on average closer to its peak efficiency. It also allows operating the vehicle with the IC engine turned off for short periods if needed. It facilitates regenerative braking to recover some of the kinetic/potential energy and store it for later use by operating the electric motor as a generator to slow the vehicle down. Though beneficial on a flat road, simply running the engine at its peak efficiency still does not guarantee the most efficient operation [16], [17]. Tie and Tan [10] state that a full hybrid electric vehicle with a high-capacity energy storage system can achieve a fuel saving of up to 40 % without compromising performance. The true benefits of hybrid systems however can only be achieved if the energy flow is optimised for the vehicle, in real-time. Hu et al. [18] proposed a novel unified cost-optimal control scheme to consider all the contributing factors in HEV cost, able to yield a 28 to 40 % cost reduction. They considered mains grid charging cost, power management during driving, fuel cost and battery life models using rapid and efficient convex programming (CP).

A further step from hybrids is to go fully electric. In an electric locomotive, the IC engine is removed and only the electric motor, motor drivers and controllers and a significantly larger battery are used.

2.2.2. Route identification and currently applied methods

To optimally manage the energy usage of a vehicle it is necessary to account not just for the instantaneous power usage but also strategize a plan for future consumption based on the route travelled as it can yield significant advantages [8], [19], [20]. This can only be done if the vehicle has detailed information about the route parameters, like topography, required speed profiles, vehicle mass, distances, fuel/charging station locations etc.

The most general way of determining the location of an above-ground vehicle on a specific route is by making use of GPS. GPS's in ideal conditions can state the absolute location of an object within a couple of metres (<5 m), though this is not good enough for road grade/incline measurement. If the GPS has less than ideal conditions the GPS data may be substantially less reliable and sometimes even corrupted [21]. Many of the current route identification studies make use of GPS data for their algorithm [13], [22], [23]. Some authors supplement GPS data by feeding additional information into the route identification strategy. These additional parameters include inertial information obtained from accelerometers and gyroscopes, wheel speed sensors, steering angle, power demand levels, stop times, standard deviations for power demand, road slope, GPRS communication, laser displacement sensors to scan distances to objects and one used magnetometers to measure heading for an indoor tour robot application [22], [23], [24], [25], [26], [27]. Yokoi et al. [20] found that velocity profiles vs. time are not very useful in route identification as too many factors influence the time. Bender et al. [28] noted that the quality of the route prediction greatly affects the energy resource consumption of the vehicle.

GPS cannot work underground, so a RIS making use of other senseable data is required. From the references studied it was found that two variables can be measured with relative ease, which are route-specific variables, and are not dependent on the vehicle, driving style or time. These are the magnetic heading and barometric altitude, as a function of travel distance. Making it a function of distance also removes the travel time from the equation.

2.2.3. Topography/incline data acquisition

The real benefit of hybrid and electric systems is exploited when the vehicle travels over hilly terrain, where uphill and downhill sections are frequent or when accelerating and decelerating regularly [13], [14], [15]. Flat routes' fuel consumption can be 15 to 20 % lower than that of hilly routes for non-hybrid vehicles showing that road grade plays a significant role in vehicle fuel consumption and emissions. However, some optimisation strategies do not even take topography into account, in which case the solution to the problem is over-simplified. [5], [12], [29], [30], [31]. Road-grade data with a high enough resolution to be able to be used in simulations is hard to find. Li et al. [30] state that a barometric pressure sensor works better than accelerometers and other types of sensors for incline estimation, the latter of which's data are overwhelmed by the noise caused by vibrations during driving. However, they do note that barometric altitude is sensitive to weather. Boroujeni et al. [32] showed that a distance increment of approximately 0.1 mi = 160 m is an appropriate segment length for quantifying road grades as individual runs.

A barometric pressure sensor was used in this study to obtain an estimate of the altitude. From there the topographic profile and inclines for sections of the route are determined. An inexpensive open-source GY91 module sensor board and software were implemented. The GY91 comprises a BMP280 barometric pressure sensor and an mpu9250 3-axis accelerometer - 3-axis gyro and 3-axis magnetometer. These module boards retail for ±10USD (R200).

According to the datasheet, the BMP280 sensor can measure pressures ranging to a maximum of 1100 hPa (1.1 bar), which is sufficient for all above-ground applications but needs further testing to determine if it will survive in a mine. In ultra-high resolution mode, the barometric pressure sensor has a resolution of 0.0016 hPa, equating to a 12.6 mm altitude, with the RMS noise of the signal being able to go down to 0.1 metres [33]. The altitude value can be improved even further by using a low-pass filter. The barometric sensor's pressure reading is converted to an effective altitude using the International barometric formula, Eq. [2.1].

$$Altitude = 44330 \left(1 - \left(\frac{P}{P_0} \right)^{\frac{1}{5.255}} \right) \quad [34] \quad (\text{Eq. 2.1})$$

where P is the measured pressure, P_0 is the pressure at sea level and the altitude is then given in metres. In this equation, a pressure change of 1 hPa equates to a change in altitude of ≈8.4 metres at sea level.

Three simple indoor tests were performed to confirm the sensor's accuracy and usability for obtaining usable data for the incline and topographic profile.

2.2.3.1. Test 1: Barometric altitude vs. Time

(To investigate drifting of the altitude values over time)

For test 1 the sensor was left stationary on a table and the data was recorded over 24 hours. The resulting plot can be seen in Figure 2.1, top left. It is noticed that the absolute altitude varied by as much as 60 metres over 24 hours. This means that the absolute value for barometric altitude is unsuitable for determining the inclines of normal road profiles. When viewing the barometric pressure over a shorter time frame, Figure 2.1, top right, one can notice that the effect of the drift is much less (typically around 0.1 m per minute). From this result, it is decided to use the relative altitude between two adjacent points rather than their absolute altitude value, due to the time frame for travel usually being rather short compared to the drift rate.

2.2.3.2. Test 2: Barometric altitude during an elevator ride

(To verify the sensor's ability to detect small altitude changes smoothly)

In Figure 2.1, bottom left, the altitude is recorded when riding in an elevator, from floor 9 up to floor 15, down to floor 3 and back to 9 again. The red line is the raw data and the black line is a digital low-pass filter applied to the data afterwards. It is noted that noise in short time frames for the low-pass filtered data amounts to less than 0.75 metres. For a 100-metre stretch, the noise equates to an incline angle error of less than 0.5°. This shows that the barometric data is sufficiently sensitive to determine an

approximate topographic profile for a route as the ranges typically experienced are significantly larger than the sensor's resolution on barometric altitude. It is noted that this sensor's minimum detectable height limits the shortest distance increment for which it can be used to determine the incline effectively.

2.2.3.3. Test 3: Maximum simulated depth

(To verify that the sensor can work in a deep underground mine)

The whole testing device was installed into a sealed container and incrementally pressurised up to a maximum absolute pressure of 1.5 bar. A pressure of 1.35 bar absolute pressure equates to the pressure that it will experience at the bottom of the deepest mine in the world at the time of this study; -2500 metres below sea level, so slightly exceeding this deepest effective pressure would suffice to the proving of the device's ability to function as required [35]. The results for test 3 are shown in Figure 2.1, bottom right. It is seen that the sensor behaves predictably and gives the effective depth values as expected, returning to the test location altitude again after the test, showing that no sensor damage occurred. This is a simple over-pressure test to determine if the sensor can still yield sensible data at elevated pressures and is not intended as a detailed calibration test.

The barometric pressure sensor can now be implemented in a road-going vehicle which has an optical encoder to measure wheel displacement. This data can be used to obtain an estimation for the topographic plot for a route and be used to perform route identification. To have an above-ground comparison for the barometric altitude topographic plot and verify its usability, a GPS sensor is used to log the GPS coordinates (latitude, longitude and altitude) as well. It was noted however that the barometric altitude does show jumps of tens of metres in the data when the sensor is exposed to direct sunlight, so special care should be taken to avoid these sudden changes in sensor temperature.

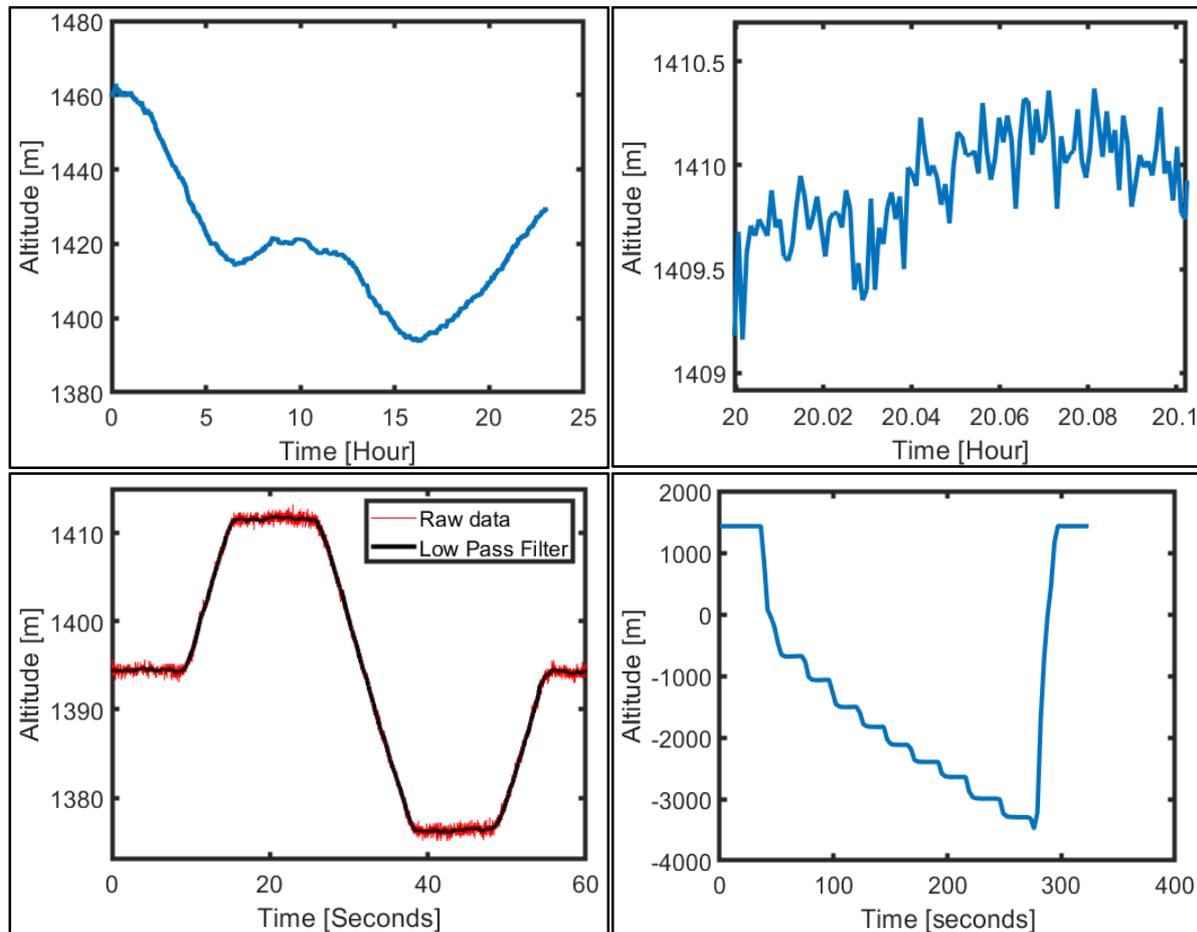


Figure 2.1: Barometric altitude sensor data

2.2.4. Heading data acquisition

Alvares-Santos et al. [26] obtained an estimate for heading from a 3-axis magnetic sensor. A magnetic sensor is a low-cost sensor that can give the directional orientation of a device. The drawback is that it is easily affected by external magnetic field disturbances like electric motors and metallic objects. Gallant and Marshall [36] propose a wheel rotation encoder to enable the heading data to be recorded as a function of distance.

The MPU9250 included on the GY-91 module board contains a high-resolution 3-axis hall-effect magnetic sensor that can determine heading accurately within 1° to 2° using its 16-bit built-in analogue to digital converter with a 4.8 Tesla maximum rating. The sensor has a high sampling rate, which allows for multiple samples to be taken consecutively to average/filter out higher-frequency noise in the sensor. The sensor can also be used in a strong magnetic field environment and still yields good heading accuracy as long as compensation is performed for the constant offsets [37] [38]. The heading noise of the raw data this sensor recorded during testing was approximately 0.5° .

2.2.5. Odometer

Initial testing was performed using an instrumented road vehicle above-ground and underground testing using an instrumented handheld surveyor's wheel. These test setups were fitted with an optical encoder

to obtain their displacements over the routes in real-time. As the odometer processor has a high-speed clock on it, it can also yield velocity, as a function of distance and time, which is useful in calculations relating to the route energy, which is covered in Chapter 4.

2.2.6. GPS as reference for odometer, velocity and barometric altitude

To prove that the barometric pressure sensor can obtain a topographic plot for a route accurately and to be able to calibrate the odometer, it is required to verify the results using a good reference. As all of the initial tests for the RIS are performed above-ground it is possible to use a GPS. The GPS module used can obtain a horizontal position accurately to less than 2.5 metres and velocity with an error of less than 0.1 m/s. Boroujeni et al. [32] found that GPS data is unreliable under bridges and overpasses, which was also found to be true during road tests performed in this study, especially when driving into a multi-level parking garage.

It is now possible to gather the topographic data for a route travelled by a road vehicle and compare the results to a GPS reference. The patterns in this data can be used in conjunction with patterns in heading data for various routes to implement the proposed RIS.

2.3. Calculations: modelling of the route identification system

2.3.1. Data Acquisition System

A pair of open-source Arduino prototyping boards were used to read and scale the data that is streamed from the sensors and in turn stream these values to a laptop that processes these values as input to the RIS [39]. One Arduino board was used as the master to facilitate serial communication between the sensors and a laptop computer, and the other was set up as a slave to serve as the odometer only. Once the preset distance increment for sampling of the heading and altitude is reached it transmits the total recorded distance to the main Arduino board. It was found that a distance increment of 100 metres yields good overall results. Upon receiving the distance value from the odometer board the main Arduino board requests the heading and barometric altitude 20 times consecutively in a very short time frame from each of the sensors to form a more stable averaged value for each parameter respectively, it then streams the averaged heading, averaged barometric altitude and distance to the laptop via serial communication as a string. The 20 samples were found to yield stable enough data for use in subsequent processing. The sampling rates for these two sensors are high enough to allow multiple samples to be taken consecutively without a significant calculation time penalty or distance travelled over the sampling time. Figure 2.2 shows a schematic layout of the operation of the data acquisition system. Operating the stream in this way has the advantage that the amount of data that needs to be processed is low enough that the route identification program can be run in real-time.

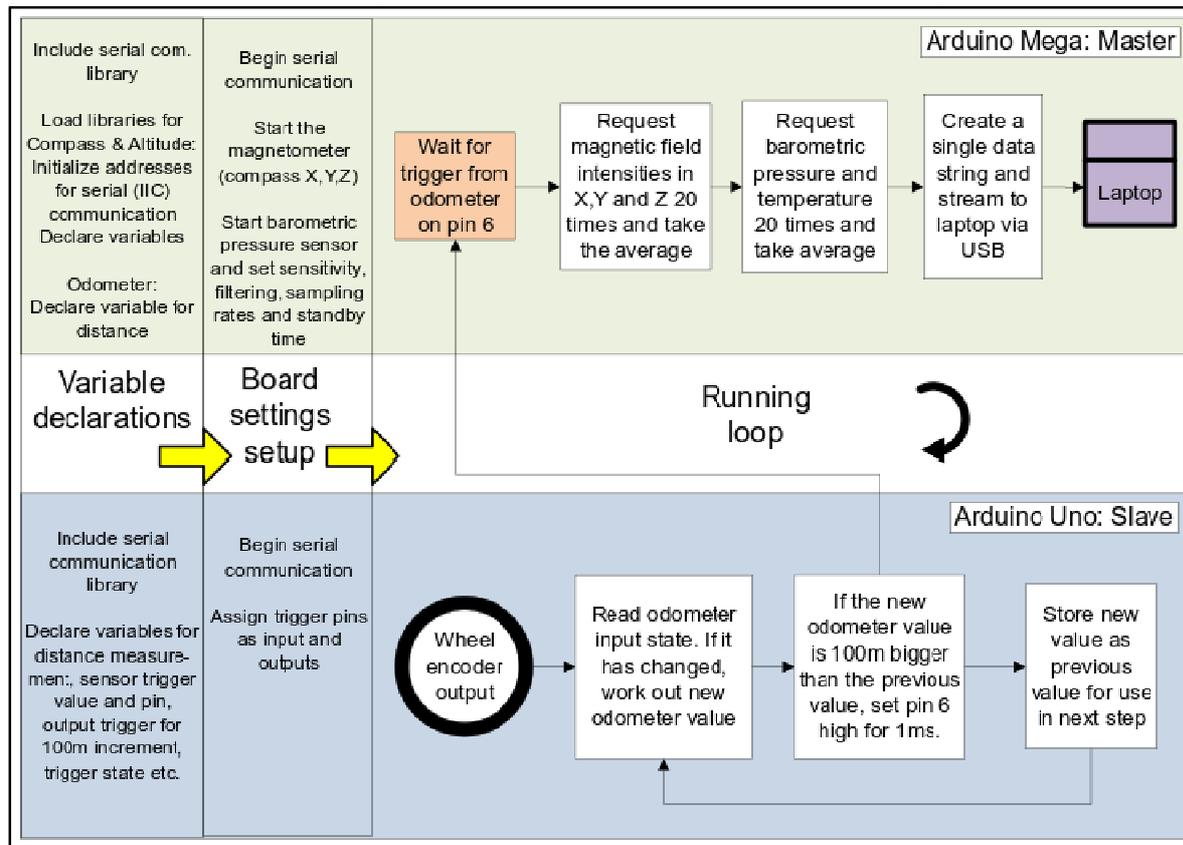


Figure 2.2: Arduino data acquisition system operational schematic

Figure 2.3a shows the GPS -altitude vs. -distance and the barometric altitude vs. Optical encoder distance. It should be noted that both the data for barometric altitude, as well as the GPS altitude, were zeroed at the start of the run to align with each other, the barometric drift was explained in Chapter 2.2.3.1. The barometric data lines up very well with the GPS data. Figure 2.3b shows the zeroed barometric altitude vs. distance for the same route driven on three separate days, showing very good consistency between the three tests.

It is noted that the GPS data signal is not without fault, as there is a jump at the end of the data, and many times missing data was observed when travelling underneath a bridge/overpass. The jump at the end of the data was when the vehicle entered a covered parking building and lost a proper signal for the GPS to function correctly.

To alleviate the issue of zeroing the barometric altitude and absolute altitude accuracy that varies with weather conditions, it makes sense rather to use the derivative of the function, i.e. the change in altitude over a predefined relatively short distance/time increment. Altitude gain/loss over a small distance in a relatively short time will not noticeably suffer from weather and other environmental effects as was shown in Figure 2.1. Figure 2.4a shows the raw incremental change in barometric altitude, termed the altitude gain, for three passes of the same route. The altitude gain also allows direct calculation of the incline of the route, which is beneficial for mass estimation performed in Chapter 3 and the energy calculations performed in Chapter 4.

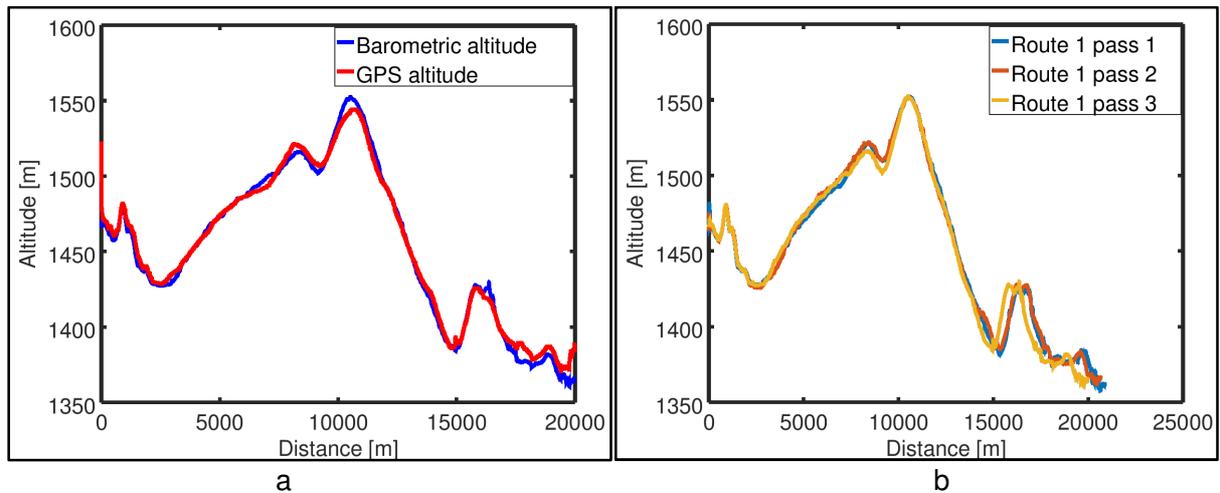


Figure 2.3: Topographic plot. a) Barometric & GPS vs. distance. b) 3 Passes of the same route

The heading was also recorded for the three passes to prove whether a good correlation between the data for the same route on different passes is found, as shown in Figure 2.4b. It is noted that the heading data shows very good consistency for the three passes of the route travelled. Headings near magnetic north cause the abrupt jumps in the data from 15000 metres to the end. As an example, a heading of 1° is only 2° away from a heading of 359° , though it isn't intuitive from observing the figure it is easily accounted for in a later step when the comparisons between routes are done in the results Chapter 2.4.

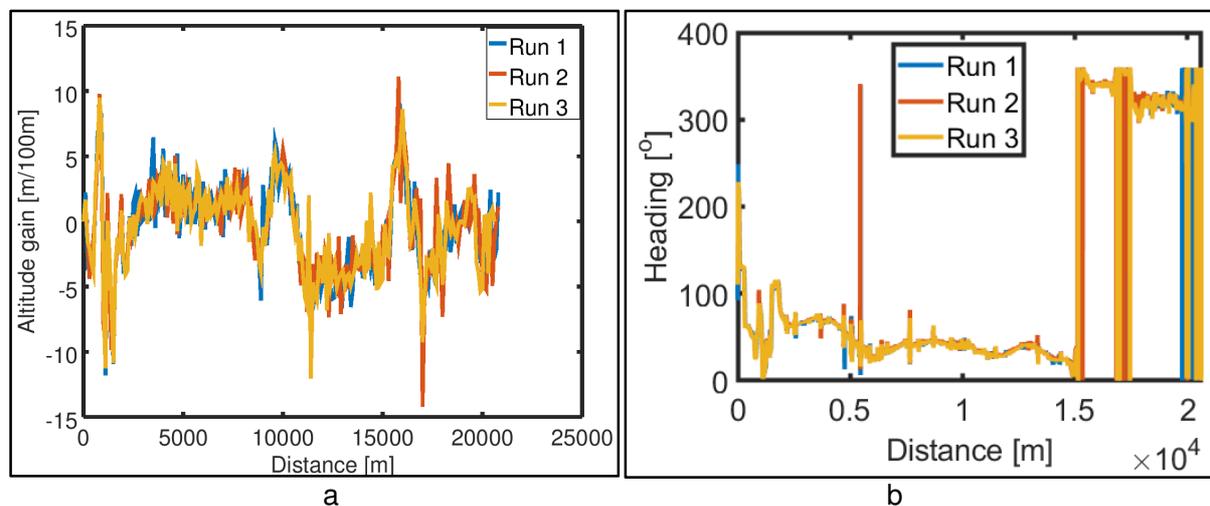


Figure 2.4: Route data for 3 passes of the same route. a) Altitude gain, b) Heading

2.3.2. Above-ground tests

The aim is to develop a system capable of identifying the route that is being travelled based on patterns in heading and altitude change data. Once these parameters were recorded above-ground and found successful in route identification it is verified for underground use. An increment of 100 m reduces the amount of data that requires processing compared to higher sampling rates and also reduces noise generated in the altitude gain data. The low sampling rate has one drawback; when the heading or

altitude changes sharply, like on a residential winding road, it will cause the ability to perform route identification to completely diminish as the exact location of the sample point around a bend is not the same from one pass of the route to the next, which will lead to unusably high differences in heading data.

If the topographic plot of a certain route and the vehicle's location on that route are known, it also allows the vehicle control system to make better decisions for its total energy management, which will be covered in Chapter 4. This study is based on an underground railway vehicle, as stated in the introduction, so the route options are fairly limited as compared to normal road vehicles which can do u-turns and take side roads, thus greatly reducing the number of possible routes to search for. This is beneficial to limiting the number of routes possible during route identification. Figure 2.5 shows a schematic diagram of how the route identification program works. The program is executed in real-time on a normal laptop. The software used is Matlab [40]. These steps are explained in the paragraphs to follow.

A real-time route comparison can only be performed if there is information available for at the very least one route previously travelled. The first route was saved from the data recorded for the feasibility study of the heading and altitude change data. This data was saved into two text files, one containing the heading and the other the altitude gain information. With the distance increment being constant at 100-metre increments it allows the saving of the heading and altitude data without their distance reference, as each new line of saved data is by design stored at the corresponding distance increment for comparison to the real-time data streaming in.

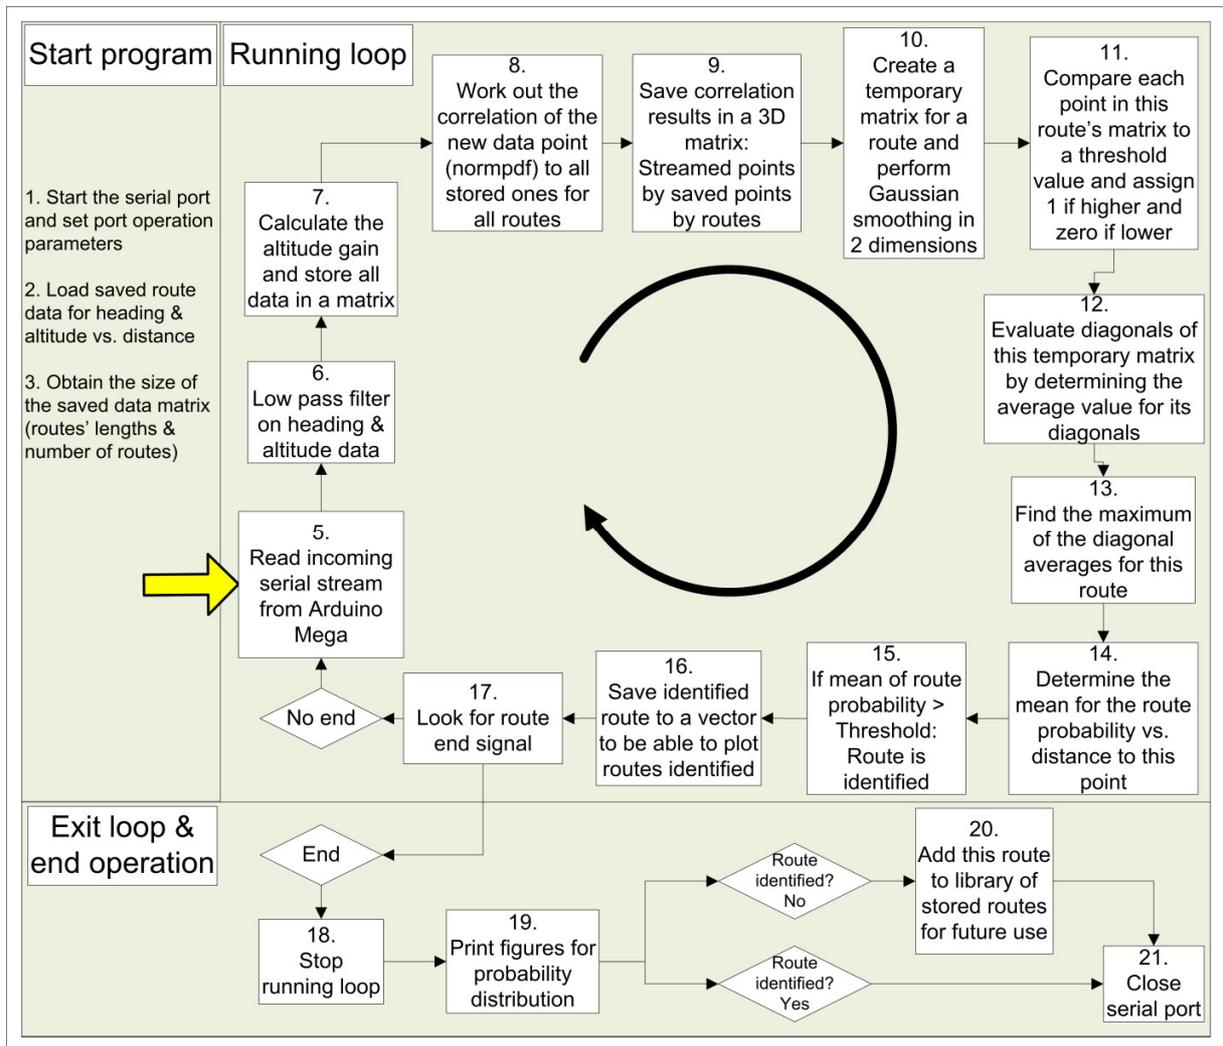


Figure 2.5: Matlab program operation schematic

With each new heading and altitude data point streamed in, the laptop in the vehicle can compare these new values to the recorded set of values for routes previously travelled and saved. The program makes use of a normal distribution function (normpdf) to yield a high value for good correlation and a very low value for bad correlation. Each new data point is compared to all the previously saved routes' data points on a point-by-point basis. Though this is a tedious process, computationally speaking it is easily accomplished before the next data points are streamed in a couple of seconds later. There is an advantage of always comparing all of the points of all of the routes as this allows the route identification program to identify a route even if the starting points do not line up, or to distinguish between routes that may have the same or very similar portions. Points are compared by calculating the value of the probability density function at point X of the normal distribution with a mean of μ and a standard deviation of σ . The variable X is the incoming data point, μ is the saved data point's value and σ is the standard deviation. The value for the standard deviation, σ , was an assumed constant value, which was determined to work well during simulations of the route identification strategy. The MATLAB function that is used for this is $\text{normpdf}(X, \mu, \sigma)$. This is not a purely statistical method, however, the normpdf

function is a very convenient method of comparing the correlation between data points. It is normalised for each data point to ensure that the maximum probability density function value given by the function can not exceed 1.0 for an exact match. A sensitivity analysis was performed to determine the best values for σ based on route identification success rates. A low value for σ will make the route identification program stricter, up to a point where no routes are ever identified as very few points will have exactly the same heading and altitude gain values, and a high value has the opposite effect where routes would be identified incorrectly. For the simulations performed a value of $\sigma=8^\circ$ was found to perform well for heading and a height change of $\sigma = 3 \text{ m}/100\text{m}$ performed well for the altitude gain.

The altitude gain correlation value for a point streamed in is compared directly to all of the points in all of the saved routes. This will yield an $N \times M$ matrix which contains the correlation values, with N being in the direction of the data for the saved routes and M being the data in the streamed in route direction. The same point-for-point method is performed for the heading data, but additional computations need to be carried out. For the heading correlation, it is necessary to account for headings close to north specifically, as was shown earlier in Figure 2.4b. Generating this correlation matrix is accomplished by taking the probability for each heading point streamed into the saved routes' values, and then doing the same again, but this time using a heading value of $(\text{heading} - 360^\circ)$. This will allow a 358° incoming point also to be calculated as a -2° . A third scenario is also applied where $(\text{heading} + 360^\circ)$ is used so that a streamed heading value of 1° can be compared to a saved value close to, but less than 360° . The biggest of these three heading correlation values for each point is stored in the heading correlation matrix and is used in further calculation steps.

Now that a correlation value is determined for both the heading and altitude gain, point-for-point for all of the saved data points, a resultant correlation matrix accounting for both heading and altitude gain at the same time is determined by simply multiplying the two correlation matrices, point-for-point. Thus, only if both the heading and the altitude gain data points correspond well to the saved data point does it register as a high correlation. These correlation values are recorded and saved in a 3D matrix, where one axis represents the number of routes evaluated, the other axis represents the streamed data and the third axis the saved data.

If the route being travelled correlates to one of the saved routes, there will be a high band in the route's plane in the 3D matrix. The streamed data and saved data will have to be in phase with each other and yield consistently high correlation values for a route to be identified. The route planes in the 3D matrix are evaluated by taking the averages of bands in these matrices, where the bands are identified as $\text{Matrix}(a+X,a)$. The value for "a" ranges from 1 up to the shortest length of the plane and the "X" is varied from the negative of half of the shortest length to the positive of half of that length. This in effect determines the averages for lines parallel to the diagonal of the matrix. The reason for not going all the way to the end of the matrix will be explained in more detail later.

Figure 2.6a shows the resultant correlation matrix plotted for a case when the data of the route travelled is compared to itself. Figure 2.6b shows the travelled route compared to a different route stored. It is noted that no pattern is visible for the mismatched route, which is exactly what we expected. The bands are only evaluated in a positive direction in both the saved and streamed data axis

as we are not evaluating travelling the route backwards at this stage. A band which has a consistently high correlation indicates a good chance of travelling on the specific saved route to which it is being compared. Not only the central band but also bands parallel to that need to be evaluated as the exact start points for the saved and streamed routes do not have to match. The parallel bands are important to evaluate as this will allow route identification even when the starting point of the routes is not the same or when there exist slight distance discrepancies between the saved and current test stream. There is however a limit to the minimum distance that a route needs to correlate with before it is claimed as identified. This restriction is necessary to neglect the false positive route identifications if only a couple of consecutive points line up close to the start or end. If the program was not able to find a good correlation between the travelled route and any of the saved routes it was set up to save this route as a new route, the data of which will then be used in the route identification steps in future passes.

It is of benefit if the route identification strategy is somewhat lenient in terms of the exact band that correlates to a route, for example, if a set of data is offset by 50 metres relative to the saved route it might cause the high band values to float between two adjacent bands, depending on the exact point being considered, which will cause the band's average value to drop drastically, even though the same route is being travelled. This situation can occur due to wheel slippage, or slight offsets in the starting points for rail-based vehicles, with lane selection playing a role in road-based vehicles. The strategy implemented to make route identification more forgiving and robust is to make use of a smoothing filter on the route's correlation matrix, which means that a route can be identified even if there does not exist a single band with high correlation values. A Gaussian filter was implemented for this purpose. This filter will also pull up local low points which might have been caused by a local error, for example when travelling around a sharp corner and the two data points which are being compared don't line up perfectly around the corner, yielding almost zero correlation for that local point due to the heading correlation value being close to zero.

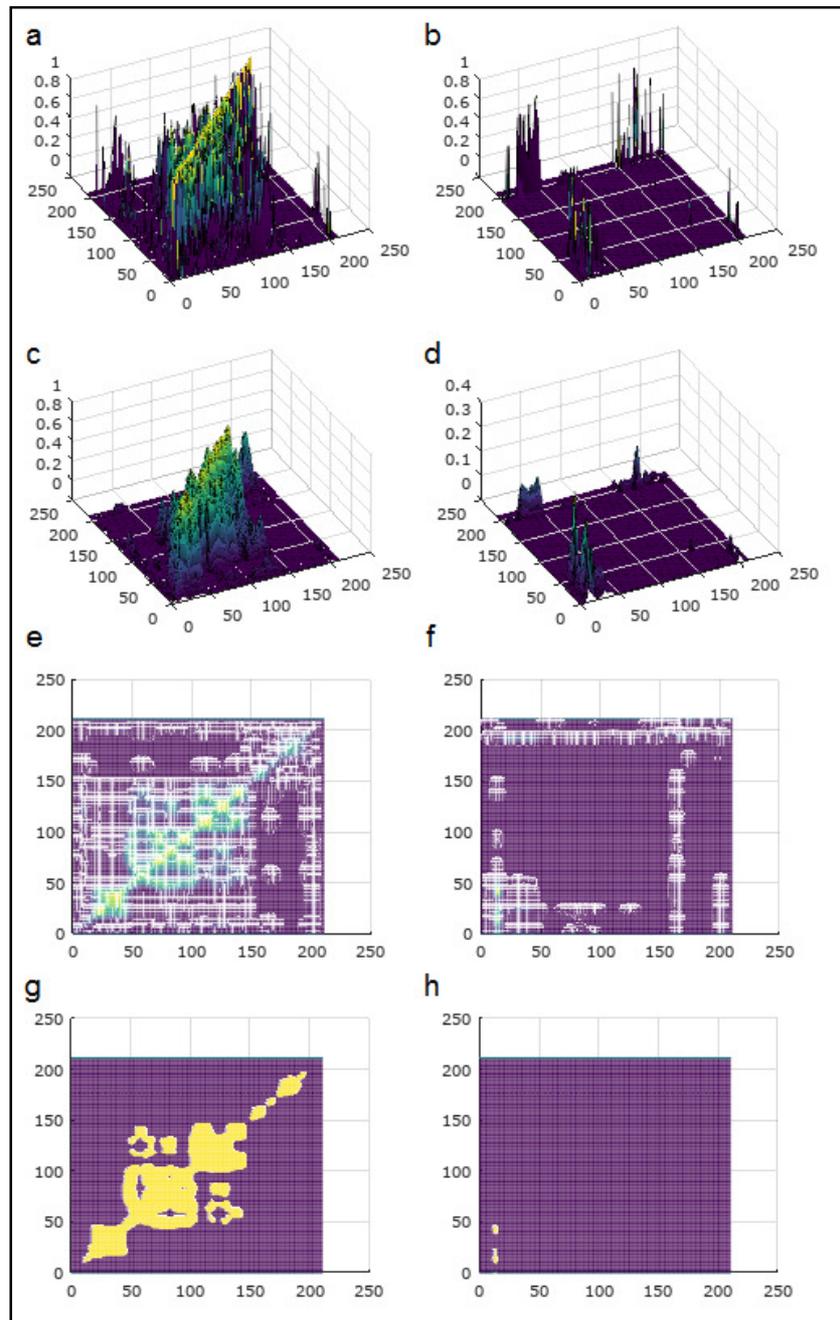


Figure 2.6: Probability matrices processing steps. Left: Same route. Right: different route

The smoothed data is shown in Figure 2.6c and d, for the same and another route respectively. This smoothing strategy will reduce the effect of starting value differences, small amounts of wheel slip and other factors that cause the exact distance correlations to deviate from the single high band. A point to note however is that this smoothing filter also reduces the maximum probability values for route points which correspond 100 %, due to the neighbouring ones that draw it down slightly. Figure 2.6e and f show plan views of the smoothed probability, for the same route (e) and another route (f). A high correlation band is clear for the similar route (e) and no high bands are visible for the other route (f). A

potential issue with evaluating the bands based only on the highest average value could be that a band with a consistent medium probability might be outweighed by a diagonal with generally low probability and some local high probabilities, or some consistently low route may even end up being the highest average probability and incorrectly identified as a route, simply because it happened to be the highest probability of the data available. To avoid this potential error in averaging of correlation values a threshold value to the probabilistic matrix was applied, causing the probability for a point to be considered as either a 0 (no correlation) or a 1 (sufficient match), based on some experimentally determined constant threshold cut-off value. This process is illustrated in Figure 2.6g and h. Now the band values can be compared and if found to exceed a preset value it can assume a positively identified route.

To avoid erroneous route identifications based on the average band approach it is important to stay in the central region of the matrix, and not to calculate the averages for the endpoints. For example, if a distance travelled equal to the final point of the route is considered, in the real time program no further information regarding the route will be available yet, and thus the diagonal average value for the endpoint of the matrix, with a length of 1 cell might have an average = 1, which is the maximum value and thus the route will be "confidently" identified based on a single data point. This will lead to many incorrect route identifications and make the whole system unstable. To alleviate this problem not all of the bands are considered, but only the central region's ones. For the tests performed it was found that using half of the matrix width in either direction for the bands mostly solved this problem of incorrect/erratic route identifications. In general route identification more than halfway through the route would not be very beneficial in an optimisation strategy in any way. If, by the end of a route no routes were identified the system automatically saves this route as a new route that will be used in future passes to perform route identification.

There were eight routes used to evaluate the effectiveness of the RIS. Most of the routes were back and forth between points A and B, as shown in Figure 2.7. These routes were plotted from the GPS data recorded during testing. It should be noted that the routes overlap close to the start and end points, so initial identification of a route may be incorrect, as it is impossible to tell where you are going to turn off and continue with the planned route. It is important to remember this when the route optimisation is performed later on as it directly influences the route optimisation that needs to be performed and can be detrimental to optimising if the incorrect route is optimised for initially. Most of these routes are tarmac public roads. Route R5 is a tabletop test, where the distance travelled is theoretically simulated, although no real movement took place. The advantage of using the constant valued route R5 is that it facilitated fault finding and refinement of the program during tests/simulations performed without needing physical testing on the road. The trip distances of the routes between A to B are stated in Table 2.1. Route 2 was an around-the-block route to verify sensor data and was not travelled for route identification purposes.

Route	Trip distance (km)
R1	20.62
R2	2.6 (point A - point A)
R3	21.3
R4	22.1
R5	N/A
R6	19.5
R7	21.2
R8	22.6

Table 2.1: Test route lengths

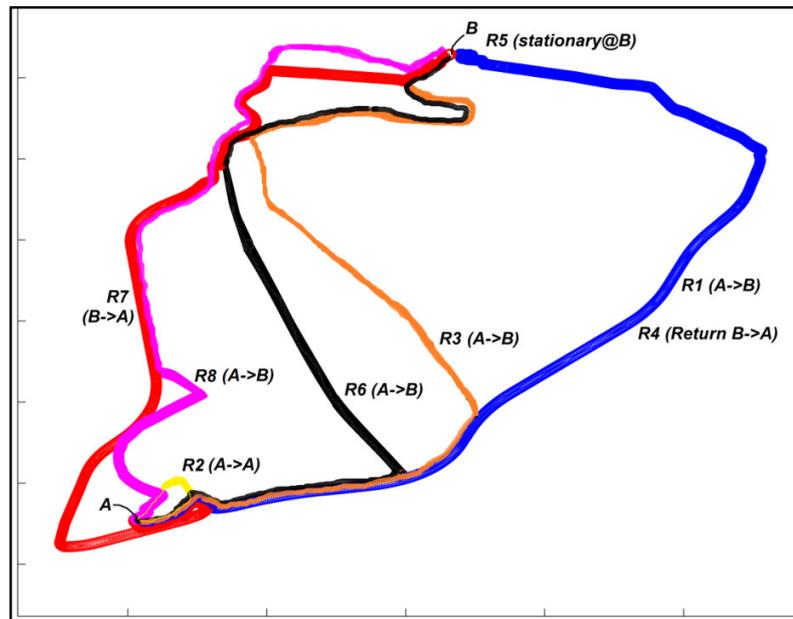


Figure 2.7: Above-ground routes travelled

A sensitivity analysis was performed through simulations and the most usable and safe values for accurate route identification are as follows: $\sigma_{\text{Altitude gain}} = 3\text{m}/100\text{m}$, $\sigma_{\text{Heading}} = 8^\circ$ and threshold cut-off value = 0.2. These values are then used to perform above-ground real-time route identification for the 8 routes. The results of the above-ground route identification system are presented in the results section.

2.3.3. Underground tests

This project found its application in an underground locomotive, which justifies the necessity to evaluate the RIS underground as well. Obtaining access for gathering of vehicular test data in an underground mine was not possible due to security and safety concerns. The best alternative was to build a small handheld device that could be used to record heading vs. distance and altitude vs. distance data in an underground environment to prove that these parameters can be effectively measured underground and that it forms patterns able to yield the same RIS success as it did above-ground. The handheld device used for the underground tests was an instrumented surveyor's wheel, fitted with an optical encoder, magnetometer and barometric pressure sensor, but in this case, the real-time laptop communication is replaced by recording the data on an SD card and processing afterwards due to space and practical reasons. Figure 2.8 shows the handheld device. At the bottom of the device, just above the wheel, is the battery. The battery needed to be placed far away from the magnetometer as it was found that close proximity causes a catastrophic offset in the heading data, rendering the whole test unusable if not specifically recalibrated for this offset. The magnetometer/barometer were mounted on top of a skeletonised soft foam suspension block to reduce accelerations on the sensor, which were found to increase data noise in the barometric pressure sensor. The surveyor's wheel has a factory-calibrated distance readout, which makes calibration and quality checking of the recorded data much easier. For the above-ground tests, the distance increment was set to 100 metres, but for practical reasons, the handheld device had the distance increment reduced to 5 metres, as walking many kilometres underground is not very practical. It was recognised that the altitude data will now yield noisy incline

data due to the minimum elevation accuracy mentioned in Chapter 2.2.3. It has already been proven that the heading and altitude gain strategy works above-ground, so in-detail underground tests are not deemed essential, rather, simply proving the device's ability to detect appropriate heading and barometric values underground to sufficiently facilitate route identification is performed.



Figure 2.8: Surveyor's wheel for underground testing

2.3.3.1. Basement parking lot test

With a working data acquisition system, some route identification tests were performed underground. The first underground tests were performed in the basement parking of the Engineering 3 building, at the University of Pretoria's main campus, on level P1, two levels below ground level. The building is a steel-reinforced concrete structure with no external steel structural members. Mentioning this is important as it could affect the magnetic data. A circuit on level 1 was walked four times with the device reset at approximately the same point on each pass. The circuit started on a flat section, proceeded down a slope, back up the slope again and flat to the start point of the route. The results for route identification in the basement tests are included in the results section.

2.3.3.2. Gold mine underground testing

The testing of the device to determine if it functions in an actual gold mine was performed at Gold Reef City, an amusement park located in Johannesburg, South Africa. It has a small decommissioned gold mine on site, which is now only used for tourist and educational visits. The purpose of performing a test in a real mine, though it is very small compared to the commercial active mines is to verify that the sensors will yield information that can be used to perform route identification successfully. The same surveyor's wheel setup that was used for the basement test was used for the underground mine test. The mine's tourist level is located approximately 76 metres below the surface of the ground. Only two passes were granted for testing purposes. The results for route identification performed in the Gold Reef City mine are presented in the results section.

2.4. Results and Discussion

2.4.1. Above-ground route identification

Eight different routes were evaluated. Route 2 was just a short test route to evaluate the sensor function and the data being streamed. Route 5 was a theoretically generated constant route to allow

finer testing and calibration of the system. In Figure 2.9 the route correlation values are plotted against distance for all of the saved routes when travelling on route 1. It is noted from Figure 2.7 that routes 1, 3 and 6 all have the same initial portions, which is why it makes sense that routes 3 and 6 both also had increasing correlations up to the point where they deviate from route 1. This proves that the route identification strategy was able to correctly identify the route travelled with high confidence by 4.6km, equating to less than 20% of the total route length and with moderate confidence by 2.5km (25 data points).

Route identification tests were performed multiple times for every real route to evaluate the route identification strategy's ability to perform satisfactorily. Table 2.2 summarises the results showing the route travelled, the route(s) identified along with the distance where it was identified and the number of times each route was tested.

From the test results shown in Table 2.2, it is seen that the route identification strategy can identify the correct route in every test performed, in real-time, with small errors occasionally found in the initial parts of the route. In the case of route 7, route 5 was identified twice initially for small parts of the route due to similarities between the routes' data, though this error was corrected within 400m and the correct route was maintained until the end.

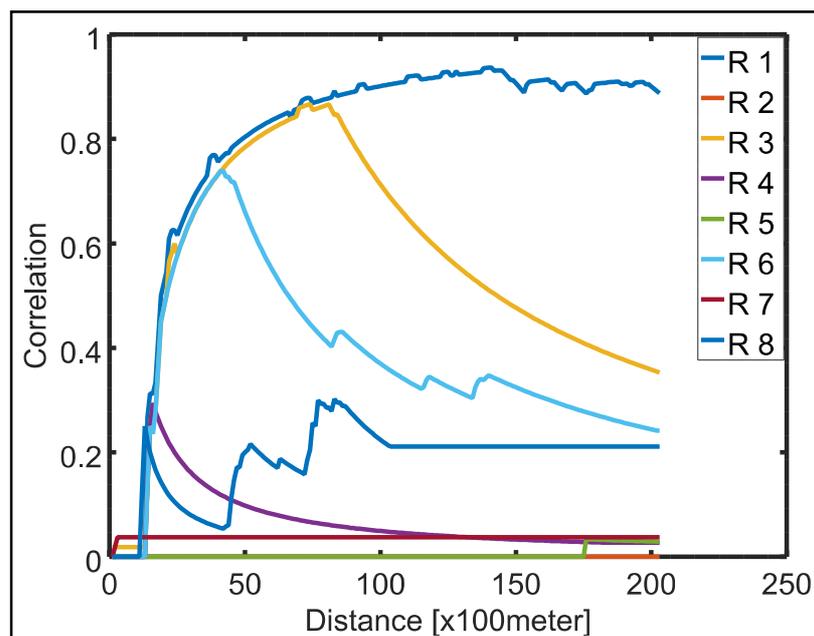


Figure 2.9: Correlation vs. Distance when travelling Route 1

Route number	Route identified & distance to successful identification	Number of hits
1	R3@800 m, R1@8.5 km R1@800 m R1@4.2 km	3
3	R6@4.4 km, R3@8.5 km, R1@4.6 km, 6.8 km R3, R3@700 m, R6@4.3 km, R1@4.4 km, R3@8.4 km to the end	4
4	R4@500 m R4@500 m R4@800 m R4@1.1 km	4
6	R6@4.6 km R6@3.4 km	2
7	R5@4.0 KM, R7@4.1 km R7@3.5 km R7@3.5 km R7@1.7km (R5 from 2.8 km - 3.2 km), R7@ 3.3km to end R5@3.0 km, R7@3.4 km to end	5
8	R8@500 m R8@3.1 km R8@800 m	3

Table 2.2: Summary of route identification results using $\sigma_{alt\ gain} = 3m/100m$, $\sigma_{heading} = 8^\circ$

The reason why the identification correlation value starts at zero is due to the program being set up such that it starts at zero for all routes to not favour any route above another without proper knowledge of the route being travelled. With it using an averaging program the initial zero correlation gradually grows as the correlation of the route with a stored one increases over distance. The routes are usually more winding at the start and towards the destination due to parking lot driving and sharp corners followed in residential areas, where the under-sampling of the route data at 100-metre increments reduces the route identification strategy's ability to lock onto a route as well.

Due to the ability of this system to run in real-time, it allows the route identification strategy to not only identify the route being travelled but also to know the location of the vehicle on that route (route progress), which is essential when optimising the energy management of a hybrid/electric vehicle, performed in Chapter 4 and 5.

2.4.2. Basement test results

The heading and altitude data for these passes can be seen in Figure 2.10a and 2.10b respectively. Good consistency is noted in the data for the four passes. It should be noted that being hand operated this data will not be as consistent as the data obtained by a vehicle driving on a public road with lane directions clearly indicated and closely adhered to, and even less consistent than what a rail-bound locomotive will see.

Four passes of the same route were used for the test, the first one is saved as the reference, with all 4 routes compared to this stored route. Comparing the route to itself is handy for consistency checking. Figure 2.12a shows the correlation as a function of distance samples (5-metre increment). From Figure 2.12a it is seen that the route that was walked 4 times in the basement parking level 1 yields data that is usable for the route identification strategy proposed. The noise introduced due to it being hand-operated did cause a reduction in the smoothness of the correlation graphs, as is expected. The reason why a 100 % probability is not achieved with route 1 compared to itself is due to the Gaussian filter as was explained in the above-ground tests in Chapter 2.3.2.

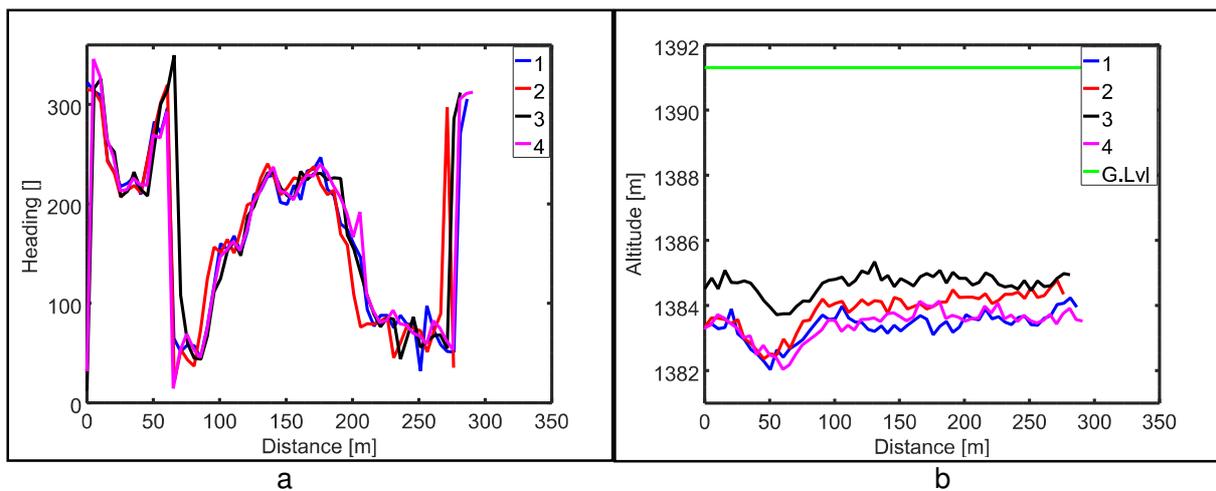


Figure 2.10: Basement parking level 1 test. a) Heading-, b) barometric altitude vs. distance

It can now be concluded that the route identification system yields satisfactory results for an underground basement test using the handheld data logging device. The next step is to perform tests in a mine, located at deeper depths and in the typical mineral rocks that mines are excavated from to confirm the applicability of this system in underground environments.

2.4.3. Gold Reef City Underground Test Results

Only two passes were granted in a small underground gold mine on the tourist level to determine whether the heading and altitude data could be used to identify a route being travelled. In the top of Figure 2.11, the sensor data can be seen for a trip from the surface, through the mine's underground visitor's loop twice, and back to the surface. It is noted that the initial change in heading was as we approached the elevator shaft, and the jump on the vertical black line is due to us walking out of the elevator in the opposite direction with the test data for the first pass starting directly after that. There is also a jump present in the altitude data as the lift travelled down without us walking. The data is then split into the two individual passes plotted over each other. The progress stages are indicated by the three vertical lines in the top figure. The altitude is shown in this case instead of showing the altitude gain as the noise is too large on the small distance increments used resulting in figures which are hard to see any patterns on.

In the bottom left of Figure 2.11, it can be seen that the heading data shows a very close correlation between the two passes. The altitude data shows a good correlation as well, although the data is much

more noisy when one zooms into the local values. The noise in the data amounted to around 0.5 metres, which corresponds to the accuracy stated in Chapter 2.2.3. The route identification algorithm can be set up to allow for higher noise in the barometric data by using less strict statistical parameters (larger σ value). Though it is noisy this does still enable comparison of patterns in the data.

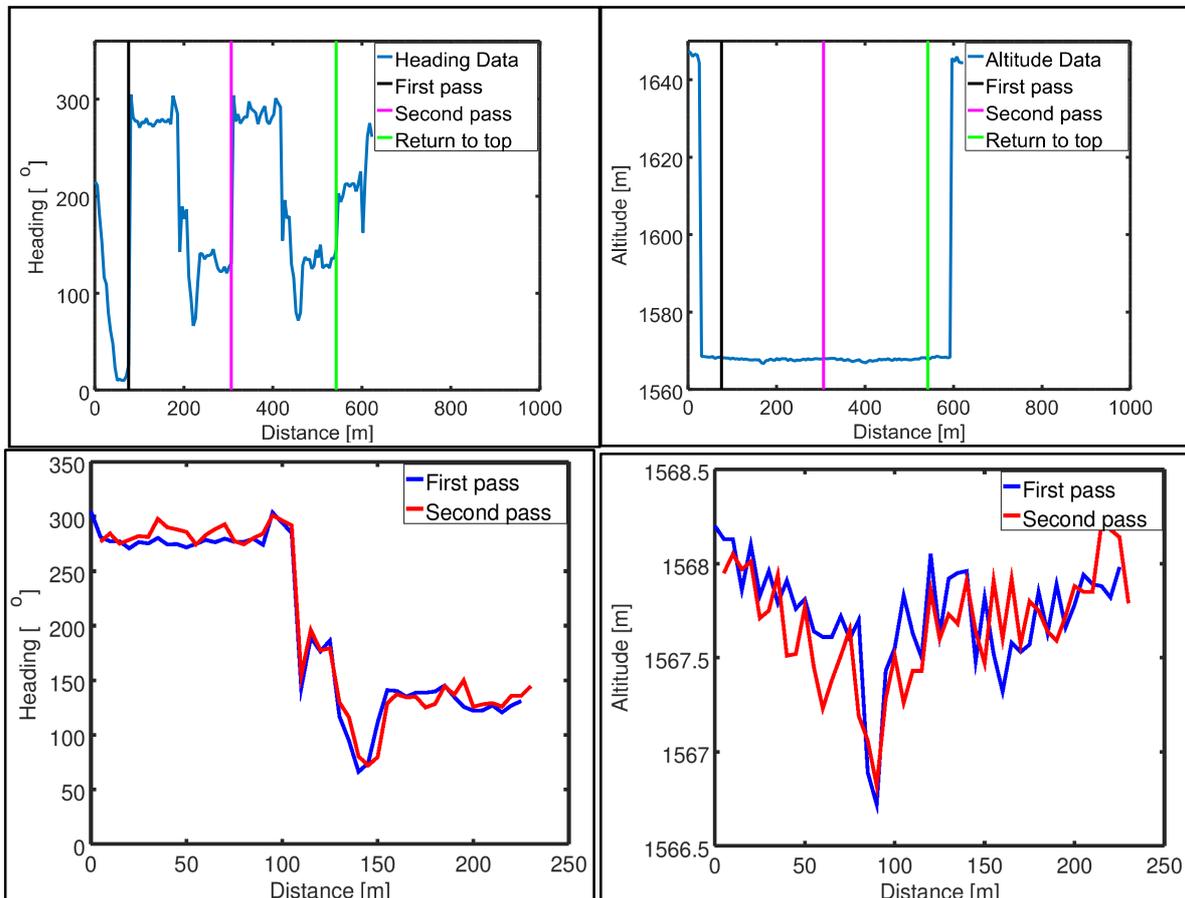


Figure 2.11: Underground gold mine test correlation

The correlation vs. distance plots for the underground gold mine tests are given in Figure 2.212b. Pass 1 is compared to itself and then the second pass is compared to the first as well. The reason why it does not give 100 % all the way through is due to the Gaussian filter, as was discussed previously. It can be seen that the RIS very quickly identifies pass 2 as the same route as pass 1, taking only around 4 samples, which equates to a 20-metre distance, or less than 10 % of the loop distance and consistently identifying the route from there onwards with high confidence.

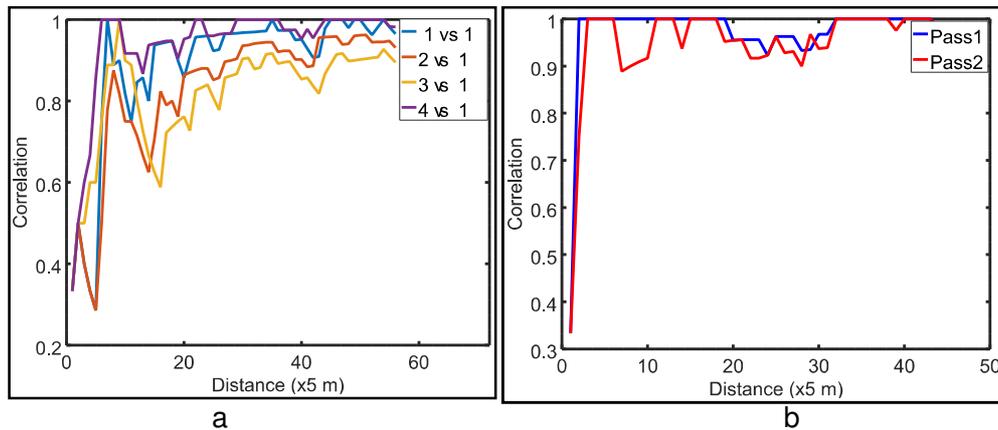


Figure 2.12: Underground RIS results. a) Basement tests, b) Gold mine tests

2.5. Conclusion

A novel method of identifying routes underground and finding the location on that route is presented. This method makes use of easily and affordably obtained magnetic heading and barometric altitude values measured as a function of distance whilst driving the route. This automatically also yields a topographic profile for the route, which is used in later processing to obtain the vehicle mass (Chapter 3) and route optimisation (Chapter 4). A pattern recognition program is developed that can compare patterns in data for a route being driven to a set of routes that were driven previously and stored into a simple small text file in memory. This strategy is able to successfully identify a route being travelled by a vehicle as well as knowing the location of the vehicle on that route, for a total of 6 real-life route options investigated. The location on the route is known simply by the most recent data point streamed in, once the route is identified. If the route is not known to the system it automatically saves this route as a new route to the small text file, which will then be used for route identification comparisons in future driving cycles. The system proves to function both above-ground and underground based on the tests performed. This proves that the system allows route identification to be performed even in an underground environment where a GPS is not available. A sensitivity analysis was able to yield the optimal parameters for which the route identification strategy yielded the best results. The route identification strategy was able to always converge to the correct route within a short distance, typically within the first 20% of the route travelled.

2.6. Future work

Improvements can be made to the system by using better filters for the incoming magnetic data, which will yield more stable heading values, allowing the route identification strategy to identify a route sooner than before. The constant value for the standard deviation, σ , used in the normpdf equation can be adapted to be determined from route data, if more consistent data is expected in a certain region, this value can be reduced to perform more accurate correlations to the data.

Now that a method has been developed that is able to easily identify a route by, without the use of external communication, the next step is to develop a method that is able to estimate the mass of an entire vehicle rig in real-time, which is required to estimate the energy requirement of a vehicle. This is presented in Chapter 3.

3. Real-time gross mass estimation

3.1. Introduction

3.1.1. Background

With limited environmental and financial resources, it has become more and more essential to reduce the consumption of fossil fuels in the operation of heavy commercial vehicles to move towards a cleaner and more sustainable future. Hellström et al. [41] reported that up to 30% of the life cycle cost of operating a large diesel truck can be attributed to diesel consumption, proving that a small change in fuel consumption can have significant financial and environmental effects as these trucks travel hundreds of thousands of kilometres annually and obtain a fuel consumption of around 3km per litre of fuel. In an attempt to reduce the consumption of fossil fuels an approach has been to implement a hybrid power train strategy where an electric motor, electric motor drivers and an energy storage device are used together with the internal combustion (IC) engine to improve fuel consumption. Hybrid power-trains can store and convert energy that would have been lost in a normal fossil fuel system. These energies could be from having to reduce speed (kinetic energy) or to change elevation (potential energy). Usually, these excess energies will be converted into heat using some braking mechanism. Hybrid systems are also able to power the vehicle or supplement the fossil fuel IC engine. Many other strategies where more efficient consumption is possible using the hybrid system are also applied to hybrid vehicle systems [16]. Tie and Tan [10] reported that a full hybrid electric vehicle with a high-capacity energy storage system can achieve a fuel saving of up to 40% without compromising performance. A step up from this would be fully electrifying the vehicle and running it solely from a battery. In a fully electric vehicle, the benefit of having a real-time mass estimation mechanism is of similar benefit to that of a hybrid vehicle system.

In the equations used to determine the energy of a vehicular system, the mass plays a very large role, with the energy, whether kinetic or potential being directly proportional to the mass of the vehicle. The mass also influences the traction abilities of the vehicle. Vahidi et al. [21] reported that the mass of a heavy-duty vehicle can vary by as much as 400%. With this vast variation in mind, it would be almost impossible for a hybrid/electric vehicle control system to optimise energy usage without having a good estimate for the vehicle's overall mass (gross mass or gross combined mass). The overall mass includes the mass of the vehicle and any form of trailer(s) that are being towed behind the vehicle. It will thus be of great benefit for hybrid/electric vehicle optimisation if the mass of the entire vehicle rig is known so that better decisions can be made regarding the energy flow in the vehicle.

3.1.2. Motivation for this project

To be able to optimise a hybrid/electric vehicle's energy system it is required for the vehicle energy management system to know the current state of the system's variables, including fuel level (in the case of a hybrid), battery SOC (state of charge) etc., it also needs to know what the current load conditions are, and what they are expected to be in the future, both near and further down the route. This optimisation can best be done if the vehicle control system has knowledge of all the energy

requirements to travel its route [19], [20]. Knowing the vehicle mass with good confidence is of great benefit as it allows a more accurate calculation of the required energy to complete the route.

3.1.3. Project requirements and beneficiaries

It is required of this study to deliver a system that can estimate a vehicle's mass with good confidence by using simple-to-implement sensors and software. It will be of great benefit if the system is capable of not only estimating the vehicle mass but also being able to account for any towed loads as well, as this is the actual mass where the energy will be spent. If proven successful this system will allow better control of hybrid/electric vehicle power systems yielding more efficient overall operation. This will not only be of financial benefit to the companies running hybrid/electric vehicles but will also be of benefit to the environment as less energy is wasted.

In this project, a torsional load cell is implemented on the driven live-axle of a road vehicle to measure the driving torque whilst in motion. An optical encoder is implemented to measure the wheel rotational displacement and a barometric pressure sensor is used to measure the barometric altitude. With barometric altitude measured as a function of distance, it is possible to obtain a good estimate for the incline of a section of road. These values are then used in a force balance equation and filtered using specific criteria to obtain an estimate for the vehicle mass that continuously updates as more data is available. Test results showed that the proposed system is capable of estimating the vehicle's gross mass accurately to within 10%, with the standard deviation for the mass estimation error being 5.2%. These mass estimations also held accurate with trailers added behind the vehicle with the only exception being that if the cross-section of the trailer is bigger than the vehicle the area used in the formula for aerodynamic resistance must be corrected. A PCT patent is filed for this system, with patent number PCT/IB2022/061330.

3.2. Literature study

The motivation for hybrid/electric vehicles was already mentioned in Chapter 3.1.2 above. Below follows some additional relevant research into the topic of hybrid/electric vehicles, when they are of advantage, their shortcomings and optimisation that was performed on them. This will motivate the need to know the mass of a hybrid/electric vehicle during use.

The real benefit of hybrid systems is exploited when the vehicle is travelling over hilly terrain, where uphill and downhill sections are frequent or when accelerating and decelerating regularly [13], [14], [15]. Flat routes' fuel consumption can be 15%-20% lower than that of hilly routes for non-hybrid vehicles showing that road grade plays a significant role in vehicle fuel consumption and emissions [12]. They also state that vehicle weight influences fuel consumption significantly.

Boretti [42] states that battery-based hybrids have a full regenerative cycle efficiency of around 36%, and proposed a mechanical energy storage system to use in regenerative braking to store vehicle kinetic energy in another form of kinetic energy elsewhere rather than storing the energy in an electric storage unit. This kinetic energy storage system is said to have an overall efficiency of 70%, almost twice that of a battery hybrid system. An estimate is made of a 30% CO₂ reduction over frequent start/stop driving cycles.

Johannesson et al. [13] propose speeding up when going downhill is preferred to regenerative braking due to less wear on the battery and hybrid system. Cruise controllers also allow a loss in speed when climbing up a hill to reduce peak power demand, it then engages neutral over the crest and recovers speed when going downhill. With the weight of heavy vehicles, the kinetic energy in movement becomes a big store of energy itself. Additional savings can be achieved by accelerating up to the maximum over-speed velocity and then coasting down to a lower speed again, running the motor in its more efficient region, though this can cause discomfort to the passengers.

Lajunen [14] reported that payload-specific-fuel-consumption goes down as the payload goes up, meaning that high loads are more efficient than smaller loads that result in more trips. Boroujeni and Frey [43] give an equation for VSP (vehicle specific power) [kW/ton], power required for acceleration, drag, friction losses and aero losses. If the mass increases the VSP decreases.

The road incline greatly influences fuel consumption and the vehicle's power requirements, a study into how incline is determined is therefore performed. In many cases, researchers made use of detailed as-built plans for a specific section of road to be travelled [21]. This gives very accurate topographic data but means that the route will have to be measured and then programmed before the hybrid system will be able to use the information to optimise energy usage. Vahidi et al. [21] used the as-built plans for the route together with the vehicle's onboard engine controller to log the engine torque, wheel speed and transmission retarder torque. They also inserted a pressure transducer in the brake system. They digitised civil maps (road plans) to find road grades. They tried using GPS, but the data was found to be too noisy for practical use. GPS was just used to determine the starting point on the map. They found that obtaining a model that accurately correlates brake pressure to stopping force is very difficult, so instead, they opted for brake pressure transducers in the brake lines, and these were used to identify when braking takes place and neglect portions of the data in which service brakes were applied. To obtain road grade information accelerometer and tilt sensors with OBD CANBUS data as well as civil drawings were used to determine the road grades. Boriboonsomsin and Barth [12] used a comprehensive modal emissions model (CMEM) to simulate the road grades, not actual topographic data. Hellström et al. [41] make use of an on-board road grade/slope database together with a GPS unit for positioning. Using this data the information about the route that lies ahead is extracted and better decisions can be made regarding the driving style, affecting fuel consumption. They determine the slope by differentiating the topographic route profile. Their mathematical model makes use of real-time dynamic programming. Zhang and Frey [44] make use of LIDAR to determine the topographic profiles of routes. They state that road grade/incline has a significant effect on vehicle specific power (VSP) and emissions. VSP is the instantaneous power per unit mass of the vehicle based on speed, acceleration and road grade and accounts for power demand and movement losses.

To qualify the idea that accurate mass estimation is of value to energy optimisation a search is done into work related to the topic. Hellström et al. [5] proposed an optimised control strategy to reduce the amount of fuel used in a heavy diesel truck. This strategy utilises information about the road that lies ahead, specifically looking at the topography of the route. They minimise the fuel mass used whilst taking the maximum route time allowed into account. They state that mass is the most important parameter in their generic analysis of a heavy vehicle.

To obtain the incline of a stretch of road most authors used either the as-built plans of a road, GPS or a barometric air pressure sensing device to determine the altitude variation. Boroujeni and Frey [43] state that barometric altitude estimation is accurate to ± 1 metre relative altitude over a 1km distance. They found that results within $\pm 0.5\%$ can be achieved by averaging the results of between 22 and 50 runs of the same segment.

Estimating vehicle mass from within the vehicle is not a new topic of study. For civilian electric vehicle transport purposes, most authors assume the vehicle mass to be constant [45], [46]. This is reasonably accurate for small city cars, but not accurate for larger commercial vehicles where the total mass varies significantly [21].

Wilhelm et al. [47] found that the mass of an electric vehicle can be estimated sufficiently accurately using Newton's second law of motion if sensor data is accurately gathered. They found that the best accuracy ($<3\%$ error) is obtained during wide-open throttle (full power). During normal driving the model accuracy deteriorates to an error of around 10%, which is not sufficient for their study as they wanted to determine the number of passengers in the vehicle based on their estimated mass. They also concluded that real-time mass estimation can significantly improve distance to empty estimation.

Vahidi et al. [21] reported that it is essential to know the vehicle mass during operation if one is to optimise its energy management over a driving cycle accurately. They report on a patented approach in which the mass is estimated by measuring the vehicle velocity loss during gear changes. This assumption can hold seeing as a gear change is usually short in duration, resulting in quite constant road loads during the calculation. Their incline estimations were obtained from road profile plans.

Fiori and Marzano [48] use the instantaneous total weight of the delivery truck to model the energy consumption of the electric vehicle. The weight is estimated by a person on board who manually subtracts the weight delivered at a delivery point from the total load in the vehicle. Meaning that the real mass is used, although it is estimated by manual methods. They estimated the topographic profile of the route using GPS data sampled at 1Hz. They recorded the vehicle's power-train data by recording the battery voltage and current obtained from the vehicle OBD port, also at a 1Hz sampling frequency. They found that up to 25% of energy could be recovered by regenerative efforts.

In the electric vehicle energy consumption model proposed by [49], they assumed a constant mass for the electric vehicle. They used acceleration, road grade, vehicle speed etc. to determine the electric vehicle's instantaneous energy consumption. Their model has an average error of 5.9%. They reported that a heating or air conditioning system greatly reduces the electric vehicle's range.

A patent [50] made use of electromagnetic load sensors that are placed between the vehicle suspension and the load to yield an estimate of the load on the vehicle. Yang et al. [51] implemented strain gauges on the leaf springs of a commercial vehicle to measure suspension loadings at each wheel and then sum the results to obtain an estimate of the vehicle mass. They made use of a neural network model to account for the nonlinearities encountered in measuring the leaf spring strains. Though the direct measurement strategies provide good mass estimates they only work for loads on the vehicle, not including something that is towed behind it.

The research aspect that this paper wants to address is thus to obtain a method of obtaining an estimate for the mass of a vehicle (and possibly adding a trailer) by making use of simple-to-implement sensors and low computational requirements. This process can be easily applied to a great variety of vehicles by simply characterising a couple of parameters and making use of data from low-cost and easily applied sensing and computing devices. The proposed system should not be sensitive to weight distribution, which will make it more robust when applied in practice.

An optical encoder will be used to measure wheel rotational displacement and thus also velocity, a patented torsional load cell will be used to obtain the driving torque to the driven wheels and a barometric pressure sensor will be used to obtain an estimate for the incline of a stretch of road. These parameters will then be used to solve the mass estimation equation based on Newton's second law of motion[7]. If a vehicle is fully electric the motor torque can be accurately estimated from the electric current to the motor, and with an accurate model of the drivetrain efficiency, the wheel torque can be estimated allowing the vehicle mass to be determined from there. This was cited in the literature presented above. The advantage of the load cell presented in this thesis is that it can work on conventional IC, hybrid and fully electric vehicles that make use of a conventional differential, removing the need for drivetrain efficiency models and allowing a wider application field than simply using electric motor current in fully electric vehicles.

3.3. Theory/calculation (theoretical modelling)

To estimate the total mass of the vehicle a free-body diagram of the vehicle, showing all external forces is constructed, see Figure 3.1. For a vehicle travelling on an incline, the force balance on the vehicle is comprised of the propulsion force, aerodynamic resistance, rolling resistance, acceleration and incline forces. A sum of forces balance is performed to solve the mass estimation problem. Chapter 3.3.1 summarizes how each of the forces is determined for the analysis.

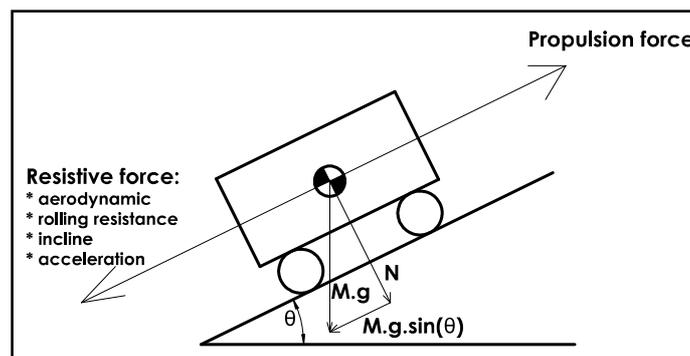


Figure 3.1: Free body diagram of a vehicle on an incline

3.3.1. Determining forces for mass estimation

Next follows a description of how each force in the free body diagram is determined.

3.3.1.1. Acceleration force

The wheel rotational displacement is logged as a function of time. The module responsible for measuring the speed of the vehicle measures the number of microseconds between two consecutive

pulses on the wheel's rotary encoder, and from this, an instantaneous velocity can be obtained. By calculating the change in velocity over a known time period one can find the acceleration that the vehicle experienced during that time period. Determining the acceleration in this way provided more stable values than what an accelerometer can give as an accelerometer is also susceptible to vehicular vibrations. This corresponds well to the conclusion by [21]. The force, F , generated by means of accelerating, a , the vehicle of mass, M , is then determined by Newton's second law, Eq.3.1:

$$F_{acceleration} = M \cdot acceleration \quad (\text{Eq.3.1})$$

3.3.1.2. Aerodynamic force

Aerodynamic forces are modelled as a function of vehicle geometry, size, air density and velocity. The general aerodynamic force equation Eq.3.2 is applied to the model with initial estimates for the vehicle used for testing purposes obtained from online sources. These initial estimates are then calibrated using the results obtained from our tests, as the vehicle has a slightly modified outside geometry from the standard one described in the auto data sheet. Air density is represented by ρ , the drag coefficient by C_d , the frontal area by A_f and velocity by V .

$$F_{aero} = 0.5 \times \rho C_d A_f V^2 \quad (\text{Eq.3.2})$$

A difficult phenomenon to account for in aerodynamic loading is wind loading, which changes the effective velocity and angle at which the vehicle passes through the air. This is assumed to be a small portion of the overall effect as the geographic location of the testing is not in an extremely windy zone, and thus the effect of wind is neglected. The average annual wind speed for the area where tests were conducted is 7kts (3.6m/s) [52]. To evaluate wind speed sensitivity a simulation was run with a constant offset in the forward velocity for the aerodynamic drag equation of 3.6m/s. During this simulation, a mass error of 10% is noted. This shows that the mass estimation algorithm is very sensitive to variance in this parameter value. Measuring the true wind speed can be done using a wind speed meter mounted on the vehicle. This was however not done for this project as mounting of the sensor would involve extensive testing to determine the correct placement to account for aerodynamic effects. The 3.6m/s offset simulation assumed the vehicle constantly experiences a wind force opposing the direction of movement, which would not be the real-world case.

3.3.1.3. Rolling resistance

The resistance to motion as a result of the friction of a tire rolling over a road can be presented in several ways. The first and most simple way is by simply multiplying the normal force on the road by a constant factor, termed the rolling resistance coefficient. This is the most elementary form of the rolling resistance equation. For vehicles travelling at a relatively low and constant velocity, this approach works well. It can be represented by the following equation:

$$F_{roll} = C_{rr} N \quad (\text{Eq.3.3})$$

Where F_{roll} is a force that always opposes the direction of motion, C_{rr} is a constant known as the rolling resistance, usually between 0.012-0.015 for vehicle tires on tar surfaces and N is the normal force of the vehicle on the road [53]. It should be noted that the experiment that was performed in this study was done using a road vehicle, so this factor should be updated accordingly for different vehicle and running surface types.

At low speeds, a more accurate version of this equation incorporates a linear velocity-dependent term.

$$C_{rr} = 0.01 + (1 + V/100) \quad (\text{Eq.3.4})$$

Where V is the velocity in mph [53]. Variations of this linear velocity-dependent rolling resistance coefficient also exist, but work in the same manner where a constant term is added to a term multiplied by the velocity, some models even account for the type of tyre construction, like radial or bias.

In [53] it is presented that over even broader speed ranges the rolling resistance factor rises in a manner that is more closely approximated by a quadratic relationship concerning velocity. The Institute of Technology in Stuttgart developed such an equation. The factors used in their equation are then also pressure-dependent, making it very hard to implement on a vehicle when the exact details of all their coefficients are not known [53].

Due to the moderate speed levels that the test vehicle will experience and the additional complexity of the non-linear equations, only the constant and linearly varying rolling resistance functions will be applied and updated if proven not to yield satisfactory results.

3.3.1.4. Incline force

In Chapter 3.2 it was reported that a barometric pressure sensor combined with an odometer was capable of accurately representing the topographic characteristic of a route travelled, with the typical relative error of the sensor being less than 1 metre, resulting in an incline error over a 100 metre distance of less than 1°. By taking the derivative of the topographic plot with respect to distance one can obtain the incline angle data for a route travelled. With the incline angle known throughout the route, it would be possible to determine the incline component of the forces imposed on the vehicle, as indicated in Figure 3.1. When the vehicle is travelling up an incline the incline forces act to reduce the velocity of the vehicle, and when travelling downhill the incline force direction switches around such that the decline force works in the same direction as the propulsion force wanting to increase the vehicle velocity. The incline force is calculated by taking the weight component parallel to the direction of travel for a vehicle located on an incline, thus Eq. 5 gives the incline force:

$$F_{incline} = M \cdot g \cdot \sin(\theta) \quad (\text{Eq.3.5})$$

Where M is the mass, g is gravitational acceleration and θ is the incline of the road, in radians.

3.3.1.5. Propulsion force

A proposal was made to measure the torque on the rear axle of the vehicle using strain gauges. The torque values obtained from the calibrated strain gauges will give repeatable and stable results with minimal noise. With the torque easily and accurately measured directly using strain gauges, one can divide the axle torque by the effective wheel radius, R_{wheel} to find the effective forward propulsion force applied by the vehicle drive train, as shown in Eq.6.

$$F_{propulsion} = \frac{T_{axle}}{R_{wheel}} \quad (\text{Eq.3.6})$$

3.3.1.6. Braking force

To obtain an accurate model of the braking force of a vehicle is very difficult to achieve, as factors including brake bias (which is load dependent for many commercial vehicles), brake fluid pressure, temperatures, velocity, the type of brakes on each wheel (front disc, rear drum), brake wear and potentially other influences all have a drastic effect on the velocity retardation. The approach followed was to neglect calculations of vehicle mass whilst the brake pedal is pressed, as this is for a small fraction of the total travelling time. This is a similar concept to what [21] performed. A pressure transducer was installed into the right rear wheel's brake line and the output voltage of the transducer was measured using an analogue pin on the Arduino microcontroller. If the voltage reached a certain threshold value (exceeded a preset pressure value) the Arduino code paused until the brake pressure is relieved. This pause/discontinuity does however give an error in the mass estimation results, as the vehicle velocity has a discontinuity which causes the program to assume the abrupt change in velocity was due to excessive mass, usually a ridiculously high value, to the order of millions of kilograms. These intermittent excessive values had to be discarded by the mass estimation algorithm.

3.3.2. Sum of forces

Now that all the forces acting on the system are characterised it would be possible to implement a force balance equation and solve for the only unknown, the vehicle mass. This is performed as follows:

$$F_{propulsion} - F_{roll} - F_{aero} - F_{incline} = F_{acceleration} \quad (\text{Eq.3.7})$$

$$\frac{T_{axle}}{R_{wheel}} - C_{rr} \cdot g \cdot M - \frac{1}{2} \rho C_d A_f V^2 - M \cdot g \cdot \sin(\theta) \quad (\text{Eq.3.8})$$

$$= M(a)$$

$$M = \frac{\left(\frac{T_{axle}}{R_{wheel}} - \frac{1}{2} \rho C_d A_f V^2 \right)}{(a + C_{rr} \cdot g + g \cdot \sin(\theta))} \quad (\text{Eq.3.9})$$

Where:

T_{axle} = Torque on axle measured by strain gauges, Nm

R_{wheel} Wheel radius = 0.39m (235/85R16 tyre)

ρ = Air density determined from barometric pressure, $\sim 0.98 \text{ kg/m}^3$

C_d = Drag coefficient, 0.4 (from vehicle datasheet)

A_f = Frontal area of vehicle, 2.33 m^2 (from vehicle datasheet)

V = Vehicle velocity from odometer sensor, m/s

a = Acceleration, determined from the change in velocity, m/s^2

C_{rr} = Rolling resistance factor, speed-dependent (more accurate, option applied)

typically between 0.012 and 0.015 for radial rubber tyres on tarmac

= $0.013(1+V/42)$ {found by experimentation to yield best fit to data}

θ = Incline angle, radians

This equation does however pose two risks, namely:

If the propulsion force equals the aerodynamic drag force the numerator in Eq.3.9 approaches zero, yielding zero mass calculated. Also, if the sum of the acceleration, the rolling resistance term and the incline term approaches zero, the denominator causes the mass value to approach infinity (\pm). Eq.3.9 still holds during these cases, but the inaccuracy of the sensor data causes the massive error. This needs to be carefully accounted for during the mass estimation calculations.

3.3.3. Sensors implemented and vehicle used to facilitate mass estimation

The test vehicle used for the project was an Isuzu Frontier. This vehicle was instrumented with sensors to facilitate mass estimation. The sensors and how they are applied are explained in the following chapters. Figure 3.2 shows the test vehicle on a weighing bridge. The vehicle had a mass of 2090kg with a full fuel tank, no payload and no driver. The accuracy of the weighbridge was stated as within 20kg, which translates to 1% of the vehicle mass. A laptop was used together with two Arduino prototyping boards to facilitate the data recording of the sensor output values.



Figure 3.2: Test vehicle on weighing bridge, Mass = 2090kg, full fuel tank, no payload, no driver. Accuracy within 20kg (1%)

3.3.3.1. Odometer sensor

An optical encoder was fitted to the right rear wheel of the test vehicle so that the rotational displacement of the wheel could be accurately determined. From the rotational displacement, the linear travel distance can be obtained by scaling the rotational displacement by the wheel radius. The main hardware for the optical encoder consists of an infrared LED and an infrared photodiode which oppose each other. A ridge is added on the outside of the brake drum which is used to obstruct the infrared light beam, causing a change in voltage on the optical encoder and allowing the rotational displacement to be counted in discrete increments by the Arduino board. An important point to remember is that wheel spinning will induce an error in the odometer reading if it is installed on one of the driving wheels, so care should be taken in the initial testing to minimise wheel spin, or to place the encoder on a non-driving wheel. Using a wheel to measure the displacement instead of the propeller shaft does induce small errors when the vehicle is travelling around corners, but this effect was neglected for the purpose of this study as the radii of roads are generally quite large compared to the width of the vehicle.

A GPS coordinate logger was used as a reference to compare with the optical trigger odometer values. From the test performed the GPS data and the calibrated odometer data match extremely well. By the end of the 20 km route, there is an odometer difference of 60 metres as compared to the GPS and online map information, yielding an odometer error of only 0.3%. This verifies the usability of the proposed device.

3.3.3.2. Barometric pressure sensor for incline estimation

A method of obtaining the incline characteristics of the driven road is required. One method to obtain altitude is to use a barometric pressure sensor, and from the air pressure one can determine an estimate of the altitude, from which the incline data can be derived if used together with an odometer. In an attempt to save cost an inexpensive open-source sensor board was implemented, the GY87 consists of a BMP085 barometric pressure sensor, an MPU6050 3-axis accelerometer, a 3-axis gyro and a HMC5883L 3-axis magnetometer. These module boards retail for around ±150 Rand (±10USD).

Next, it will be researched if the altitude accuracy of the barometric sensor is as good as was stated by [43] to allow its use in determining the incline for the mass estimation strategy. From the datasheet the "ultra-high resolution" mode is capable of measuring a 0.25 metre altitude change, with the RMS noise of the signal being able to go down to 0.1 metre, the resolution of the output data is 0.01 hPa (<0.1 metre at sea level) [34]. An initial test was performed by applying a change in altitude and determining the barometric altitude from the change in measured pressure. The pressure reading is converted to an effective altitude using the International barometric formula, Eq. 10. Where P is the measured pressure and P₀ is the pressure at sea level. In this equation, a pressure change of 1 hPa equates to a change in altitude of 8.43 metres at sea level.

$$Altitude = 44330 * \left(1 - \left(\frac{P}{P_0} \right)^{\frac{1}{5.255}} \right) \quad [34] \quad (Eq.3.10)$$

For the gathering of initial test data, the sensor was moved up and down in a building using the elevator, starting on floor 9, riding up to floor 15, down to floor 3 and then back up to floor 9. The raw altitude data was downloaded and can be seen in Figure 3.3. The red line represents the raw data, and the black line is a simple software low-pass filter used to smooth out the higher-frequency noise. The raw data has a noise range of around 2.5 metres, with the filtered data having a noise range of 0.4 metres. Tests showed that an altitude drift due to weather can however cause a change in barometric altitude of up to 0.5 metres in one minute. This does deteriorate the accuracy of the incline estimation. Even with the drift of 0.5 metres and noise amplitude of 0.2 metres, it will still yield incline estimations within 0.4° over a 100 metre distance. It can thus be assumed that the barometric altitude is usable as a means to determine the incline of a section of road. It should be noted that this sensor's minimum detectable height does place limit the shortest distance increment for which it can be used effectively.

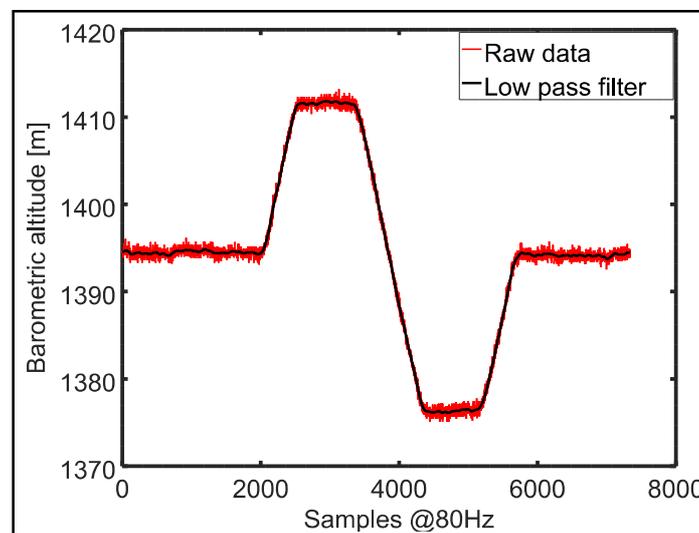


Figure 3.3: Estimated barometric altitude by riding the elevator from floor 9 to 15 to 3 to 9

3.3.3.3. Axle torque sensing

Strain gauges were applied to the axle of a live axle motor vehicle for driving torque measurement. The eight strain gauges were placed four at the front of the axle and four at the back of the axle, on the centreline of the axle. Placing them on the centreline means that they will not see a bending moment due to the vehicle's weight. Placing them in front and at the back of the axle allows to wire them such that the bending moment in the forward and backward direction is cancelled out as well. Because the vehicle uses a live axle that is supported on the suspension at the ends of the axle using leaf springs it means that the torque transfer from the differential to the wheel will be measured by these strain gauges as a reaction force of the centrepiece of the differential will torque against the suspension. See Figures 3.4 and 3.5 for reference. The strain from these 8 strain gauges is measured with a 24-bit HX711 analogue to digital converter. The torsional load cell was given an initial calibration by applying a torque to the differential with a handheld torque wrench. The torque was applied through knuckle joints so that no vertical forces/torques were applied to the tube of the axle by the torque wrench during

calibration. Figure 3.6 shows the axle torque result obtained when accelerating from a standstill from first gear to fourth gear. The driving torque measured on the rear axle yields stable and sensible results and can now be used in the mass estimation strategy [7].

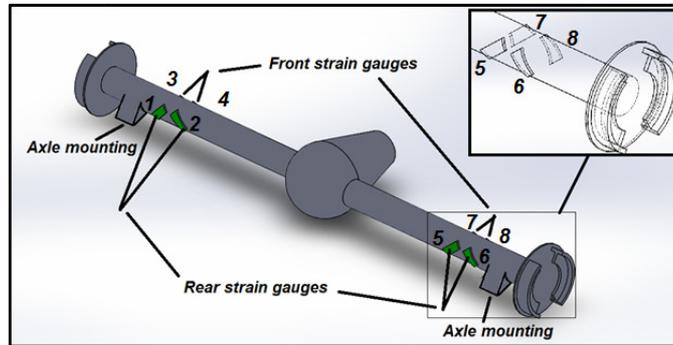


Figure 3.4: Live axle with 8 strain gauges mounted (numbering is consistent with figure 3.5b)

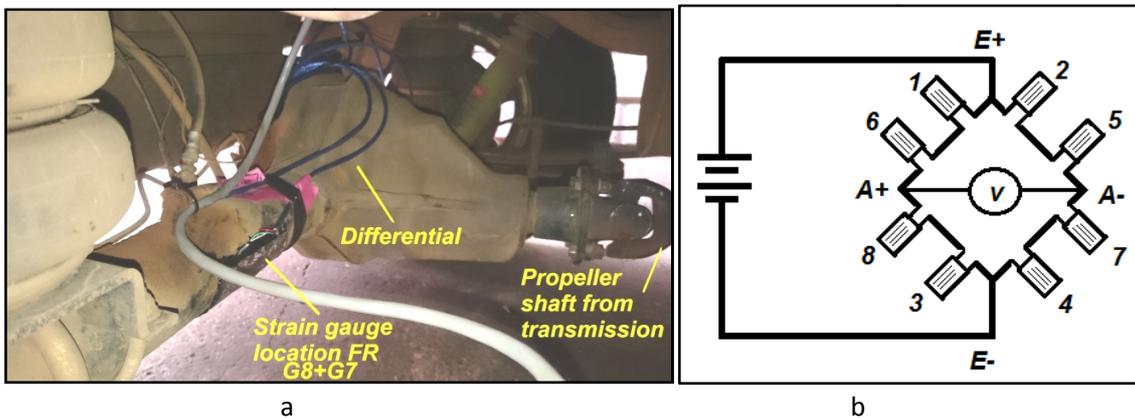


Figure 3.5: Strain gauges mounted to the rear axle. a) front right mounting point, b) Wheatstone bridge wiring

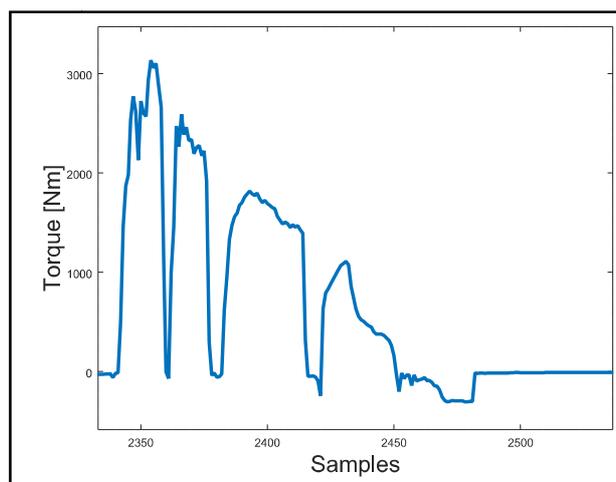


Figure 3.6: Axle torque whilst accelerating in gears 1 through 4

3.3.3.4. Brake system fluid pressure sensor

To facilitate pausing the program during brake events the brake pressure was measured using an analogue pressure transducer. The signal was processed using a 16-bit ADS1115 analogue to digital converter. It will now be possible to know when the brakes are applied if the pressure in the brake system is found to exceed some minimum threshold value, and the mass estimation program can then be paused until the pressure drops below a threshold value. The pressure sensor used is an analogue sensor that outputs a voltage from 0.5 v to 4.5 v for a pressure of 0 MPa to 6.9 MPa (1,000 PSI) respectively.

3.3.3.5. Schematic layout of sensing system

Figure 3.7 shows a schematic layout of all the sensing and recording components used for mass estimation. The Arduino Mega board is the master board that streams the data to the laptop for processing. The Arduino Uno board counts the optical encoder pulses to function as an odometer. The HX711 board is the 24-bit ADC for the axle torsional load cell. The ADS1115 is the ADC for the brake pressure. The GY87 contains the barometric pressure sensor for altitude estimation and the GPS is used as a reference to calibrate the optical trigger and verify the barometric data.

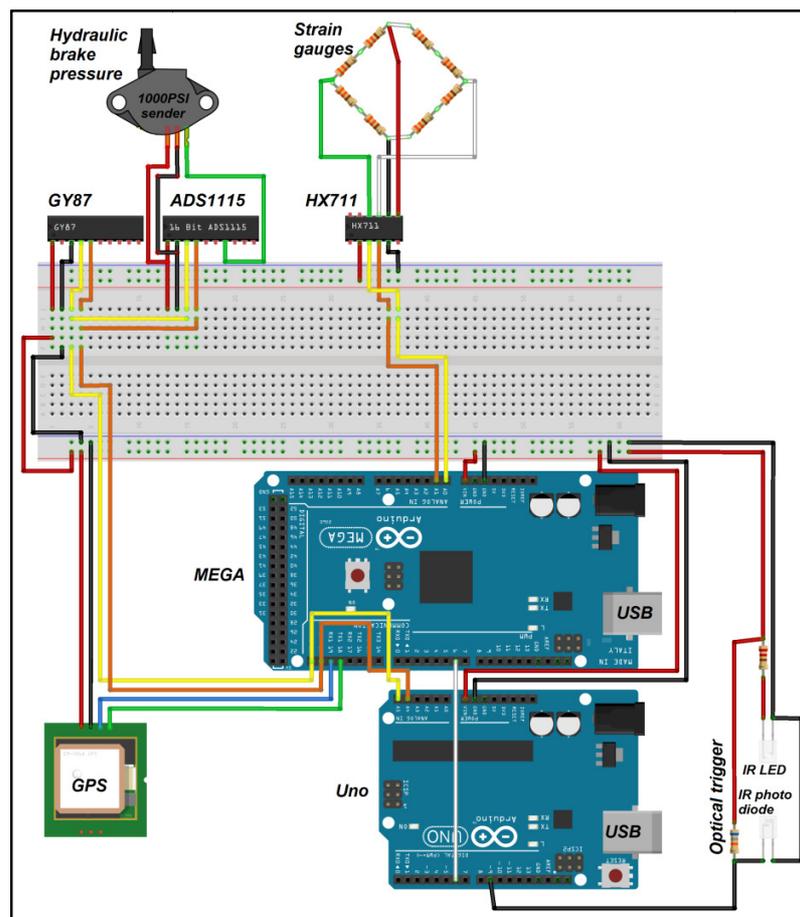


Figure 3.7: Vehicle mass estimation sensing circuit diagram

3.3.3. Octave program structure schematic

A schematic diagram showing the calculation sequence for determining the vehicle mass is shown in Figure 3.8. An advantage of this method is that it can be run post-test or updated in real-time, as the mean value will be updated continuously if run in real-time. The smoothing filter that is applied reduces the errors in mass estimation by reducing higher-frequency variations in the sampled data.

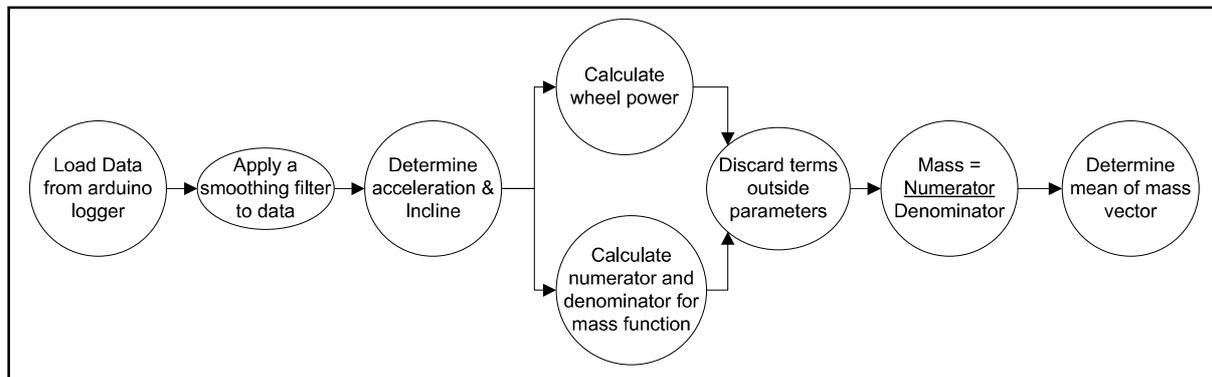


Figure 3.8: Octave program schematic for mass estimation

3.3.4. Experimental setup for mass estimation

The approach followed in this work is a novel way of determining vehicle weight using a torsional load cell mounted on the differential. The advantage of this approach is that it is not sensitive to weight distribution on the vehicle and may even work to estimate the combined mass of the vehicle and a trailer.

The parameters listed in Chapter 3.3.3 were recorded at a frequency of 5Hz, which was deemed high enough to pick up all the vehicle's dynamic movements but low enough not to generate excessive data. Vahidi et al. [21] made vehicle mass estimations based on the assumption that the forces on a vehicle are constant through the time a gear change takes place, which is longer than 0.2 seconds (5 Hz) as was used in this study.

For all the tests performed the fuel mass was estimated by taking the average fuel consumption for the vehicle and the trip odometer reading to estimate the fuel amount used and subtracting the mass of that fuel from the vehicle's gross mass.

Several driving tests were conducted with varying payloads placed in the vehicle. Tests were also performed with the vehicle towing a small trailer laden with 75 x 5 litre (5 kg) water bottles. The small trailer had an empty mass of 230 kg determined by using a scale. The water bottles were each weighed and filled to within 1% of 5 kg using a digital scale. The vehicle and small trailer are shown in Figure 3.9. A larger trailer was also used to test the system for larger gross combined masses. The larger trailer was towed in both empty and laden form. Figure 3.10a shows the vehicle with the larger trailer, as well as the inside laden with 90 x 5 kg water bottles, Figure 3.10b. It should be noted however that the larger trailer increases the frontal area, which affects the aerodynamic forces that need to be overcome by the propulsion force, the effect of this will be assessed in the results section.



Figure 3.9: Test vehicle with a small trailer, GCM 2755kg (75x5kg+230kg+2150kg [driver and full tank]). $A_f \approx 2.33\text{m}^2$



Figure 3.10: Test vehicle with a large trailer, GCM 3270kg. $A_f \approx 3.3\text{m}^2$

3.3.5. Testing Procedure

To determine the mass using Eq.3.9 all sensors and constants must be extremely well calibrated and defined as slight variations in the force balance between terms in the numerator and denominator can cause severe variance in the estimated mass of the vehicle. Fine calibration of the sensor's data and constants, like drag coefficient and rolling resistance, is however quite difficult to perform theoretically or by simulation, so these parameters were calibrated based on test results. An iterative method was used to perform fine calibration of the sensors such that the mass equation yields the known masses for the vehicle rig in different loading conditions and for different routes travelled. This is discussed in more detail in Chapter 3.4.

To gather data the vehicle was fitted with the Arduino sensor board to stream the various parameters to the laptop. Normal driving routes on public roads were driven as a normal driver would do, so that realistic trip data may be obtained. Wilhelm et al. [47] mainly made use of maximum power scenarios, which limit the real-world applicability of their system. The driving cycles for this paper comprised urban and highway routes so that the accuracy of the system in both of these scenarios can be evaluated. The

data stream was paused when the vehicle brakes were applied, due to the difficulties mentioned by Vahidi et al. [21], thus simply neglecting when the brakes influence the vehicle's movement. The results of the tests are discussed in the Chapter 3.4.

3.4. Results and Discussion

Here follows the results obtained from testing of the vehicle mass estimation strategy. Figure 3.11a shows the raw mass estimation derived from Eq.3.9 for a route of ± 20 km travelled. This is the data before imposing any limits or cutting/discarding of data.

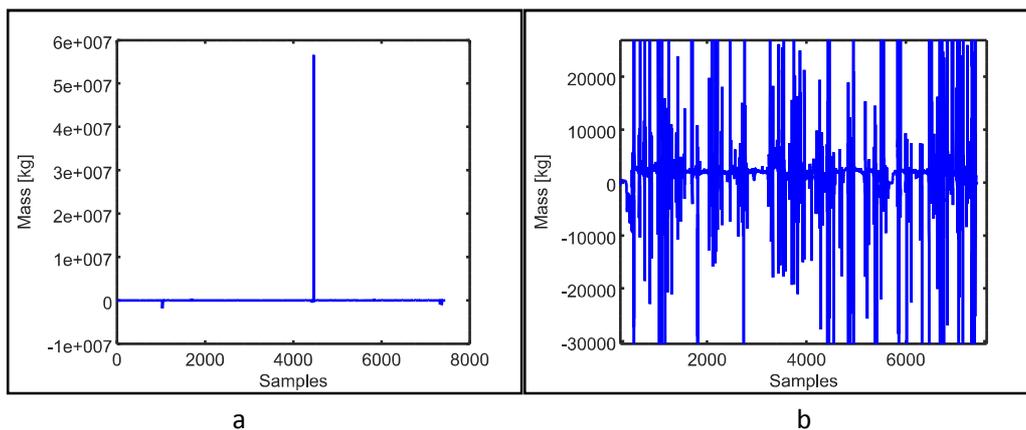


Figure 3.11: Raw mass estimation, sampled at 5Hz. a) The total mass plot, b) zoomed in

From the data shown in Figure 3.11, it can be seen that the mass results obtained by implementing Eq. 9 yield non-stable and completely unrealistic results. The vehicle had an unloaded mass of 2090 kg with a full tank of fuel and no driver so this is the typical order of magnitude where the mass estimations are expected. The mass estimation equation is very sensitive when the forces are very small, and any noise/variation in any of the parameters can potentially yield these unrealistic results. To avoid such unrealistic mass estimates a strategy needs to be implemented to filter out the unrealistic data or avoid using the scenarios where such data are expected to occur. To find the sources of unrealistic data the mass was plotted as a function of various measured parameters. Figure 3.12 shows the calculated mass plotted as a function of a) wheel power, b) velocity and c) acceleration. From these figures regions may be identified where the most stable data is found, and the unrealistic areas can then be avoided/discarded. It is noted that the calculated mass has large errors at low power levels, and seems less noisy at absolute power levels greater than 20 kW, Figure 3.12a. Power, P , is calculated by taking the product of propulsion force and velocity $P = F.V$. There doesn't seem to be any assumptions that can be made for the data as a function of velocity, Figure 3.12b. Figure 3.12c does show that higher accelerations yield more stable results, except for cases where the acceleration was calculated in the discontinuous sections experienced when pausing the data stream during braking and calculating the acceleration based on discontinuous velocity data, so data points for unachievable acceleration are also neglected (acceleration $> 3\text{m/s}^2$ discarded).

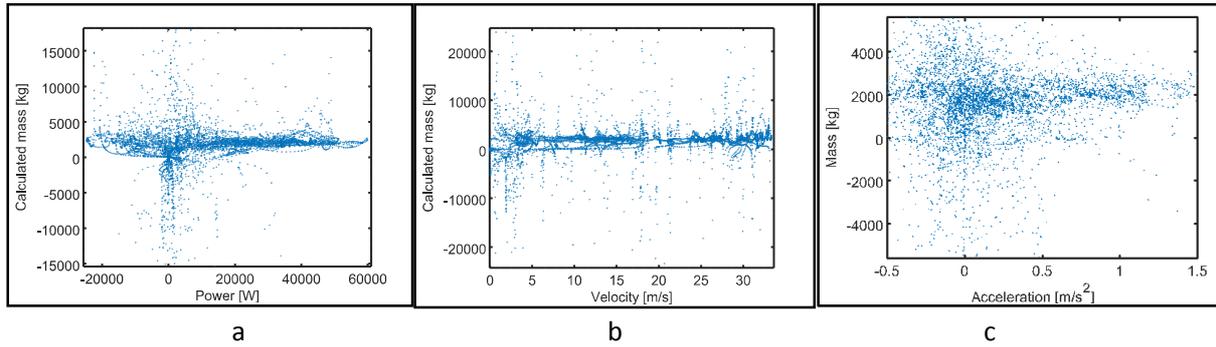


Figure 3.12: Test data for calculated mass versus: a) Power, b) Velocity, c) Acceleration

The first strategy implemented was to remove data that did not meet a minimum amount of absolute wheel power, thus discarding the very low velocity conditions as well as coasting, where the terms in the mass equation become small and any variation in any parameter causes vast changes and instability, and thus errors in the mass calculated. This will also remove data when the vehicle is stationary and most of the terms in Eq.3.9 fall away.

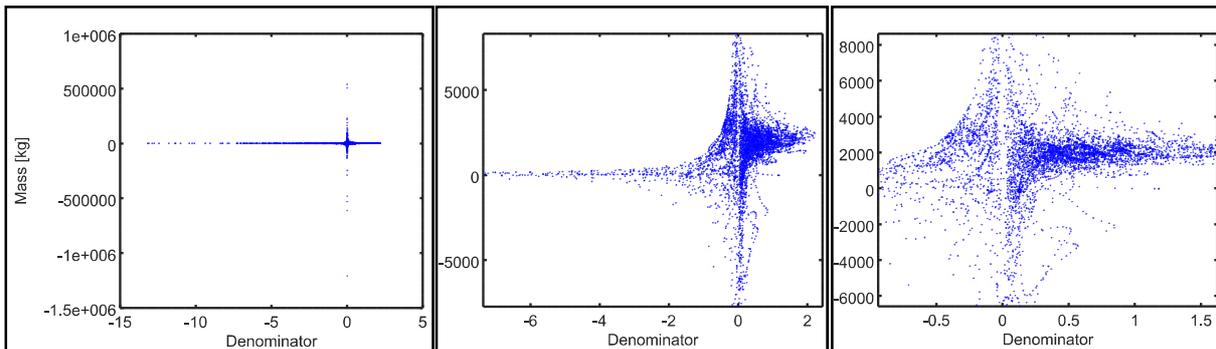


Figure 3.13: Estimated Mass for denominator value

From the data shown in Figure 3.13, it may be noted that as the denominator in Eq.3.9 approaches zero the calculated mass shoots out of reasonable bounds. It can also be seen that a negative denominator yields mass results close to zero, which is also not applicable. The second strategy to remove unrealistic results was to limit the minimum size of the denominator then, and impose a maximum on the absolute value of the estimated mass that may be considered for further steps. A suitable denominator minimum was determined through simulations using actual test data to be 0.2. This was set at a point where the final result was not very sensitive to changes in this value. Making it too small yields mass estimates that shoot up too high, and making it too large discards a lot of usable data. The test vehicle had a nominal mass of 2090 kg, so the setting of the mass cut-off needs to be significantly higher than the true mass so that the cut won't end up discarding usable data that a low-pass filter strategy would be able to use sensibly. The mass cut-off value for the absolute value of the calculated mass was set at 10,000 kg. This value was small enough to filter out the extreme data, but large enough to not cut off usable data that can be filtered in another step. It was also found that small changes in this mass cut-off value did not influence the estimated mass significantly. This is important to make sure that the value did not yield a satisfactory result by accident. Negative mass estimates were not directly discarded as it was noted that

strong lows would often have counteracting strong highs to the other end, yielding on average an acceptable result. This was done to not overly limit the results and overwhelming the freedom of calculation by too strict pre-conceived bounds. It doesn't help to limit the values so much that the answer is almost pre-generated and fixed, independent of the data gathered, rather than determined by well-considered measurements, especially with data that has such a wide range of results. Reference [47] excluded more than 75% of their data, mostly looking at high to maximum power delivery. From the simulations performed it makes sense why so much data was discarded. Using only maximum power data is not ideal though, as a vehicle doesn't necessarily see full power during every run and maximum power would potentially negatively impact vehicle efficiency, which is against the whole goal of an optimisation project. With the proposed parameters to cut the unrealistic mass values the strategy proposed in this study also excludes close to 75% of the data points. Figure 3.14 shows the selected mass estimate data based on the above threshold parameters, in blue, with the continuously updating mean of the mass values shown in red. It can be noted that, when there are fewer data points available the mass value can vary quickly as the selected mass values are streamed in, but as the number of available data points goes up the mean value for the mass stabilises at the desired true mass value. The answer for this test run was spot on to the estimated real mass of the vehicle. The tests contained a combination of highway and urban driving, with the middle section of Figure 3.14 being the highway part from around point 500 until point 1200. It is observed that highway driving sometimes yields slightly lower mass averages than urban driving.

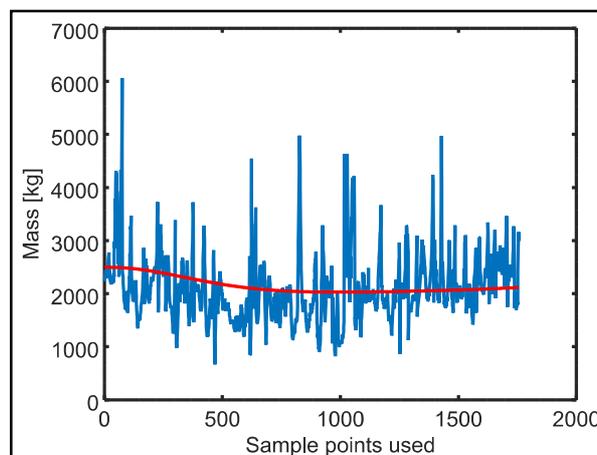


Figure 3.14: Mass estimated with parameter limits imposed. Blue = calculated real-time mass cut data, Red = LPF

Now that data is generated for the mass, it is necessary to calibrate and characterize all the parameters to ensure the reported results are usable and accurate. These include torque cell and drag coefficient calibration, coefficient of rolling resistance characterization, and finding the best values for the boundary parameters; power cut, mass cut and minimum allowable denominator. Over 2000 simulations were performed to iteratively find the values for these parameters that best fit the true data. These values are summarized in Table 3.1. In application, this characterization process will only need to be done once for a vehicle type, and calibration of the torsional load cell will be done once for each load cell fitted to a vehicle.

Parameter	Value	Unit
Torque calibration factor	1.04	-
Wheel radius (235/85R16)	0.39	Metre
Drag coefficient	0.41	-
Frontal area	2.33 (vehicle, vehicle with small trailer) 3.3 (vehicle with large trailer)	m ²
Rolling resistance	0.0129 (1 + Velocity/36)	-
Mass threshold	10000	kg
Minimum power	20000	kW
Denominator minimum	0.2	-

Table 3.1: Mass estimation parameter calibration values

Table 3.2 summarises the test results for several road tests (27) showing the actual mass of the vehicle during the test, the mass estimated by the mass estimation strategy for that specific test and the error in the estimated value compared to the correct value. Test 26 and 27 have their simulation results presented twice, part "a" shows the estimated mass when the frontal area is left as the vehicle frontal area, as with the other simulations, where part "b" accounts for the increased frontal area of the large trailer, as was shown in Figure 3.10a. Green highlight indicates <5% absolute error, yellow 5% <absolute error<10%.

Test #	Payload	True mass	Estimated mass	Error %
1	Driver + Fuel	2150	2049.9	4.66
2	Driver + Fuel	2140	2326.9	-8.73
3	Driver + Fuel	2120	1974.7	6.85
4	Driver + Fuel	2115	2115.9	-0.04
5	Driver + Fuel	2150	1949.2	9.34
6	Driver + Fuel	2140	2178.5	-1.80
7	Driver + Fuel	2150	2239.7	-4.17
8	Driver + Fuel	2140	2221	-3.79
9	Driver + Fuel	2140	2242.3	-4.78
10	Driver + Fuel	2135	2204.8	-3.27
11	Driver + Fuel	2125	1954.3	8.03
12	Driver + Fuel	2120	2095.6	1.15
13	Driver + Fuel	2115	2152.6	-1.78
14	Driver + Fuel	2112	2115.5	-0.17
15	Driver + Fuel	2130	2053.5	3.59
16	Driver + Fuel	2135	1982.2	7.16
17	Driver + Fuel	2145	1997.5	6.88
18	Driver + Fuel + 330kg water	2430	2534.7	-4.31
19	Driver + Fuel + 330kg water	2430	2291	5.72
20	Driver + Fuel + 330kg water	2425	2658.3	-9.62
21	Driver + Fuel + 330kg water	2420	2265.7	6.38
22	Driver + Fuel + 330kg water	2490	2365.8	4.99

23	Driver + Fuel + small trailer + 375kg water	2745	2869.6	-4.54
24	Driver + Fuel + small trailer + 375kg water	2735	2758.3	-0.85
25	Driver + Fuel + small trailer + 375kg water	2730	2782	-1.90
26a	Driver + Fuel + large trailer	2825	2992.6	-5.93
27a	Driver + Fuel + large trailer + 445kg water	3270	3375.4	-3.22
26b	Driver + Fuel + large trailer (Improved frontal area)	2825	2840.1	-0.53
27b	Driver + Fuel + large trailer + 445kg water (Improved frontal area)	3270	3246	0.73

Table 3.2: Mass estimation results

A variety of loading cases were evaluated, as shown in Table 3.2. This includes an inside-vehicle payload variance of 330kg ($\approx 15\%$ change) from fully laden to empty. A small or large trailer with various loads was also attached to the vehicle in tests 23 to 27, equating to an up to $\approx 52\%$ Gross Combined Mass (GCM) variance. The standard deviation for the error in the data presented in Table 3.2 is 5.2%.

3.5. Conclusion

The proposed model uses a simple force balance equation using Newton's second law of motion and easily implemented sensors to obtain an estimate of a vehicle's total (gross combined) mass. This mass value is necessary to facilitate proper optimising of the energy usage of a hybrid vehicle during operation. Still, it can also benefit legislation compliance monitoring on commercial vehicles.

Simple sensing devices were implemented to gather information on vehicle forces and the route's incline. With these parameters known a program can be executed that estimates the vehicle's mass. Due to the difficulty in obtaining an accurate braking force value from the simple sensing data it was decided to exclude data when the brakes are applied.

The initial estimates for the vehicle mass yielded unrealistic values, which led to a data filtering strategy that only used data when it fell within parameters determined to yield the most likely accurate mass estimates. This involved setting a minimum amount of wheel power, an upper limit to the mass usable and minimum denominator size in the mass estimation equation. With the mass estimates now in reasonable bounds simulations were run to obtain accurate calibration for the torque load cell, the drag force and an empirical equation for the rolling resistance as a function of vehicle velocity. It was found that the mass estimation may be sensitive to wind, which can be accounted for more accurately by implementing a wind speed sensor on the vehicle.

Rather noisy real-time mass estimates were still found. Still, it was found that averaging the mass estimations yielded an estimate for the vehicle mass that was accurate to within 5.2% on average of the actual mass of the vehicle/ and trailer. In general, estimation of the mass yielded accurate results when more than 500 usable data points were evaluated. This is equivalent to around 5 minutes of driving.

The ability to accurately estimate the gross combined mass of a hybrid/electric vehicle will allow the vehicle control system to make better decisions and energy usage estimations, thus further improving the vehicle's overall energy consumption, reducing cost and emissions.

It is important to note that this proposed mass estimation strategy is very sensitive to the surface on which is being travelled. It was proven successful on hard tarmac surfaces, with good control over tire pressures. Still, this model will lose applicability once soft surfaces with varying rolling resistances are travelled over. Rolling resistance coefficients on tarmac vary from 0.012 to 0.015, a reasonably narrow band. In contrast, the factor for rolling resistance on softer surfaces, like gravel and sand, varies from 0.02 up to 0.3, which will yield completely incorrect mass estimates as the instantaneous coefficient of rolling resistance cannot be determined during driving [53]. The vehicle speeds reached during this test were moderate enough that only applying a linearly varying rolling resistance equation proved successful and implementing a quadratic equation was not deemed necessary.

3.6. Future work

By evaluating the data certain situations proved more suitable than others at estimating the gross vehicle mass. For future studies it should be considered to build a statistical model that estimates how trustworthy the data is at a point and then apply a weighting factor to the answer provided in that situation, to improve the rate at which an estimate for the mass is obtained and also improve the accuracy of the value itself.

With a method that proved successful in estimating the mass of the entire vehicle, including trailers, if added, it facilitates the development of a robust velocity profile optimisation method, which is covered in Chapter 4.

4. Velocity profile optimisation

4.1. Introduction

Improving operational efficiencies of machines may result in increased profit, reduced wastage of resources and a smaller carbon footprint. The latter has become essential as we need to conserve the limited resources available to yield a more sustainable future. This study aims to develop a robust method of finding an optimal velocity profile at which a vehicle can travel a known route such that the overall energy consumed to traverse the route is minimised. This study forms part of a larger study in which the total cost of ownership of an underground electric locomotive is minimised. In underground mining electric locomotives are often used to haul the ore to the shaft, ready for elevating up to the surface. These electric locomotives travel known routes until the shift changes, or the battery requires recharging.

For this study, the work is all performed above-ground using an internal combustion (IC) vehicle travelling known routes. A system has been developed as part of the larger study that can identify a route being travelled in real-time without the use of external information sources like GPS, as GPS does not work underground [6]. With an underground mine being a very harsh environment the practical use of beacons and similar devices for locating purposes is not feasible. A method for estimating the vehicle's mass while it is driving, in real-time, was also developed as a part of the larger study, as the true real-time updated mass is required to be able to perform meaningful optimisation and route energy estimation [7]. The optimisation focus of this study is concerned with finding an optimal velocity profile for a known route, which means that the vehicle energy source is not of concern at the moment, and modelling thereof falls outside the scope of this study, allowing the use of an IC vehicle instead of electric. A first-principles model of the system is developed that can accurately model the vehicle dynamics and energy requirements as the vehicle completes its known route. A first-principles model is beneficial as it considers most of the system variables and is less reliant on assumptions. Developing an accurate model that is easy to solve is however not always possible. The vehicle model is constructed in Matlab & Simulink [40]. The reason for using a model in Simulink is that the built-in battery and motor models may be used for future simulations in the larger study instead of having to spend a significant amount of work developing those from scratch.

The vehicle is instrumented with sensors to enable logging of the instantaneous driving power, velocity and route topography in real-time. With this data available comparative simulations and optimisation can be performed to compare new route velocity profiles to when the vehicle was driven manually. This can then be used in future applications to inform the vehicle controller of better choices in velocity and power usage profiles when completing its route.

The quality of an optimisation solution depends on the model quality, data and optimisation strategy used, as is shown in Figure 4.1. This study focuses on developing a robust strategy to obtain an initial guess that facilitates finding consistent optimisation results. For the solver, a standard *fmincon* solver was implemented to find the local minimum of the cost function, and it was noted that the initial guess for the solver plays an important role in finding a global minimum of the cost function. Initially, sensible

first-order guesses were implemented, which proved to have limited success. A study was done to find better ways of defining the initial guess so that a more robust and faster-running optimisation strategy could be developed. The optimal velocity profile was found to have a shape that closely resembles an inverted version of the topographic profile of the route being travelled. An optimal initial guess strategy was developed and proven successful in improving the robustness of the optimiser to find the optimal solution.

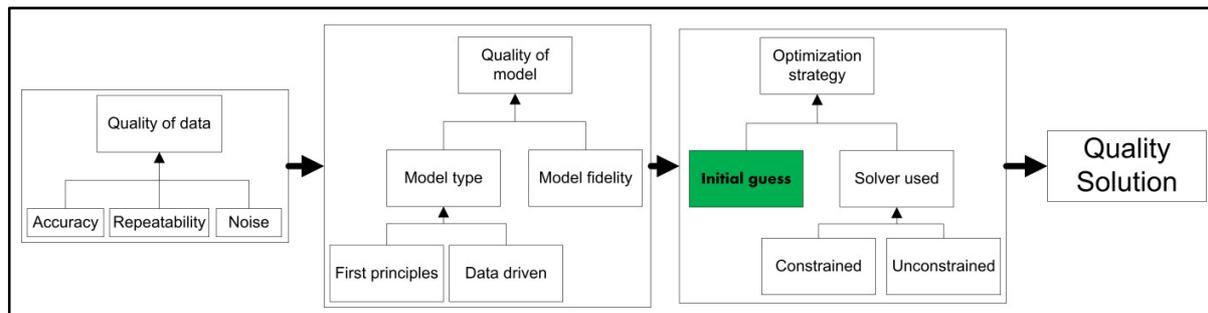


Figure 4.1: Quality of solution in optimisation

The initial guess method developed in this study also proved successful in finding improved velocity profiles in a significantly reduced time due to the low fidelity of the strategy only requiring optimisation of two variables to deliver a velocity profile for the entire route, based on the route's topographic profile. This approach typically runs for less than one minute, whereas the full optimisation runs for around 3 hours for a model fidelity of 50.

A study into the effect of regeneration was performed and noted that the ability to perform regeneration can reduce the overall energy consumption of a vehicle significantly, although the effect of an improved velocity profile proved to have a larger effect on route energy consumption. A study into the effect of mass was performed and it showed that mass only plays a role on the energy consumed but does not influence the optimal velocity profile to complete a route.

4.2. Relevant literature

To optimise vehicular energy consumption, it is important first to understand and characterise all the role players in the energy consumption. Factors such as the type of vehicle, the mass, the route and velocity profiles influence the energy required to complete a route in the given logistical constraints [31].

4.2.1. Vehicle type

The first variable is the vehicle type/size, which directly influences the aerodynamic and rolling resistance, maximum acceleration and braking, and cornering characteristics [53]. The aerodynamic drag force is expressed by Eq.4.1. The term C_d is the drag coefficient, an experimentally determined factor relating to the aerodynamic smoothness of the vehicle. The term A is the vehicle's frontal area, ρ is the air density and V is the vehicle velocity. The number of axles (driven and free) also influence the vehicle's energy consumption.

$$F_{aero} = \frac{1}{2} C_d A \rho V^2 \quad (\text{Eq.4.1})$$

The vehicle type relates to the expected minimum and maximum values for the vehicle mass. The vehicle's mass plays a role in the rolling resistance equation but will be covered separately in the following sections.

4.2.2. Vehicle total mass

The mass of the vehicle, and trailer, if equipped, contributes significantly to the energy required to complete a route. The mass is somewhat related to the vehicle type but also accounts for the payload, which can contribute significantly to the gross mass of the vehicle. Vahidi et al. [21] reported that the mass of a commercial vehicle can vary by as much as 400% during operation. Both kinetic and potential energy equations are linearly proportional to the object's mass. Basso et al. [31] generated a model where the energy is linearly proportional to the mass. Many researchers use a constant mass for the vehicle, which, in most cases is acceptable, but for commercial vehicles where the overall mass can vary significantly this assumption cannot be applied [54]. Boriboonsomsin and Barth [12] state that weight affects fuel consumption significantly.

4.2.3. Route

The route's surface influences the rolling resistance and maximum tractive abilities of the vehicle on that route. Rolling resistance is a factor that relates the amount of force needed to move a vehicle over a specific terrain as a function of the vehicle's weight. Harder surfaces generally require less effort to move a vehicle over, while softer surfaces like sand can greatly increase the force required. In its simplest form, the rolling resistance force can be approximated by a coefficient multiplied by the vehicle's weight, with C_{rr} for a typical car or truck on tar surfaces having a range of 0.012-0.015. A more accurate version of the rolling resistance equation for low velocities includes a linear velocity-dependent term, as is shown in Eq.4.2 [53]. The velocity-dependent term for high velocities can take on a quadratic form, though this was found unnecessary for this study.

$$F_{roll-basic} = C_{rr}N \quad F_{roll} = \left[Constant1 \left(1 + \frac{Velocity}{Constant2} \right) \right] N \quad (\text{Eq.4.2})$$

The friction coefficient is a measure of the limit of tractive effort on a specific surface. This relates to the maximum force that can be transferred from the wheel to the ground to accelerate/ decelerate the vehicle. The typical values for rolling resistance and friction coefficients/coefficients of adhesion (for railway applications) are stated in Table 4.1. Considering that these vary significantly, it is important to correctly account for this in analytical models.

Vehicle type	Typical coefficient of Rolling resistance	Typical coefficient of friction (adhesion)
Car [53]	0.015 (concrete) - 0.3 (sand)	0.8 (dry) - 0.3 (wet)
Heavy truck [53]	0.012 (concrete) - 0.25 (sand)	
Railway [55]	0.002 - 0.0125 (for bad quality wheels and rails)	0.25 (dry) - 0.1 (wet)

Table 4.1: Typical rolling resistance and friction coefficients

Route topography plays significantly affects the amount of energy required to complete a route. Boriboonsomsin and Barth [12] reported that fuel consumption in a normal IC vehicle can decrease by 15-20% if a route is flat rather than hilly [29]. If a vehicle is fitted with a powertrain capable of storing energy it can reduce the effective fuel consumption significantly [13], [14], [15], [30]. Liu et al. [30] report that topography has a significant influence on vehicle speed, acceleration, emissions, fuel consumption, driver behaviour, road capacity and congestion.

Hellström et al. [41] use a GPS and database of road slopes for their incline/topography estimates. This look-ahead information is used in an optimisation of the velocity trajectory concerning a criterion formulation that weighs trip time and fuel consumption. A dynamic programming algorithm is devised and used in a predictive control scheme by constantly feeding the conventional cruise controller with new set points. The algorithm is evaluated with a real truck on a highway, and the experimental results show that the fuel consumption is significantly reduced.

Basso et al. [31] propose an energy consumption model incorporating vehicle velocity and topographic profiles. Their literature research showed that most energy models are simply linear with respect to distance, though this can be dangerous, as the route topography and speed profiles are not considered in that assumption. They state that accurate prediction of the energy consumption is very difficult if the route is not known and state that second-by-second energy estimates cannot be accurately done without the route info. As an example, a freeway leading to a mountain pass will not have a constant energy consumption rate.

Tie and Tan [10] mention that global optimisation is based on knowledge of the energy requirement, both in the past and the future, proving that knowing the route is highly beneficial to optimisation.

To optimally manage the energy usage of a vehicle it is necessary to plan not just the instantaneous power usage but also strategize a plan for future consumption based on the route travelled [19], [20]. It is thus of great benefit if it is possible to identify not only the route being travelled but also the approximate location of the vehicle on this route.

4.2.4. Velocity

All transportation activities are governed by time. The logistical goals set for the transportation activity influence the time available to complete a given route. If more time is available the average velocity can be reduced, which will lead to a reduction in energy required. Zhou et al. [56] report that aggressive driving styles can increase fuel consumption by 68% for urban and 47% for rural roads. This shows that the velocity profile is crucial in the amount of energy required to complete a route. There now exists a trade-off between route completion time and route energy. Speed limits, traffic and compulsory stops,

such as traffic lights and compulsory weigh bridges also influence the velocity profile by adding additional constraints to a route. These constraints must be considered when generating a velocity profile for a vehicle on its route [13]. The vehicle type, mass and the route to follow are considered unchangeable for this study, leaving the velocity profile as the only variable that it is possible to improve on to reduce the energy consumed during a route. Johannesson et al. [13] reported that speeding up when going downhill and losing speed when going uphill will consume less energy, as you are using an efficient transformation of kinetic and potential energies between each other, reducing the peak power demands on the vehicle power source. Boriboonsomsin and Barth [12] reported that going downhill however doesn't produce as much kinetic energy as what is consumed when going uphill.

4.2.5. Optimisation Techniques

The field of optimisation is a very broad field with many different approaches and strategies that apply well to one type of problem but are ill-suited to another, it is thus important to first better understand the requirements for the problem at hand before diving into strategies to solve it. A summary of the optimisation problem for this study is given in Figure 4.2. The problem has many constants, like the route, comprising topography, surface finish and length. The vehicle type is also constant, so the aerodynamic characteristics will only be a function of velocity, with the rolling resistance force being a function of mass. Variables in this study include the vehicle mass, the route itself and the velocity profile. The constraints include factors like speed limits and completion time. All of these are input to the optimiser to try and locate a global minimum that yields the lowest energy possible to complete a given route in the available time constraint. The main variable to solve for is the route's velocity profile such that all constraints are met and a minimum amount of energy is consumed.

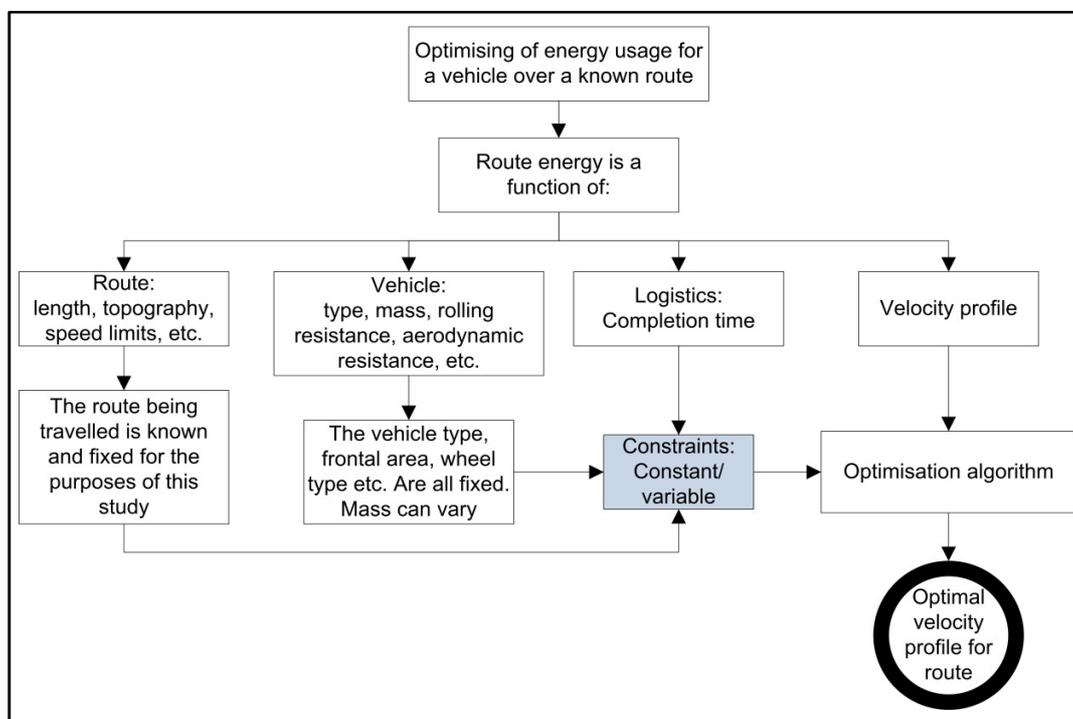


Figure 4.2: Factors influencing the route energy

Knowing these contributing factors allows us to exclude some of the available optimising strategies. A strategy like *fminsearch* is an unconstrained solver, which means that it does not apply to our study, as many constraints must be satisfied. One example is that velocity cannot be negative, as it will result in negative time steps in the solution, which is impossible.

A first-principles analytical model of the vehicle can be constructed without excessive difficulty, which is beneficial in model accuracy, but makes finding an optimal solution more difficult. Many variables are dependent on each other and are nonlinear, thus excluding linear programming. The equations can be approximated by linear alternatives but come at a cost of potential inaccuracies. Masses of data are not available, excluding the use of artificial intelligence (AI) and other numerically based approaches. Another problem with these approaches is that all the relevant parameters, like vehicle mass changes etc need to be recorded appropriately, and if a change is done outside of the trained values an incorrect optimal may be proposed. Dynamic programming requires previous knowledge of future requirements to be able to perform an optimisation, but is also simulation time intensive, making real-time application difficult. This data can be used to train a stochastic dynamic programming strategy, but again, a significant amount of data is required to make the model usable. Genetic algorithm approaches have the disadvantage that it is impossible to see what is going on inside of the algorithm, making it hard to comment on the answers obtained. [10]. Solvers like *fmincon* can solve the optimisation problem, but require a significant amount of time to find a solution, and are sensitive to the initial guess value input to the solver to find a global solution instead of a local minimum.

Some strategies to find an optimal velocity profile include tracking the movement of the vehicle in front of you and using that data for decisions on optimal velocities, though the information for the rest of the route may still be lacking [57].

Genikomsakis and Mitrentsis [58] built an analytical model of their system using a combination of analytic and piecewise functions to simplify the variables in their optimisation. Most of the 3d characteristics of variables such as efficiency plots are reduced to simple constant look-up values.

Zhang and Xiong [59] break the optimisation into three sections, similar to what will be performed in this study: find the optimal route velocity profile, then consider internal energy management. They use fuzzy logic controllers to classify their driving patterns and identify the route in real-time. They use dynamic programming to find the optimal control strategies for the different driving blocks. The use of DP is helpful as it allows running in real-time.

Hellström et al. [5] use road topography of a known route to estimate route energy. They develop an algorithm with a complexity low enough to run in real-time. They make use of DP. They noted that formulating the problem in terms of kinetic energy in their calculations is advantageous as it avoids oscillating solutions, reducing the time to find a solution. From this, they can deliver an optimal velocity profile for the route, reducing the total energy consumed.

Hellström et al. [41] devise a dynamic programming algorithm that is used in a predictive control scheme by constantly feeding the conventional cruise controller with new set points. The algorithm is evaluated with a real truck on a highway, and the experimental results show that the fuel consumption

is significantly reduced. Shen et al. [60] provide an improved velocity prediction strategy based on a Markov model, along with Neural Networks.

Shabbir and Evangelou [61] applied a real-time hybrid electric vehicle power-train optimisation on a series hybrid electric vehicle. They use a novel comprehensive and unified power train efficiency model that considers the engine-generator combination, detailed battery models and the power electronics. A control map that contains the optimal signals is produced offline, and then operated in real-time while the vehicle is driving. By implementing their method, they report significant savings in fuel usage as well as reduced battery loading. Fuel savings of up to 20% are reported over other HEV control strategies.

Kim et al. [62] consider vehicle parameters such as mass and other coefficients, route information such as grade, terrain type and speed limits, and terminal constraints such as velocities and battery state of charge in their co-optimisation strategy. They achieved 5.3-24.2% improvements in energy usage in a vehicle with aggressive driving over hilly terrain, and 0.5 to 5.3% improvement in relatively flat roads as compared to sequential optimisation.

Optimisation problems need to be classified correctly so that appropriate decisions can be made regarding what strategy to implement and which ones to avoid. Figure 4.3 summarises the factors influencing the optimisation strategy used in this study at various levels. A large amount of unnecessary information is omitted from that figure, such that it only explains the pathway used for this study [63].

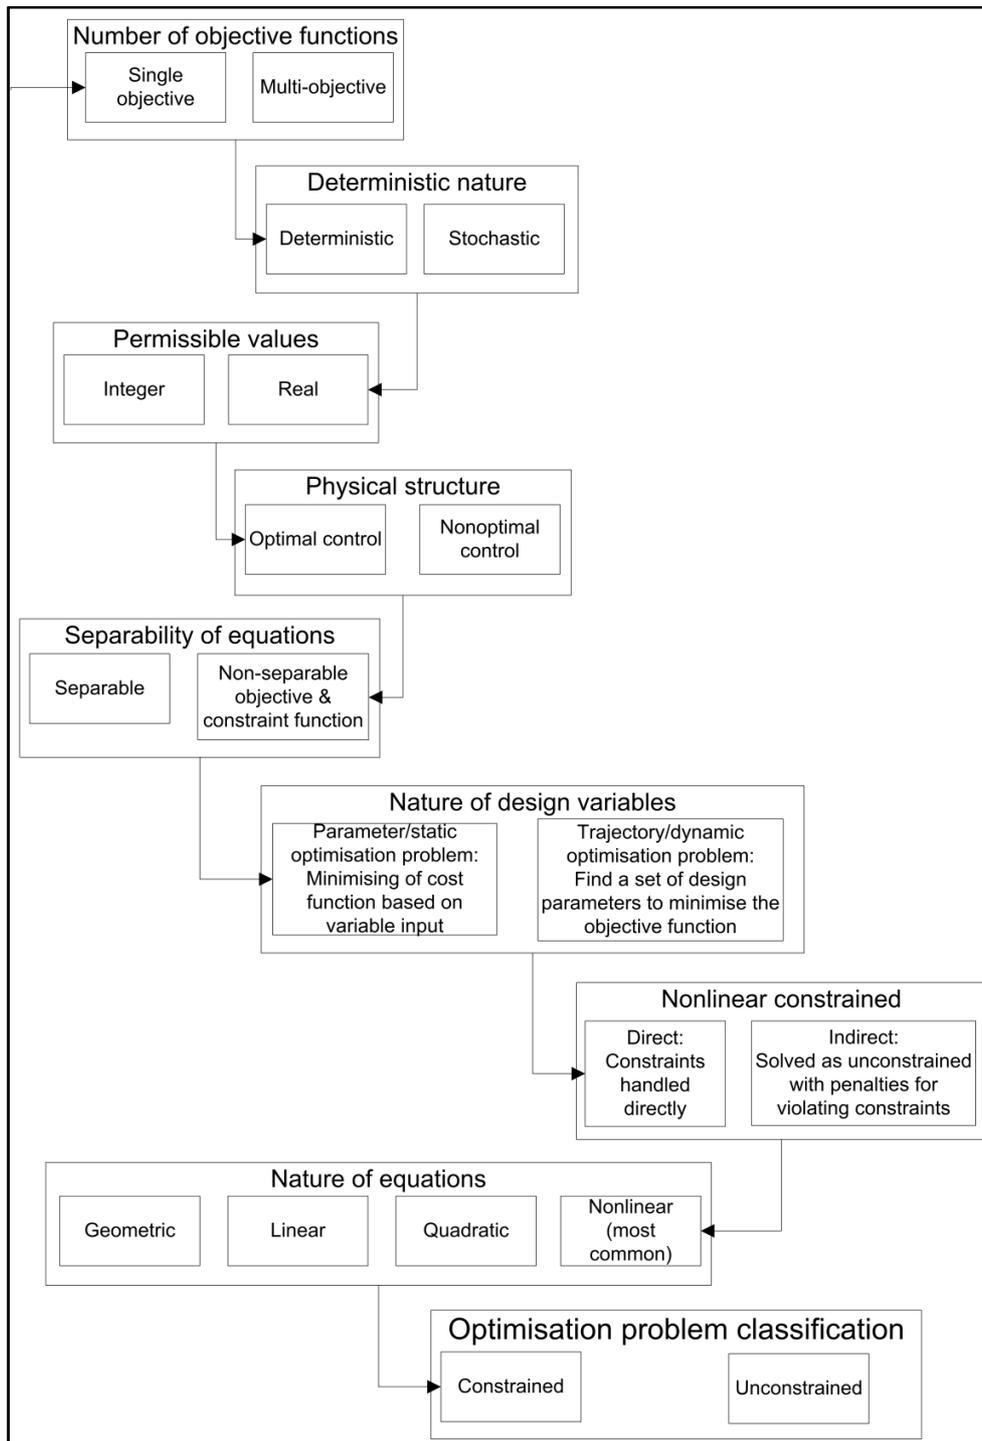


Figure 4.3: Optimisation problem classification [63]

Hu et al. [18] proposed a novel unified cost-optimal control scheme to consider all the stakeholders in hybrid electric vehicle (HEV) cost, yielding a 28-40% cost reduction. They considered mains grid charging cost, power management during driving, fuel cost and battery life models using rapid and efficient convex programming (CP). They performed comparative studies on vehicle-to-grid (V2G) and found that

CP application to PHEV is very similar in fuel economy to dynamic programming (DP) while running about 200 times faster.

Back et al. [19] mention that the route can't be simply optimised for topography and loads as many times the optimal plan will be inhibited by external influences, like traffic. Taking all the internal and external factors into account yields a very complicated optimisation problem that can't be solved in real-time. They propose Model Predictive Control (MPC) to simplify and solve this problem.

Zhang et al. [64] present a solution to the conflicts in vehicle management systems where, in one case one wants to optimise the vehicle's economy versus battery life, and in the other case compromise between the true global optimal execution versus the ability to implement it in real-time. The proposed system firstly comprises an internet-connected vehicle which makes use of powerful cloud computing to optimise the system using machine learning and stochastic dynamic programming, after which an onboard computing system with limited capabilities tries to adapt the driving in real-time to match the global optimal. The Internet Distributed Energy Management System (ID-EMS) outperformed a rule-based energy management system in terms of overall cost by 6.8% but was outperformed 7% by an acusal EMS with DP.

Mathematical optimisation is another term for nonlinear programming. The choice of algorithm is important as it governs whether a solution will be found quickly, slowly or even at all. Many times algorithms are not even applicable to a specific optimisation problem, and this is only found out later in the process after a lot of time has been invested into it. Gradient-based methods can be of great value whilst the other more general methods like particle swarms, genetic algorithms, etc. are often too computationally expensive to make their use feasible. Gradient-based methods potentially fall prey to only finding local minima, but this can be circumvented by clever use of the gradient-based methods, and at a fraction of the computational cost of methods like genetic algorithms (GA's). The global minimum is generally found by using local minimum finders multiple times with random starts or sweeping a grid of starting values [65].

4.3. Model

The total energy to complete a given route is mainly dependent on the vehicle and drivetrain type, the mass being transported, the route length and topography, and lastly the velocity profile followed throughout the route. The vehicle type is usually known in advance and can be characterised for use in the simulations. For this study, the mass is also known, as a detailed mass estimation strategy was patented as part of that work. For this study, the mass will be taken as a constant. With all the other influencers being route-dependent the best variable to optimise is the velocity profile, as the topographic profile of a given route is also not changeable. To be able to optimise the velocity profile for a route it is essential to know the exact details of the energy requirements for completing the route. Knowing the route and vehicle mass facilitates better optimisation to be performed on the energy consumption throughout the route.

Real-time route identification has been performed prior to this work in [6]. In that study, a previously travelled route was normally successfully identified in real-time within the first 3 to 4 km of travelling,

without any information known at the start of the run. The study did however mention that a route can initially be falsely identified, especially if the route overlaps with another route. As the amount of real-time route data increases the amount of certainty about the route can increase. This incorrect identification phenomenon should be considered when performing an optimisation, so that we don't optimise the velocity profile for the current route, but the velocity profile might look significantly different when travelling another possible route. This is implemented once an efficient optimising strategy is been developed and is presented in Chapter 5.4.

4.3.1 The Test Vehicle

The test vehicle used is a 1998 Isuzu Frontier 280dt, instrumented with a load cell on the rear differential to accurately measure the driving torque. This torque value is used to estimate the vehicle's gross combined mass. The exact details of the method are outside of the scope of this paper but can be found in reference [7]. An optical encoder implemented on one of the wheels measures the vehicle's displacement as a function of time. From this data, the driving power can be obtained accurately. A barometric pressure sensor is used to record the barometric altitude. The sensor has a resolution fine enough to detect relative altitude changes as small as 0.2 metres [34]. From this data the relative topographic profile and thus the incline characteristics of a route can be obtained [6]. With the vehicle's frontal area constant and known it is possible to build an aerodynamic resistance model as a function of velocity. A simple rolling resistance model is used to determine the rolling resistance force. For this optimisation study the mass was assumed to be constant.

4.3.2 Route Energy Model

A route energy model was built using Matlab/Simulink [40]. This model can determine the total energy to complete a known route based on the specific route's topography, vehicle-specific data, like aerodynamic drag and rolling resistance characteristics, the vehicle mass and the travel velocity. The advantage of using a more complicated Simulink model instead of a simpler coded model is that it allows the use of Simulink's proven built-in models for electric batteries that will be used in future optimisations to include this in the optimisation to extend the lifetime of the battery whilst optimising the velocity profile. The Simulink model used to calculate the route energy is shown in Figure 4.4. As input to the model it requires the topography of the route, simply input as an altitude vs. distance vector directly from the barometric data, the vehicle's mass, the coefficient of rolling resistance (C_{rr}), the velocity vector and the air density. The air density is assumed constant for this study, but can easily be made dependent on the altitude and temperature using the International barometric formula. The elevation of the tests is around 1600 metres above sea level, which is the reason for the low density number in the model.

The way this energy calculator is set up allows easy calculation of the total energy to complete a given route. The Simulink model also includes a regeneration model that allows calculating the total route energy if regeneration is possible. Future work involving the results of this study will find applications in electric and hybrid electric vehicles, which makes adding the regeneration model worthwhile. The regeneration efficiency can be input as a constant or even be made a function of other vehicle parameters, such as electric motor and driver efficiencies as a function of the loads. For this optimisation, the regeneration efficiency is assumed to be a constant value of zero, as this will allow a

comparison of route energies with that used by a normal IC vehicle. The model also gracefully facilitates a changing mass, instead of a constant value, by simply inserting the mass as a vector instead of a constant. This will be essential when trying to optimise the velocity profile for a delivery vehicle where significant changes in mass are expected. With the real-time mass estimation system [7] already functional, it will be easy to update the mass vector in real-time and thus perform a more appropriate optimisation.

The regeneration model allows potential and kinetic energy to be exchanged both ways without penalty, only applying an efficiency/penalty when converting any of those to another stored form of energy. When plotting the data for the vehicle travelling on the road a regeneration efficiency of 0% is used, as the IC test vehicle cannot store energy in another form than potential or kinetic, based on the vehicle mass, velocity and relative altitude.

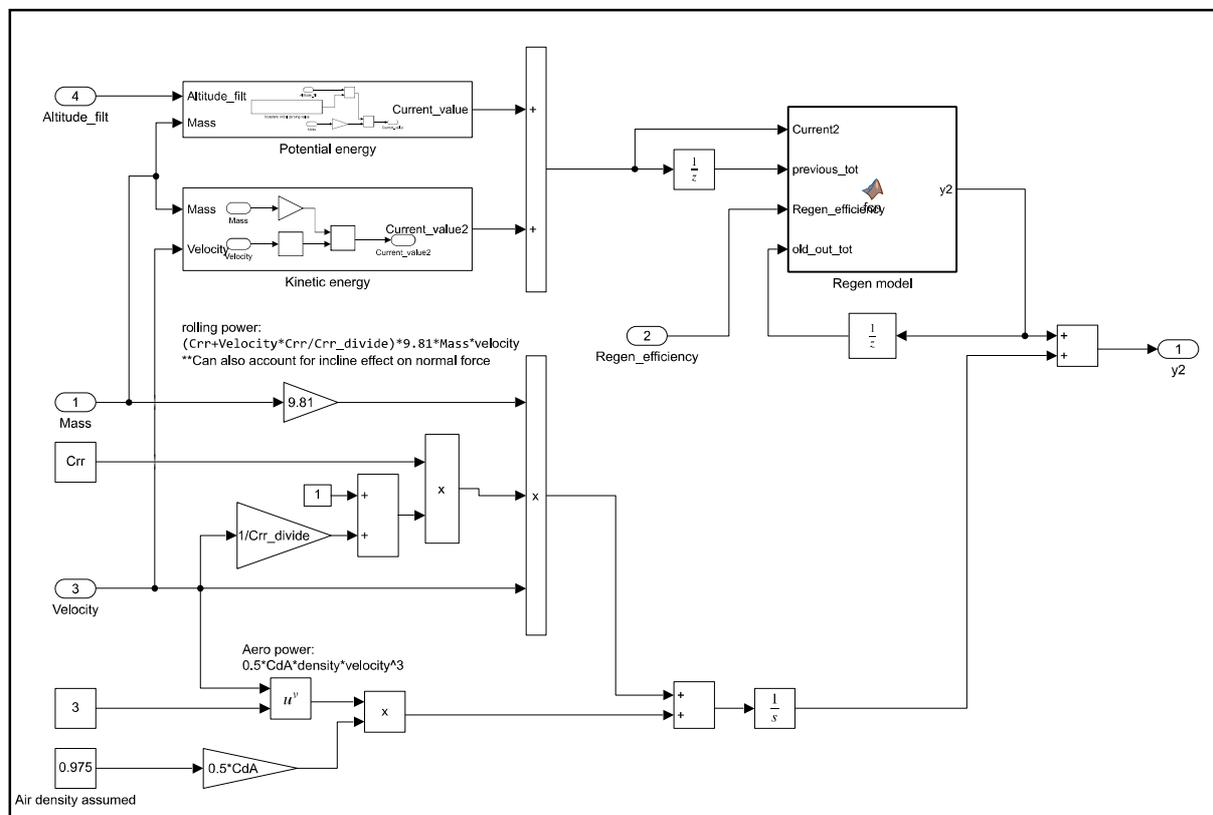


Figure 4.4: Route energy calculation in Simulink

There were four real-world test routes used in this optimisation study. These routes differ from those used in Chapter 2, as they were shorter and new sensing ability was added to the vehicle after those routes were recorded. Most of the simulations were first performed on route 4.1, with the other three used towards the end to test the ability of the optimisation strategy to work on other routes with different topographic profiles and lengths. The 4 routes' topographic profiles and original velocities are shown in Figure 4.5 a and b respectively. The routes had different total lengths as well. Route 4.1 had a length of 10 km, Route 4.2 a length of 7 km, Route 4.3 a length of 6.7 km and Route 4.4 a length of 9 km.

The routes were all loops starting and ending at the same location. The topographic data presented is based on the barometric altitude readings of the routes, which is why the final altitude may differ slightly from the starting altitude due to weather variances.

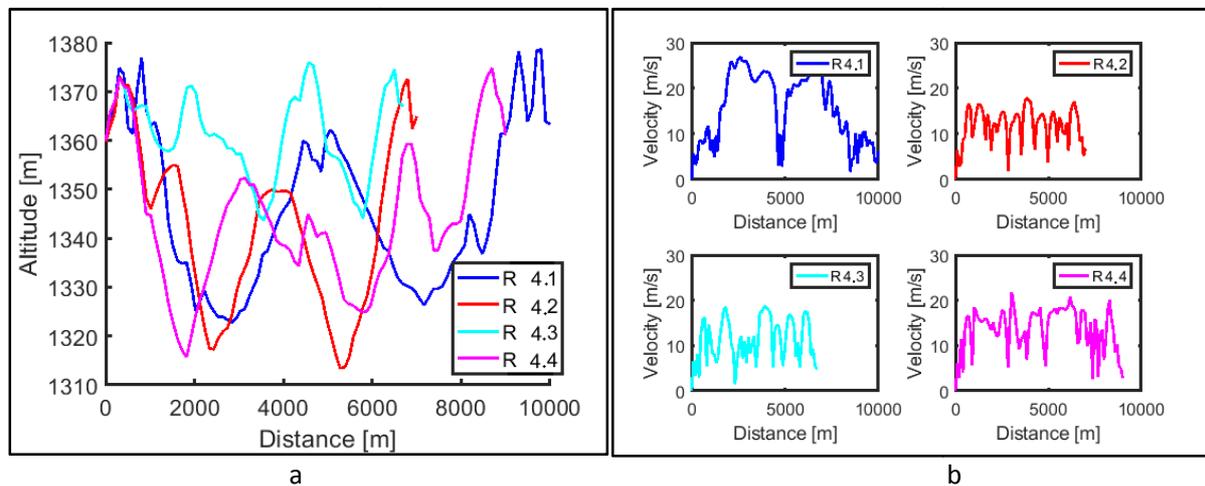


Figure 4.5: Real-world test routes. a) Topographic profiles, b) Original velocity profiles

For route 4.1, the energy consumed while travelling the route at the real-world travel time was determined using two different methods to validate the usability of the Simulink model. The first method, the green line in Figure 4.6, was a direct integration of the traction data over distance and time using a Matlab script. The data was recorded at 5 Hz. The second method, the black line, integrated the instantaneous energy in 100-metre increments using the stored route velocity vs. distance profile and the route's topography. The stored routes used for route identification in [6] use a very basic database where the data is represented in 100-metre increments, to reduce the sheer amount of data that needs processing. This database was updated with an additional column for the velocity, to use in initial route energy estimates for optimisation. Having less data to process helps an optimising strategy become feasible, making real-time optimisation a possibility. There is a big advantage in using distance instead of time on the horizontal scale. If time was used it would cause the horizontal axis to extend or lengthen based on the velocity profile, making comparisons more difficult. Having the horizontal axis in distance also allows the velocity for a section of road to be easily located relative to the actual route, the information of which is lost if time is used. Distance on the horizontal axis also corresponds with the route identification strategy presented by [6], which makes integrating this study into the greater study on optimisation of an underground electric vehicle a fair bit easier.

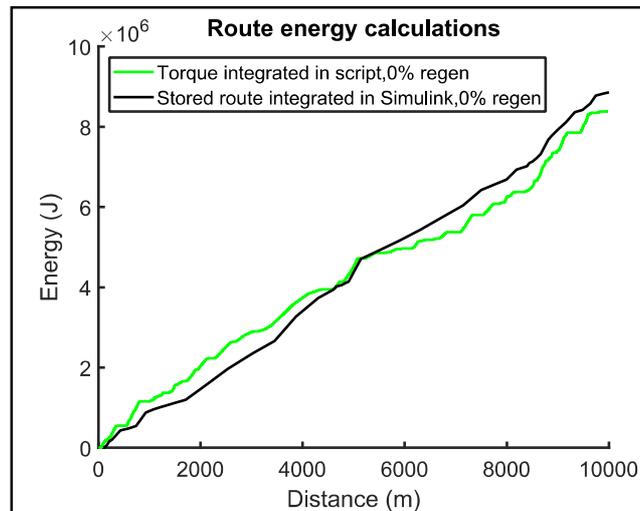


Figure 4.6: Route 4.1: Energy consumed vs. distance

From Figure 4.6 the small differences in the energy determined are due to some simplifications made in the processing, with the biggest being the reduction to only 100 data points for the 10 km long route 4.1. The green "Torque integrated in script" curve is considered to be the most accurate representation of the data as it used all of the data recorded at 5 Hz directly. It is noted that using the 100 metre incremented data does reduce the accuracy of the energy estimates, though it will still aid in the optimisation to be successfully carried out.

As an initial ballpark energy estimation, we can compare the fuel energy to the energy determined by torque measurement while travelling the route. The route completion energy from Figure 4.6 was stated at around 8.5 MJ. The calorific energy value for normal diesel is 45.5 MJ/kg. The route has a length of 10 km and the vehicle used has a typical consumption of 9 km/l, meaning that it consumed around 1.1 litres of fuel in the trip, equal to approximately 944 g of fuel: and 43.0 MJ of energy. A typical diesel engine from the 1990s would convert approximately 30% of the fuel energy to work. The drive train is assumed to have an efficiency of 70%, which means the energy from the fuel to complete the route of 10 km would need to be around 9 MJ, which, as a first-order approximation is rather close to the 8.5 MJ determined by the driving torque integration method. It should be noted that this is the tractive energy, and excludes energy to power accessories like air conditioning, electrical loads and power steering, which will all increase the vehicle's fuel consumption.

4.3.3 Sensitivity studies on route energy

To optimise the route energy for a vehicle travelling a known route it will be beneficial to study the sensitivity of the route energy requirement to various variables. The variables investigated include the velocity, mass and completion time. Other variables are out of the direct control of the operation of the vehicle and are not investigated. These include air density, rolling resistance, frontal area, drag coefficient and alternative routes, to name a few. Using the model developed in Simulink it is now possible to keep all of the variables fixed and sweeping only one variable at a time to show the sensitivity of the route energy requirement with respect to that variable.

4.3.3.1. Vehicle mass

One of the biggest contributing factors to the energy consumption of a vehicle travelling is the mass of the vehicle/rig (vehicle and trailer). In the equations for kinetic and potential energy, the energy is linearly proportional to the mass, which means that it is expected from this simulation that the overall route energy should be linear with respect to the mass of the entire vehicle. Using the Simulink model a reasonable constant velocity of 10 m/s for travelling the route was imposed and the effect of the route energy consumed was recorded as a function of vehicle mass (including trailer, if added). The simulated mass ranged from 500 kg up to 4000 kg. For this simulation the regeneration ability, like found in electric vehicles, is ignored, to keep the results comparable to the IC real-world data. The route starts and ends at the same location, making the net effect of potential energy zero.

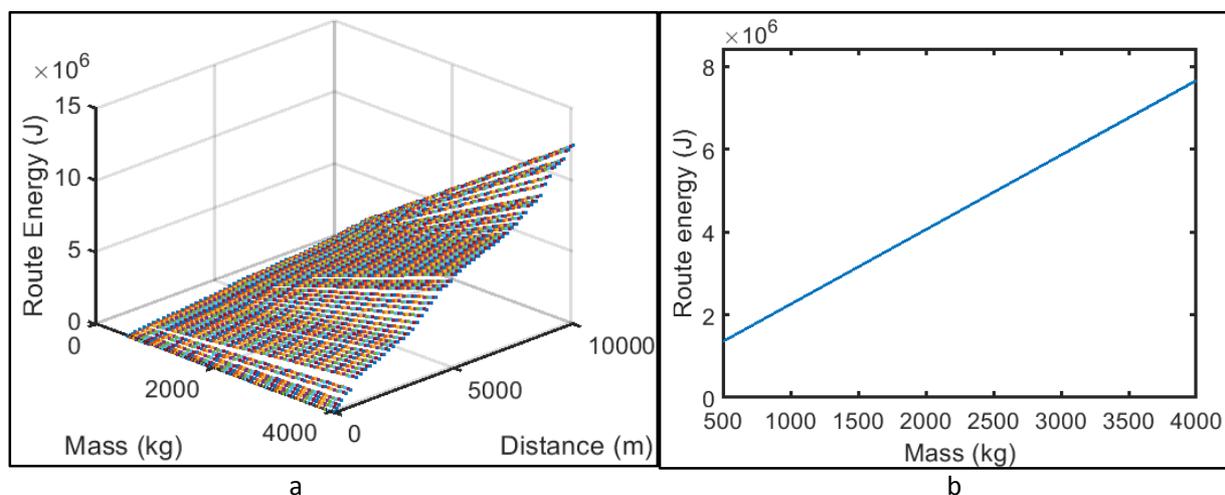


Figure 4.7: Route energy vs. Mass for real-world test route 4.1

In Figure 4.7b it is noted that the consumed route energy is a linear function of total mass, which corresponds to the work presented by [31]. The offset from zero is due to aerodynamic drag, a constant for the velocity used for this simulation.

4.3.3.2. Vehicle velocity

The kinetic energy equation as well as the aerodynamic drag are both a function of the square of the velocity term. It is expected to see a significant quadratic characteristic for the energy requirement when a sweep of velocity values are simulated. A simulation was performed for a sweep of velocity from 1 m/s up to 40 m/s to investigate the effect of vehicle velocity on the energy required to complete a route. The vehicle mass was set to a constant value of 2150 kg, a typical unladen mass for the test vehicle. Figure 4.8a shows the energy consumed as the vehicle travels the route as a function of a constant velocity for the entire route. Figure 4.8b shows the quadratic relationship between route energy and velocity.

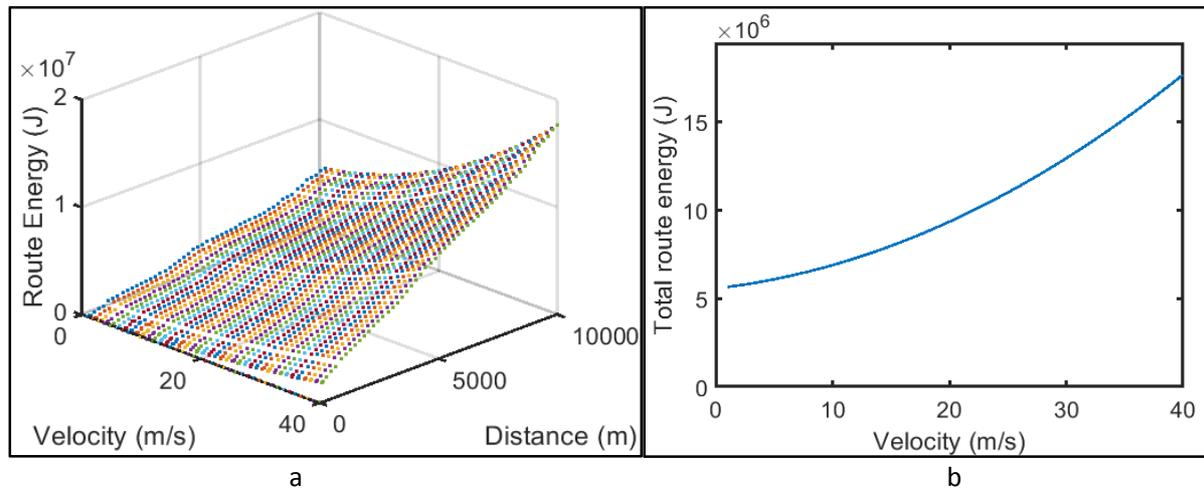


Figure 4.8: Route energy sensitivity to velocity. a) Route energy for the entire route with varying velocity, b) final route energy for varying velocity

This shows that lower velocities will consume less energy to complete the route, as the aerodynamic loading and the rolling resistance are both functions of the velocity. There is however a limit to the longest time that can be used to complete a route based on logistical limitations. It is thus important to investigate the effect of route completion time on the total energy consumed, which is performed next.

4.3.3.3. Maximum route completion time

The vehicle mass was kept constant, and the velocity swept to yield completion times from 250 seconds (40 m/s constant speed over 10 km) up to 10,000 s (1 m/s constant speed over 10 km). In Figure 4.9 the total route energy as a function of route completion time is shown. This function is expected to have a reciprocal curve shape, as route energy is inversely proportional to time.

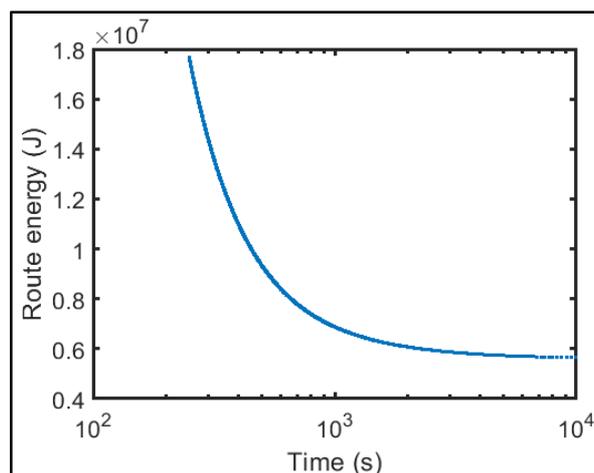


Figure 4.9: Final route energy vs. completion time for real-world test route 4.1

It is noted that the energy does not reduce down to zero but approaches a minimum threshold for times longer than approximately 1 hour (3600 s), approaching infinity. Thus, little benefit is seen for overly extending the route time. It is however noted that the route energy does increase drastically as the

route time becomes shorter than about 800 s. A clear trade-off between the maximum route completion time and the amount of energy consumed is evident. For optimisation, it will thus be essential to specify the maximum allowable time to complete the route, with the point of lowest energy consumed most likely lying close to this constraint.

4.3.4. The Optimisation Strategy

In the context of velocity profile optimisation, gradient-based optimisation techniques may or may not be suitable, depending on the specific problem and constraints. Velocity profile optimisation typically involves determining the best velocity profile for a system, often in the context of motion planning, robotics, or control systems. The objective is to find the velocity profile that minimizes some performance criterion, like overall route energy, while adhering to system constraints, like allowable completion time. Here are some considerations:

Smooth and Continuous Profiles: If the velocity profiles are represented as smooth and continuous functions, gradient-based optimisation can be effective. In such cases, you can compute gradients with respect to the parameters that define the velocity profile (e.g., acceleration limits, time durations), which allows for efficient optimisation.

Non-Smooth or Discontinuous Profiles: In some scenarios, velocity profiles may not be easily expressed as smooth functions. For example, they could involve abrupt changes or discontinuities due to specific system requirements or constraints. Gradient-based methods may not perform well in such cases, as the objective function may not be differentiable everywhere.

High-Dimensional Profiles: If you have a high-dimensional parameter space, such as optimising velocity profiles, gradient-based optimisation can be computationally expensive and may require careful initialization and tuning to find a good global minimum instead of some local one in a realistically achievable timeframe.

Safety and Feasibility Constraints: Velocity profile optimisation often involves safety and feasibility constraints, such as avoiding collisions or respecting velocity and acceleration limits. Handling such constraints can be challenging in gradient-based optimisation.

Real-Time Requirements: In some applications, velocity profile optimisation needs to be performed in real-time, and computational efficiency becomes crucial. Gradient-based methods may not be well-suited for real-time applications due to their computational overhead.

In summary, whether gradient-based optimisation is appropriate for velocity profile optimisation depends on the nature of the problem, the representation of velocity profiles, and the presence of constraints. If the problem can be formulated as a smooth and continuous optimisation task, gradient-based methods can be efficient.

Firstly, the controllable variable that has the greatest influence on the route energy needs to be found. The vehicle mass is a user requirement, so investigating changing the mass would not be helpful to the problem. The completion time is a logistical constraint, which will be used as a constraint in the

optimisation, but not as the variable to be optimised. The route selection, and thus topography, is considered a logistical requirement so that no optimisation will be attempted for that variable in this thesis. We are thus left with solving for the optimal velocity profile/driving behaviour to reduce the energy requirement for a known route, given some time constraints.

For the initial attempt to optimise the energy usage of the vehicle, it is noted that this vehicle will be modelled as a first-principles model, which will represent the dynamics and energy requirements most accurately. Once the first-principles model has been exploited to yield optimised velocity profiles it can be investigated whether surrogates might be necessary for real-time implementation of the model.

A common optimiser function, *fmincon*, in Matlab, was used in this project. The function, *fmincon*, is a gradient-based method that is used to find the local minima of a problem. One of the most important aspects of these types of optimisers is having the best initial guess possible. Having initial guess values close to the actual global minimum increases the likelihood of the solver finding the global minimum. If the initial guess is not close to the global minimum there is a significant risk that the solver will get stuck in a local minimum, or might never reach a solution if the constraint function is very noisy.

Different ways will be investigated to add diversity to the initial guess of the optimisation. The following strategy was followed to find the minimum route energy solution:

1. Set a route completion time value that corresponds to the real route travel data, so that realistic comparisons on final route energy can be made.
2. Divide the route into a number of block increments such that an optimal velocity profile can be found for the given number of block increments. This number of blocks relates to the model fidelity. In effect, the optimised value for a block will be equal to the average velocity over the length of that block increment, neglecting acceleration when there is a difference in adjacent block values. Once the optimisation solver is operational the fidelity can be investigated to find whether a fidelity is beneficial to the problem and whether there exists a point of diminishing returns. A lower fidelity however drastically shortens the simulation time.
3. Run the proven optimisation solver; *fmincon*, to find the local minimums of the constraint function.
4. Sweep the initial guess values for the local minimum solver to try and determine whether we can potentially find a global minimum. We will start with a low fidelity (number of velocity block increments) to make the sweep viable. Sweeping the velocity with a decent resolution for a high model fidelity involves a massive amount of simulations, which is simply not feasible, as is the case in most optimisations.
5. As an initial strategy, the initial guess is plausible, but completely random values. This can be performed multiple times. The results are used to log a success rate for different fidelity values.
6. After the completely random starting values, constant valued vectors are used. The elements of the vector are set equal to the average velocity that needs to be travelled at to complete the route in the longest possible time, which should be close to the lowest amount of energy.
7. Evaluate the best of the optimal vectors and determine if any trends are evident.

8. Once a good initial guess vector can be generated, the robustness can be investigated by running the optimisation with slight random variations on that initial guess, evaluating the consistency and values of the results obtained.

In the following section, the simulations performed and their results are discussed, after which conclusions can be drawn from this study.

4.4. Simulations & Results

As mentioned in Chapter 4.3, a first-principles model is used to determine the route energy when travelling a known route at a specific velocity profile. Constraints in the form of maximum allowable completion time and bounds for the allowable speeds are set. As previously mentioned, the solver used to optimise this problem was the *fmincon* solver in Matlab. This is a well-documented function that performs well at finding local minima. However, careful consideration is needed when choosing a starting point for the solver such that a global minimum can potentially be found.

Initially, a study is performed on the optimal velocity profile as a function of the route topography, starting with simplified profiles to study the correlation, then moving on to real-world route data after sufficient information has been gathered. The initial guess is then systematically adapted to a point where the strategy can prove that a dependable answer is obtained. The model fidelity is a measure of the resolution of the optimal velocity profile, with each element in that vector representing the average velocity over a set distance increment (route length/number of increments to be solved). The model fidelity is varied from low to high to investigate the advantage of spending more computational time on the simulation. The following chapters document the steps performed to develop a robust method of finding an optimal velocity profile for a vehicle travelling a known route. A schematic of the simulation approach is presented in Figure 4.10.

The constraints govern the operational range for the solver to attempt to find a solution in. These constraints were left as open as possible, to not limit potential realistic solutions. The main constraint is completion time, which is a logistical constraint that restricts how much time is spent on a route. In Chapter 4.3.3.3. we showed that the completion time has a large effect on the amount of energy required to complete a route and that longer times result in lower route energies required, with a taper-off after which a longer time doesn't reduce the energy significantly. Speed limits are necessary to ensure that the solver does not search for negative velocities, while the upper limit is left reasonably open initially, as it should avoid high velocities by default, due to the quadratic nature of energy as a function of velocity. The upper bound for the velocity can be used in future simulations to impose speed limits that may be required for certain stretches of road. The purpose of this study is thus to investigate initial guess strategies that can improve the likelihood of the *fmincon* solver to find the global minimum.

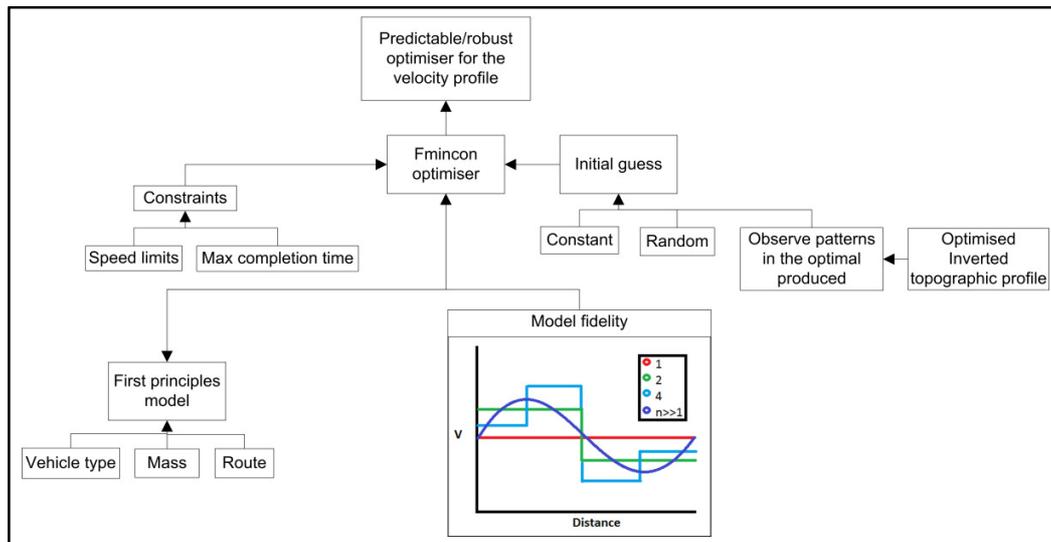


Figure 4.10: Velocity profile optimisation framework

4.4.1. Road grade sensitivity analysis

Road grade plays a significant role in the energy required to complete a given route, and will thus greatly influence the shape of the optimal velocity profile. To study the effect of road grade on route energy and the associated optimal velocity profile, a number of roads with various topographies will be simulated. An initial simulation is performed for a constant length of road with varying topographic (incline) characteristics, at a set maximum completion time value. The various simulated route profiles are plotted in Figure 4.11. The first is a flat route, then 2 %, 5 % and 10 % inclines and declines, a 5 % up-down, and a 5 % down-up with the last test routes being 7 % up-down-up-down and down-up-down-up respectively. A low model fidelity will lack the resolution required for a route with varying topographic contour, so the fidelity for this simulation will start with a value of 10, meaning that the optimal velocity will be comprised of 10 equally spaced blocks with an average velocity for each block. After this a model with a fidelity of 50 will be used as well, to be able to study the effect of fidelity on the quality of the answer. The optimal velocity profiles obtained by the *fmincon* solver for these test routes are plotted in Figures 4.12 through 4.15 for a fidelity of 10 and 50. These results can now be studied so that a better understanding of the effect of topography may be gained. For these simulations, the velocity constraints imposed were very light in that the answers should not be affected by them. The minimum velocity is set to 1 m/s, to avoid unrealistic and negative velocity answers, and the maximum was set to 50 m/s. These velocities are expected to be outside of the typical optimal answers with the maximum completion time constraint and minimising of route energy being performed. The effect of regeneration, as in the case of electric and hybrid vehicles, is ignored for this simulation and will be covered later in Chapter 4.4.6.

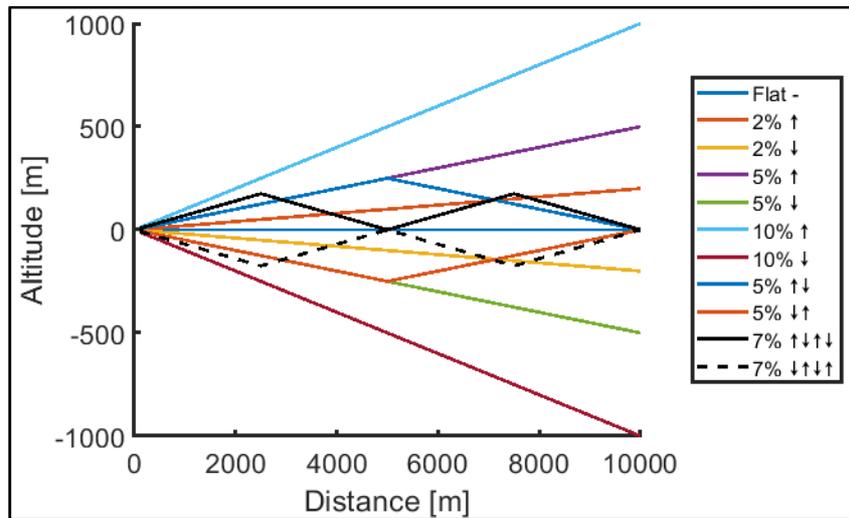


Figure 4.11: Standard topography cases for evaluating optimal velocity results

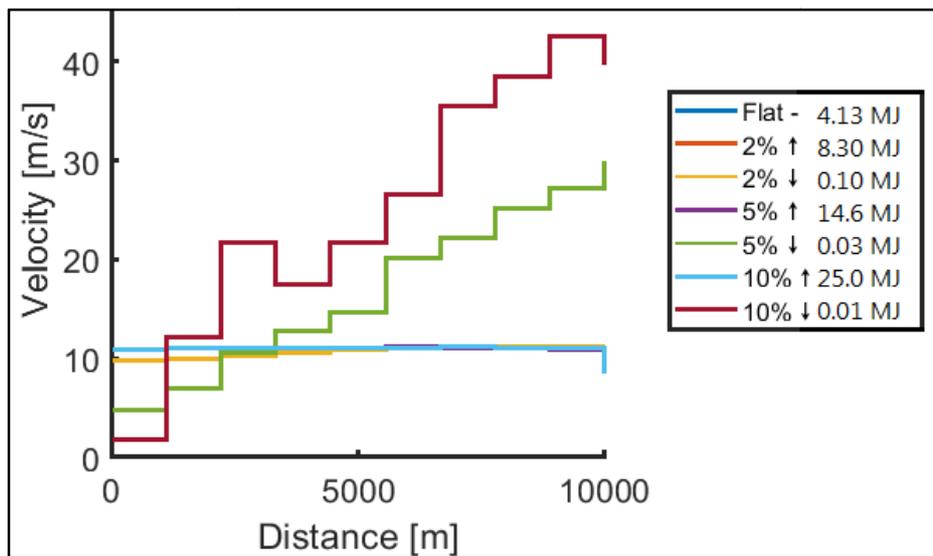


Figure 4.12: Optimal velocity profiles for straight profile routes with a fidelity of 10

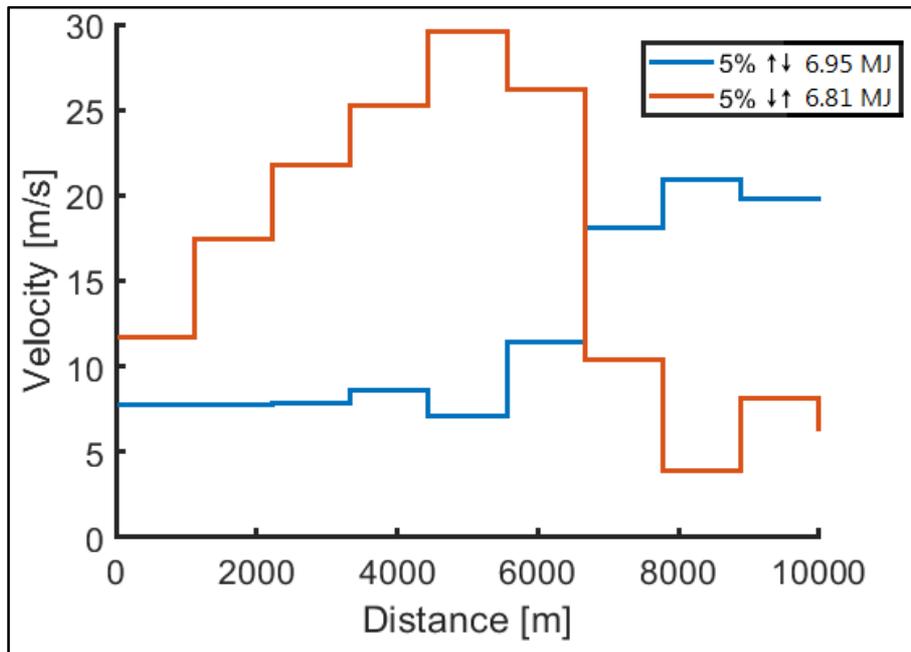


Figure 4.13: Optimal velocity profiles for 5 % Up-Down & 5 % Down-Up with a fidelity of 10

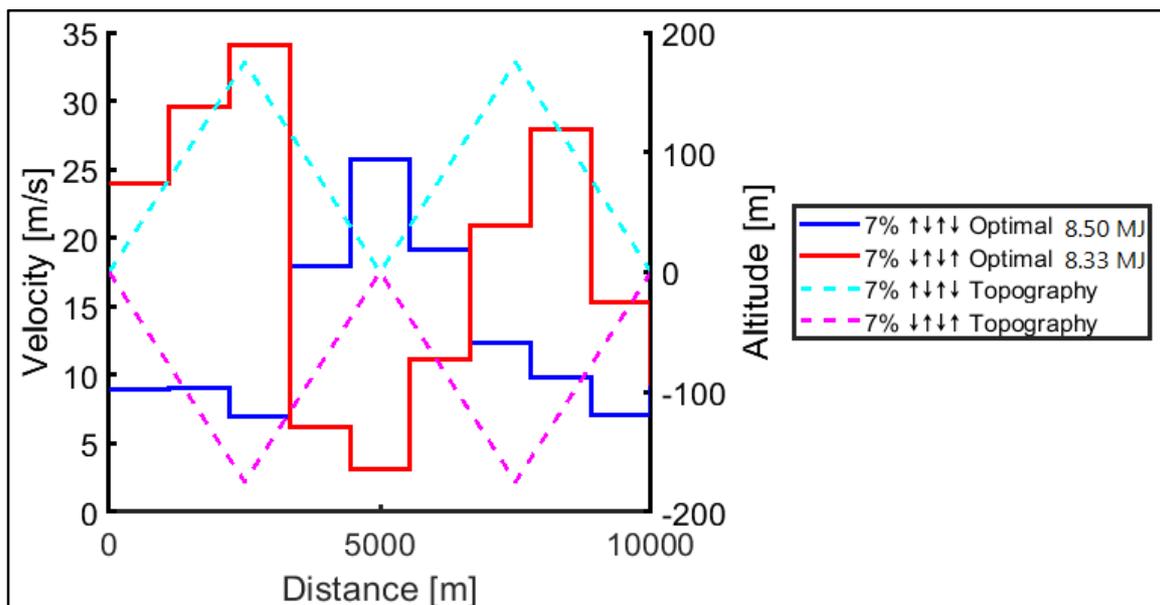


Figure 4.14: Optimal velocity profiles for 7 % Up-Down-Up-Down & 7 % Down-Up-Down-Up with a fidelity of 10

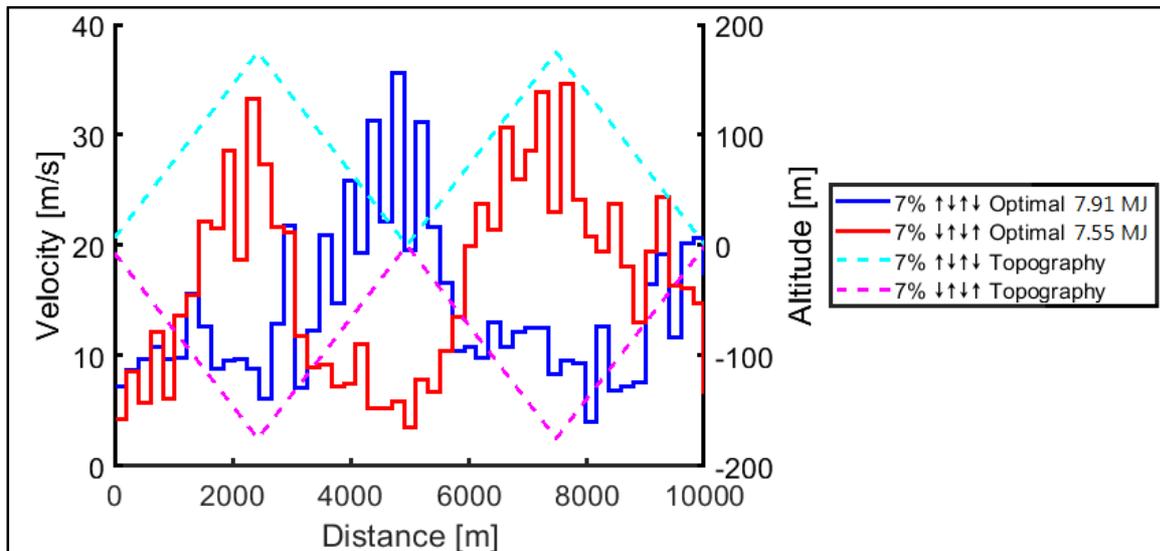


Figure 4.15: Optimal velocity profile for 7 % up-down-up-down and 7 % down-up-down-up with fidelity of 50

From Figure 4.12 to 4.15 it can be seen that the optimal velocity profiles have some correlation to the topographic profile of the route. The flat route and the constant incline routes all delivered a constant optimal profile. This makes sense as no energy can be gained from exploiting route profiles, and the vehicle is then left to travel the route at the minimum velocity required to complete the route in the maximum amount of time allowed. The final amount of energy used will however differ, based on the incline. There however is a difference when the profile has declines in it, where potential energy can be exchanged for kinetic, as reported by Johannesson et al. [13]. By allowing the vehicle to speed up when going downhill the velocity is increased with no additional energy consumed. In Figure 4.12 it is however noted that the 2% decline route still has an almost constant profile. This is due to the losses on the vehicle by aerodynamic drag and rolling resistance causing all of the forces to cancel out, yielding a close to zero total energy consumption. Any decline steeper than approximately 2% will result in an increase in the velocity of the vehicle without additional energy being consumed. It is noted from Figure 4.13 to 4.15 that the incline-velocity deviates from a constant if the route has a non-constant profile, such that velocity can be increased when going downhill and reduced when going uphill. This shows that the velocity in a specific location is influenced by future topographic changes. The routes that start with an incline will perform the initial part of the route at a constant velocity, and thereafter the profile changes based on route topography changes. From the 7% routes presented in Figures 4.14 and 4.15, it is noticed that there exists a correlation between the shape of the optimal velocity profile and that of the route topography, which is included in the figure as a dotted line as a reference, such that the one is almost an inverse of the other, apart from the constant initial velocity part if there is an initial incline. This effect will be studied further in Chapter 4.4.5. If the route has a significant decline in the first part it allows the vehicle to use the downhill to gain velocity without significant energy requirements from the vehicle energy supply, resulting in less overall energy consumed. This is evident from the 7.55 MJ required for the 7% Down-Up-Down-Up compared to the 7.91 MJ for the 7% Up-Down-Up-Down route, which both start and end on the same altitude for a fidelity of 50.

A low fidelity of 10 doesn't facilitate smooth tracing of the topography for the optimal velocity profile, resulting in a less optimal solution of 8.5 MJ for the 7% Up-Down-Up-Down route, compared to 7.91 MJ for the model fidelity of 50. A higher fidelity can better trace the optimal velocity profile as a function of route topography. Still, it may also introduce more noise, as is seen in Figure 4.15. A higher fidelity also greatly increases simulation time. The next sections will look into ways of finding better initial guesses for the *fmincon* solver to yield a more robust method of locating the minimum of the cost function.

4.4.2. Constant initial guess of distance/time

Now that an understanding of the optimal velocity profiles is obtained, studies can be performed on how to find the optimal solutions for real routes. The first simulations were performed on the test routes listed in Chapter 4.1. After the test routes' simulations are completed, optimisations can be performed for route 4.1 using an initial guess for the velocity vector with constant values throughout, equal to the total distance divided by the maximum time allowed. The maximum time allowed was taken from the real route data when driving the vehicle, as this allows direct comparison of energies between simulated and real data later on.

4.4.2.1. Constant initial guess for test routes

In Chapter 4.1. we investigated some standardised topographic profiles for test routes. The initial guess to even these simplified routes played a significant role in the solution to the optimisation. If a constant valued vector was used as the initial guess, the optimal result also often has constant sections. Figure 4.16 shows the optimal vectors for the 5 % Up-Down and Down-Up test routes when the initial guess vector had constant elements throughout. Figure 4.17 shows the optimal velocity profile found when using a constant initial guess for the 7 % Up-Down-Up-Down and Down-Up-Down-Up routes. It is interesting to note that a constant initial guess gave very similar velocity profiles for the 5 % Up-Down and Down-Up, as well as the 7 % Up-Down-Up-Down and Down-Up-Down-Up routes respectively. This does not make intuitive sense as it would be expected that the optimal profiles should be somewhat inverted with respect to each other, seeing as the topographic profiles are inverted. The constant incline routes also simply gave constant velocities, which means that the ability to gain velocity from the topography is missed. The optimal route energies for the test routes with a constant and a better initial guess are presented in Table 4.2. The method of obtaining the better initial guess is presented in Chapter 4.5.

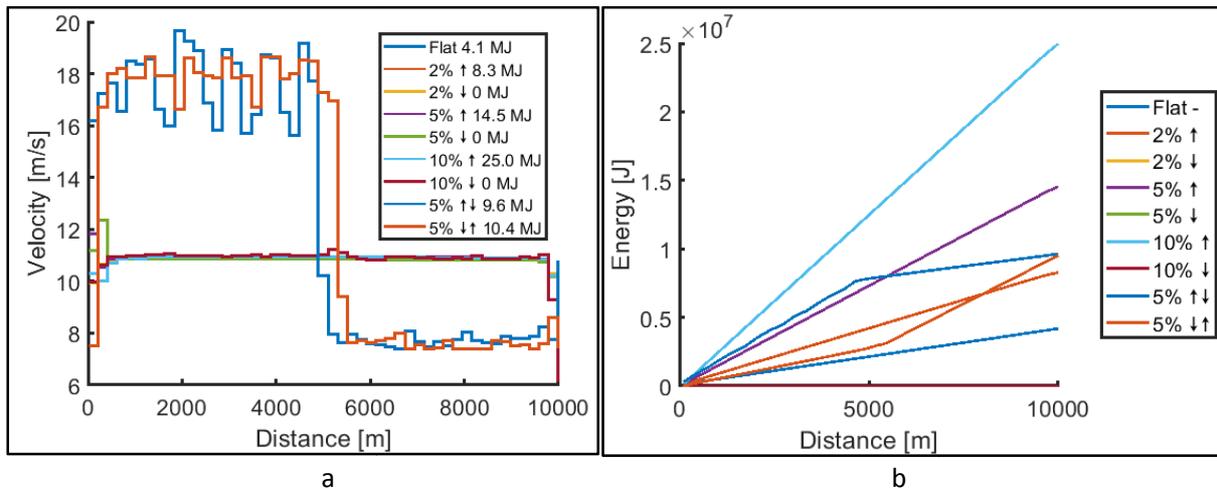


Figure 4.16: Constant initial guess for test routes, fidelity of 50. a) Optimal velocity profiles, b) Energy

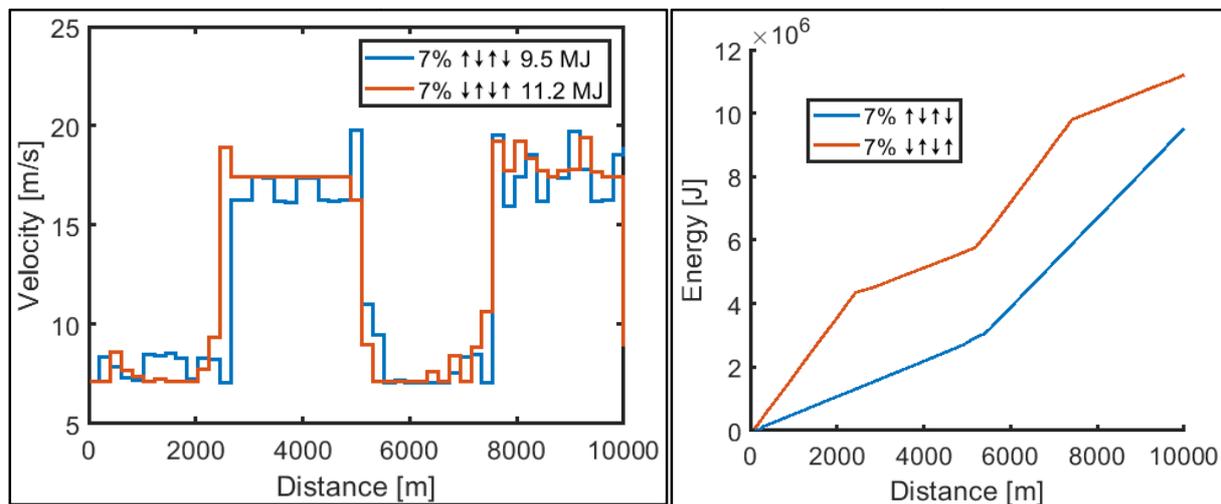


Figure 4.17: Optimal velocity profile for constant initial guess, 7% routes, fidelity of 50

Test profile shape (10 km length)	Route energy (MJ)	
	Constant initial guess of 10 m/s	Better initial guess
Flat	4.1	4.1
2% Incline	8.3	8.3
2% Decline	0.0	0.10
5% Incline	14.6	14.6
5% Decline	0.0	0.0
10% Incline	25.0	25.0
10% Decline	0.0	0.0
5% Up-Down	9.6	7.0
5% Down-Up	10.4	6.8
7% Up-Down-Up-Down	9.5	8.5
7% Down-Up-Down-Up	11.2	8.3

Table 4.2: Final route energy using various initial conditions for a model fidelity of 50

Looking at the velocity profiles presented in Figure 4.17 and the route energies in Table 4.2, the routes with constant inclines throughout yielded the same optimal, almost constant velocity profile. For the routes where the incline is not constant a significantly improved solution may be obtained from an improved initial guess to the optimisation, showing that a purely constant initial guess is not a very good start to the optimisation problem if a better initial guess is possible. The strategy to develop a better initial guess is presented in Chapter 4.5.

The zero final energy values seen in Figure 4.16 and Table 4.2 were because regeneration was ignored, thus meaning that a steep enough downhill can convert potential energy into kinetic energy at a high enough rate that traversing the route consumed no additional energy, though nothing could be stored onboard the vehicle. A decline of 2% was found to be very close to the transitional point where a route switches from requiring to not requiring energy from the vehicle powertrain, which explains the 0.1 MJ required in the better result while the constant initial guess yielded a slightly lower value. If regeneration and storing of energy were possible, these values would go into negatives, meaning that energy is recovered by travelling downhill. This phenomenon is covered in Chapter 4.6.

4.4.2.2. Constant initial guess for real-world Route 4.1

A constant initial guess strategy will also be applied to finding an optimal route velocity profile for a real-world test route, route 4.1, and the results can be compared to the energy consumed when travelling this route in a non-optimised normal driving manner. For this simulation, the optimiser didn't include any speed limits or compulsory stops, but standstill time was removed from the real route travel data so that only moving data is used for comparison purposes. The optimisation was started using a fidelity value of 1, to test whether the *fmincon* program does perform as expected and that the model functions correctly. In this study, a fidelity of 1 implies that a constant velocity will be maintained for the entire route, with a fidelity of 2 implying that the route is divided into two equal-length parts, with the average velocity for each of the parts being the solution of the optimiser. The solution to the optimisation problem for a fidelity of 1 can be calculated by taking the total distance and dividing it by the maximum allowable time, as this will be the least amount of energy possible to complete the route, based on the lessons learned from the sensitivity study in Chapter 3.3.2. It should be noted that there is a compulsory constraint that is applied to the velocity vector elements that they have to always be > 0 m/s. If the *fmincon* solver were to guess a velocity lower than zero it would cause the runtime value in the Simulink model to become negative, causing the program to crash. Once it was confirmed that the *fmincon* solver worked as expected the fidelity was incrementally increased. For this study, the velocity blocks were equally spaced throughout the length of the route, but variable length blocks can be implemented at a later stage if deemed necessary. The number of average velocity blocks used in the optimisation is the model fidelity, which was increased from 1 to 2, 4, 5, 10, 20, 25 and 50 over the 10 km route distance. The reason for the spacing is that the 100 stored route information blocks at 100-metre distance increments can be divided by these amounts of blocks without having extras, fractions or needing interpolation, making everything easier to code and interpret while still allowing the necessary conclusions to be made for the study. This can easily be adapted later to account for routes which do not have such convenient lengths.

Using the initial guess of a constant velocity for the whole route, for the various fidelity options the optimal results and simulation times were recorded and are shown in Table 4.3. A plot of the optimal velocity vectors for various fidelity values is shown in Figure 4.18.

Number of blocks	Route 4.1 final energy (MJ)	Simulation time (s)
1	7.02	0.6
2	7.29	86
4	6.79	257
5	6.66	76
10	6.51	377
20	6.21	2363
25	6.16	2873
50	5.82	7259

Table 4.3: Optimal route final energy for a constant initial guess

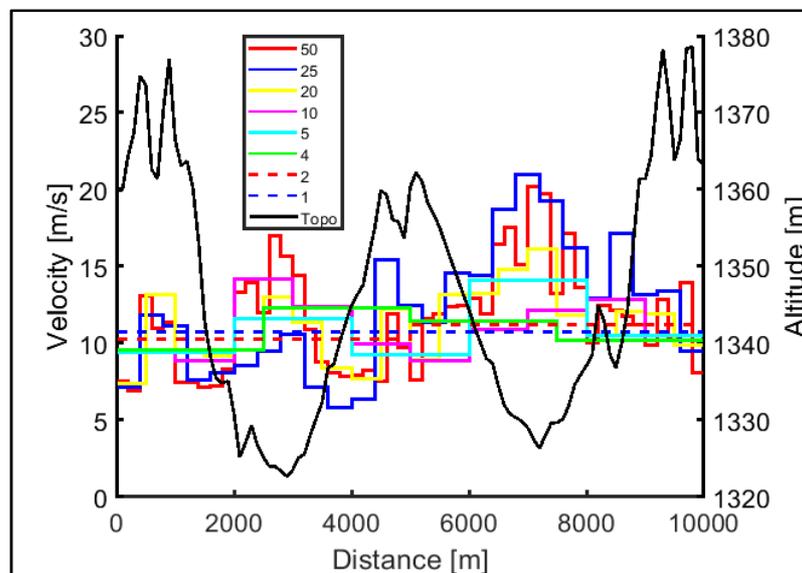


Figure 4.18: Optimal velocity vector for constant initial guess

It is seen in Table 4.3 that as the fidelity increases the solver can find a lower minimum route energy. This makes sense as a higher fidelity will allow better resolution to more effectively trace the topographic changes. This can be described as riding a rollercoaster, which corresponds to the results reported by [13]. The drawback of higher fidelity is that it increases time drastically, and sometimes the noise in the solution as well, as was seen for the 50 fidelity data of Figure 4.18. A plot of the simulation time for the various fidelities used in this study is shown in Figure 4.19.

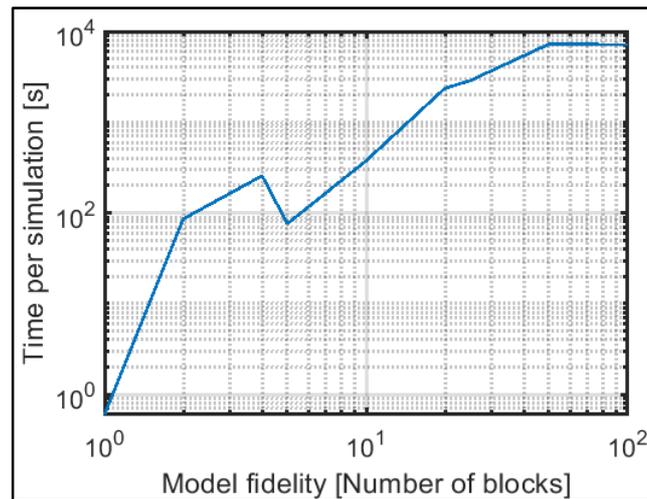


Figure 4.19: Simulation time vs. model fidelity for distance/time initial guess

From the simulations performed it is evident that there exists driving styles that show promising improvements in route energy consumption. The most saving shows an improvement of $(8.5-5.82)/8.5 = 32\%$ in route energy requirement. It should be noted that the optimal answer has not been constrained in terms of exact speed limits, maximum power or compulsory stops at this stage. The velocity was constrained to be below 40m/s, which is higher than the maximum allowable highway speed in South Africa, though the optimal values did not touch that limit. Driver comfort has also not been evaluated yet for constantly accelerating and decelerating.

Now that we have a model that seems to work, we can develop a method that can deliver better initial guess values to yield an improvement in the final answer or a reduction in simulation time.

4.4.3. Random initial guesses for varying fidelity

To try to verify whether a global minimum was found in the previous step a random start approach was implemented multiple times to see if any other lower route energy values could be found compared to the values from the constant velocity initial guess from Chapter 4.1.

The results show that the constant initial guess was not able to always find the global minimum value. Random initial guess simulations were carried out 10 times for each model fidelity value to see if any better answers existed. A higher number of random tries will increase the likelihood of finding the global minimum, but this comes at the cost of enormous processing time. A sweep function could also be implemented for lower model fidelities which is discussed in the next section, but very quickly becomes unfeasible as the fidelity increases. Table 4.4 shows a summary of the best and worst of the 10 random

initial guess simulations, along with the standard deviation for these values. The initial guess values were restricted to lie between 1 and 40 m/s, which are reasonable values for the route.

Model fidelity (Number of blocks)	Lowest route energy (MJ)	Highest route energy (MJ)	Standard deviation for 10 simulations	Average time per simulation (s)
1	7.02	7.02	0	4
2	6.70	8.01	0.37	45
4	6.572	8.081	0.47	224
5	6.513	7.375	0.34	290
10	6.443	7.984	0.49	891
20	6.38	7.71	0.42	2382
25	6.23	8.83	0.79	6038
50	6.101	8.10	0.78	14316

Table 4.4: Random starting points route 4.1 energy data

The data from Table 4.4 is represented graphically in Figure 4.20. It is noted that a higher fidelity may result in lower route energy values, but also show an increase in standard deviation of the route final energy value, and a significantly increased simulation time.

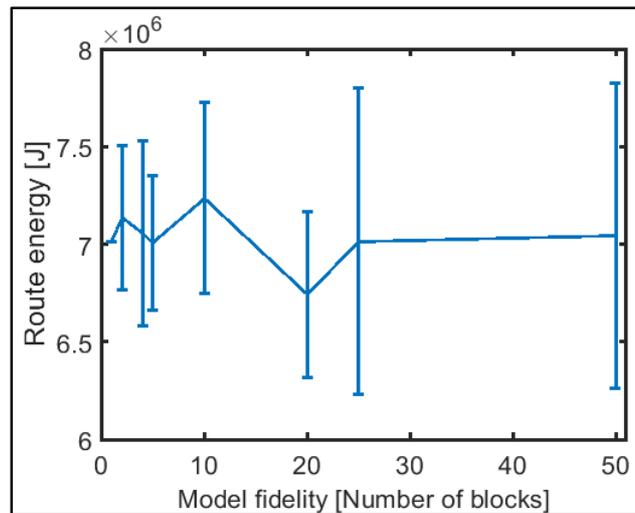


Figure 4.20: Variance in optimisation result for random starts

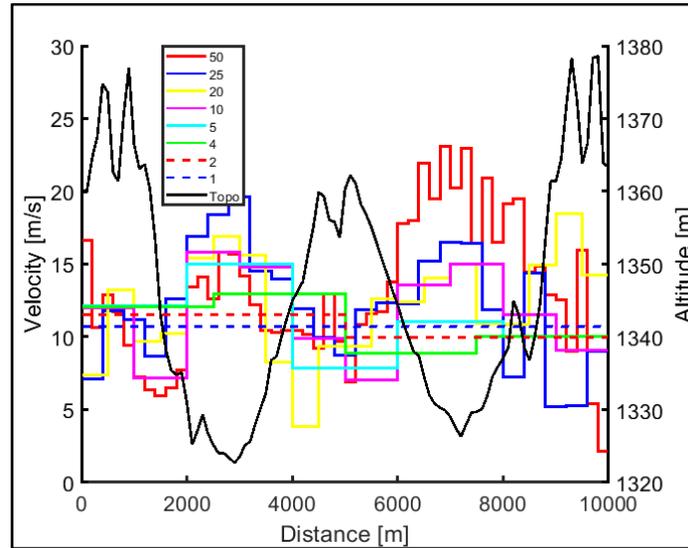


Figure 4.21: Optimal velocity vector for best of random initial guesses for route 4.1

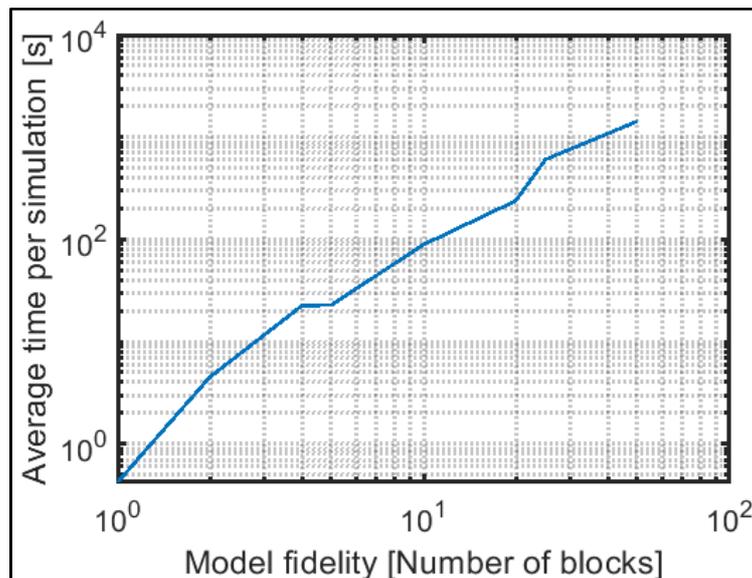


Figure 4.22: Simulation time vs. model fidelity for a random initial guess for route 4.1

Comparing the results from Figure 4.18, Table 4.3 and Table 4.4 shows that the constant initial guess method was unsuccessful in finding a good global minimum every time. Still, a random start method is not very efficient or robust either. However, it sometimes finds a value significantly lower while the best random start is sometimes worse than the constant initial guess. A better strategy still needs to be developed. A higher fidelity resulted in a lower overall route energy to be found, but the likelihood of finding a global minimum also reduced as the fidelity increased in that the standard deviation for the results increased with fidelity. A sweep of the velocity values for each fidelity case can be performed to find a global minimum, but the feasibility of this sweep becomes completely impractical as model fidelity increases. Using only random starts has the potential of yielding better results than the constant initial guess, but a large amount of simulation runs will be required to get to an optimal solution with

any reasonable confidence. The simulation time is a strong function of model fidelity, and it can quickly reach a point where optimisation becomes infeasible for a high model fidelity.

4.4.4. Fine sweep of values to find success rate

The model thus far has proven capable of finding some local minima. The task now is to obtain an idea of the likelihood of having found a true minimum. A sweep of initial values is performed for a model with a fidelity of 2, and a surface is plotted to show the optimal route energy values that were found. A broad sweep over a large velocity range with large increments of 1 m/s is performed first. The resulting surface plot of the results is shown in Figure 4.23 a.

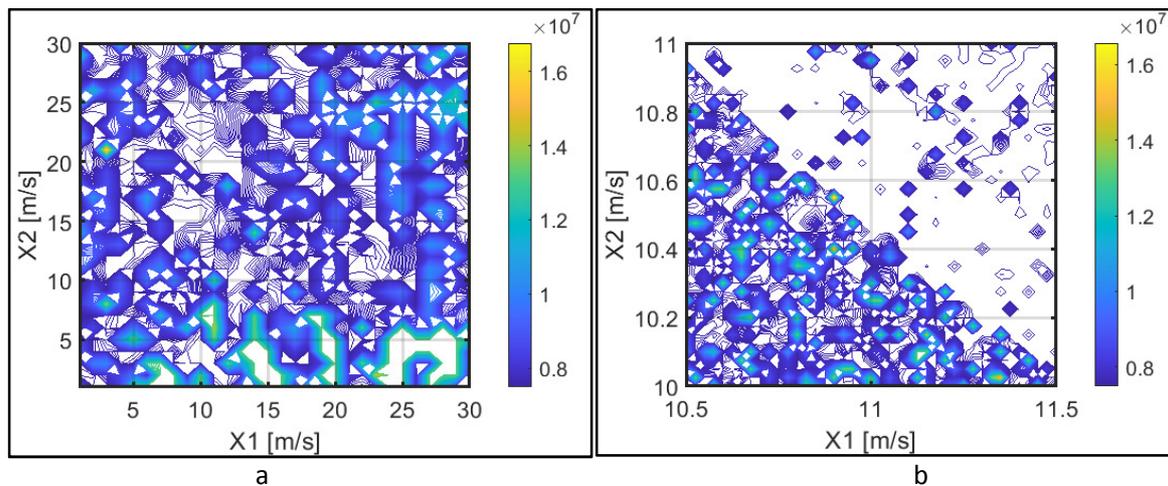


Figure 4.23: Optimal route 4.1 energy for swept initial values for X1 and X2. a) Swept from 1-30 m/s in increments of 1 m/s, b) Swept close to the optimal in increments of 0.025 m/s

The results of the course sweep showed that 64% of the optimal route energies obtained were within 10% of the lowest value, 53% were within 5% of the lowest value, and only 37% were within 2% of the lowest value.

Based on the results from Chapters 4.4.2 and 4.4.3 we have a good idea of where the optimal values are expected to lie, so a fine sweep of starting values is performed around that region. A finer resolution in the initial value close to the expected optimal showed an increased likelihood of finding the optimal value. The swept values of $X1 = 10.5$ to 11.5 and $X2 = 10.0$ to 11.0 in increments of 0.025 m/s yield the results shown in Figure 4.23 b. Here 72.3% are within 2% of the lowest value found. The best answer had $X1 = 11.24$ and $X2 = 10.24$ m/s respectively.

This shows that a finer variance in the initial value close to the optimal will result in more results of the optimisation lying close to the optimal. This allows us to make small random variations in the initial velocity vector if we know we are close to the optimal, such that we obtain a high chance of finding the global minimum. If we can find a way to obtain a good initial guess, we can apply slight variations to the values and run the optimisation several times to yield a global minimum that can be used with confidence.

4.4.5. Inverted topographic as an initial guess

It was seen in Chapters 4.2. and 4.3. that there exists an improvement in the velocity profile to travel a route, which yields a reduced overall energy consumption. By studying Figures 4.18 and 4.21 it is seen that the optimal velocity profile follows a trend which is similar to an inverted version of the route topographic profile. This can also be seen in the velocity profiles for the test routes presented in Chapter 4.1. An incline in the route resulted in a decrease in the velocity as the incline was driven up, and a decline resulted in an increase in velocity down the hill. This phenomenon makes intuitive sense, as this is the principle on which a rollercoaster operates. This is the most efficient way of travelling, converting between potential and kinetic energy directly without trying to store energy and incurring losses by efficiency effects. In this section, the development of a robust method of finding an appropriate initial guess for the optimisation is presented. If we assume an inverse relationship between the route topography and the velocity profile it allows us to generate a method to construct a sensible initial guess to the optimisation, which can yield more robust optimisation results and reduce the time to find a solution.

The topographic profile of a known route can now be modified and then used as the initial guess for the optimisation to study the benefits of using this approach, as opposed to constant valued or multiple random starts. To obtain the correct scale and vertical shift of the inverted topographic profile a quickly executable low-fidelity cost function may be optimised. A straight-line equation of the form $y = -m(x) + c$ is used. The minus sign is used to invert the profile and the constant, m , is used to scale the topographic plot $f(x)$, while c is used to perform a vertical shift, where the constants m and c are the two values optimised using the *fmincon* solver. The low fidelity of 2 allows this solver to only require a runtime of about 20 seconds to deliver an optimised, inverted, scaled and vertically shifted topographic profile that can be used as a good initial guess in the velocity vector optimiser, using most of the same scripts as for the previous optimisations performed.

The robustness of the optimised inverted topographic initial guess is investigated by incrementally adding more noise to the function values up to a point where it is almost comparable to a completely random initial guess. The final route energy is recorded for each run with the random noise added to the initial velocity vector's terms and a measure of success can be obtained from there. The optimisation with the new initial guess strategy is repeated 10 times for each noise intensity under each model fidelity. The random noise is added by using Eq.4.3. The intensity of the noise is scaled based on the *Noise_factor*. The *Noise_factor* is incrementally increased in the form of: 0.01 to 0.1, 0.5, 1, 2, 5, 10 and 16 respectively. The value stops at 16 as any value higher than this was found to have a high likelihood of going negative in the initial guess values, which will cause the program to crash due to it resulting in negative time steps, as mentioned previously. The subtracting of 0.5 adapts the normally distributed variable to be averaged around zero instead of 0.5, increasing or decreasing each term in the velocity vector by the scaled random amount.

$$\bar{V}_{initial} = \bar{V}_{opt} + Noise_factor \times (rand(1, length(\bar{V}_{opt}) - 0.5)) \quad (Eq.4.3)$$

Where $\bar{V}_{initial}$ is the initial guess vector used for the optimisation, and \bar{V}_{opt} is the optimised inverted topographic vector.

The model fidelity is increased in the following increments: 5, 10, 20, 25 and 50. The very low fidelity options are excluded as they lack the resolution to follow any route topographic curvature and do not come close to the optimal values from the higher fidelity simulations, as was seen in the preceding chapters. The following simulation results were obtained by applying this strategy to the topography of Route 4.1, a real-world test route. Figure 4.24 shows the 10 simulation results obtained for the final route energy with the various random noises added using the *Noise_factor*, for all fidelities simulated. A plot of the optimal route energy vs. *Noise_factor* as well as the standard deviation for the route energy solutions are given in Figure 4.25 a and b respectively. The mean and standard deviation of the results are also listed in Table 4.5, with the best solution energies presented in Table 4.6.

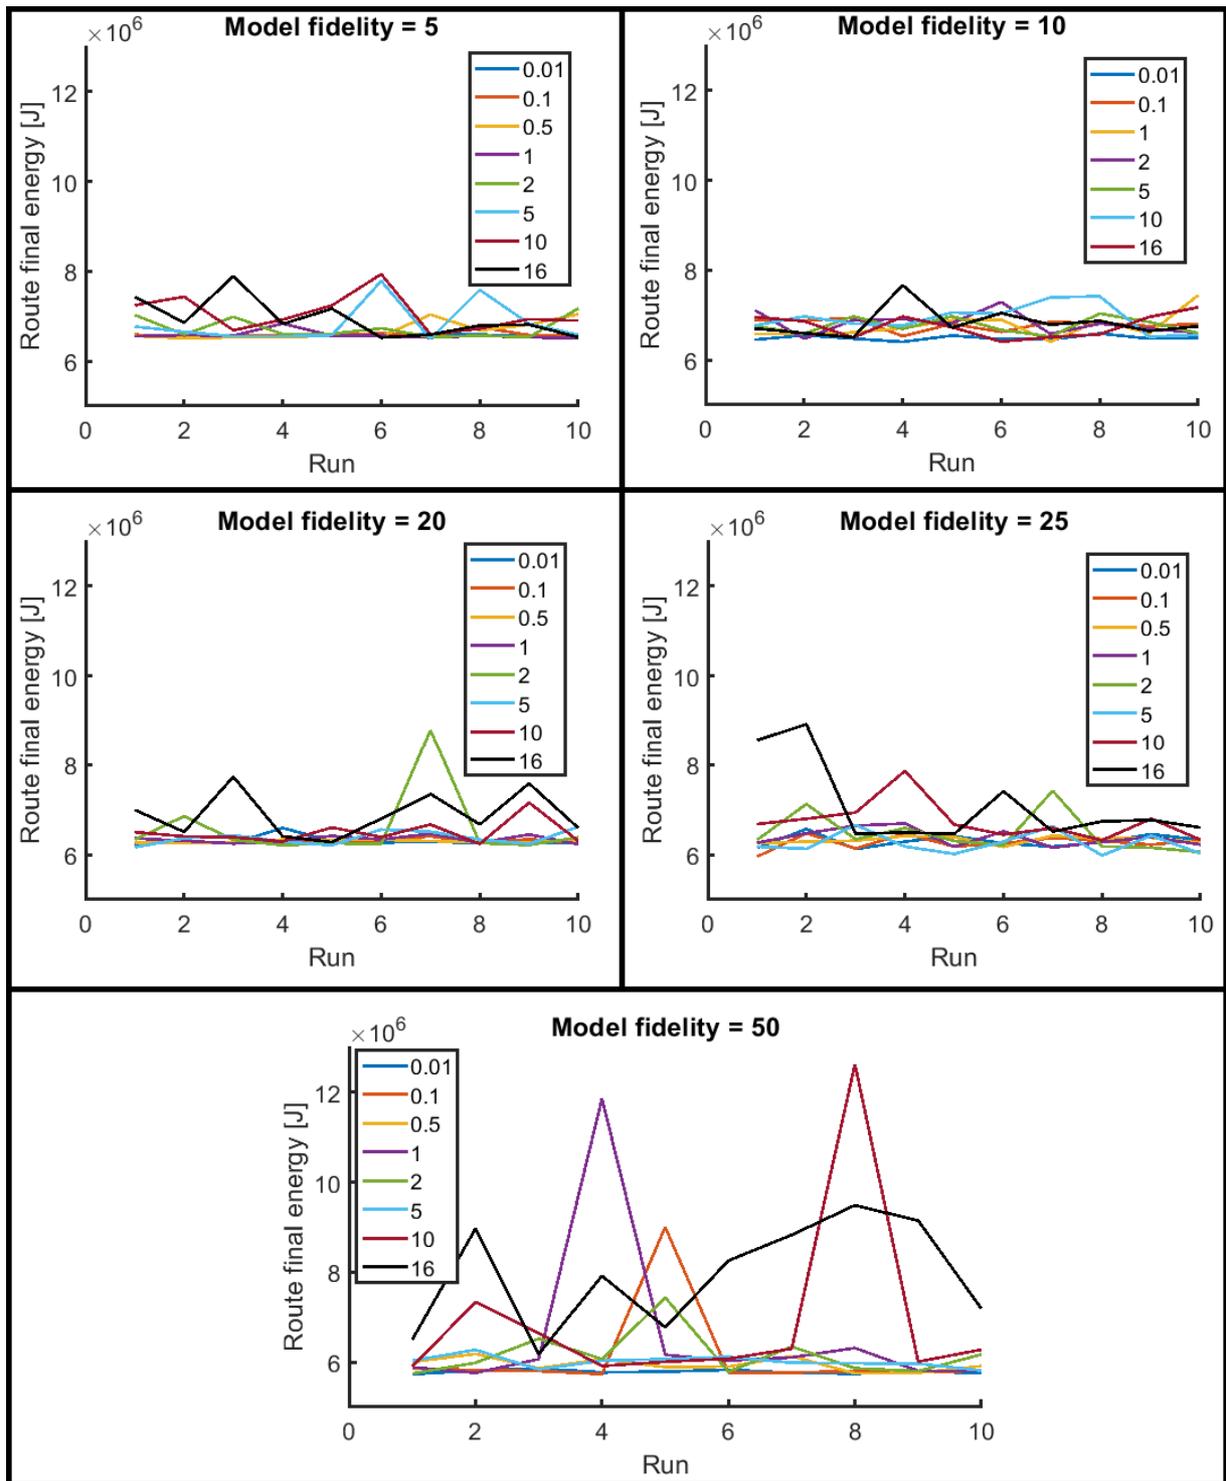


Figure 4.24: Route 4.1 simulations for varying noise factor and fidelity

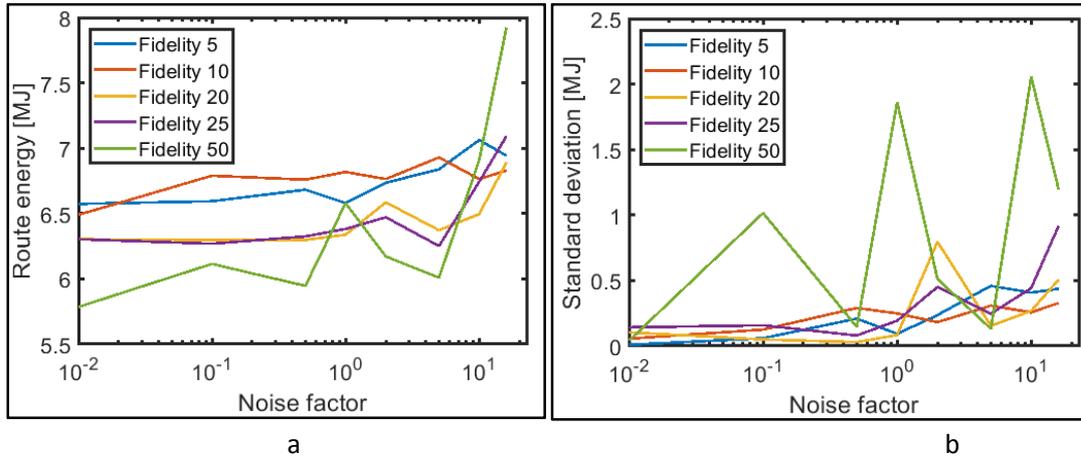


Figure 4.25: a) Optimal solutions for route 4.1 energy vs. *Noise_factor*, b) Standard deviation vs. *Noise_factor*

Route 4.1 energy mean and standard deviation vs. Noise factors & Fidelity (Best/Worst)																
Fidelity	0.01 (μ, σ)		0.1 (μ, σ)		0.5 (μ, σ)		1 (μ, σ)		2 (μ, σ)		5 (μ, σ)		10 (μ, σ)		16 (μ, σ)	
5	6.57	0.01	6.59	0.06	6.68	0.21	6.58	0.10	6.74	0.24	6.84	0.46	7.06	0.41	6.94	0.44
10	6.49	0.06	6.79	0.13	6.76	0.29	6.82	0.25	6.76	0.18	6.92	0.31	6.76	0.26	6.83	0.33
20	6.31	0.10	6.30	0.05	6.30	0.03	6.34	0.08	6.59	0.80	6.37	0.15	6.50	0.27	6.89	0.51
25	6.30	0.14	6.27	0.16	6.32	0.08	6.38	0.19	6.47	0.45	6.25	0.25	6.74	0.44	7.09	0.92
50	5.79	0.04	6.11	1.02	5.94	0.15	6.58	1.86	6.17	0.51	6.01	0.13	6.91	2.05	7.92	1.20

Table 4.5: Mean and standard deviation for optimal inverted topographic initial

Model fidelity	Best route 4.1 energy from the 10 runs [MJ]
5	6.50
10	6.40
20	6.16
25	5.95
50	5.73

Table 4.6: Best route 4.1 energy results of complete optimisation

From Table 4.6 it is seen that an increase in model fidelity results in a decrease in energy consumed. It is noted from Figure 4.25 that the chance of finding the optimal solution decreases as the noise factor increases, for all fidelities. An increased fidelity should result in a better solution, but at a significant simulation time increase, as was shown in Chapters 4.2. and 4.3. Figure 4.25 b also proves that the initial guess is successful in guiding the optimiser to a good solution to the problem, as the solution becomes worse when deviating from that initial guess. A more robust method to find a global minimum would be to run the simulation a couple of times with a small noise factor on the initial guess, to increase the likelihood of finding a global minimum. Though time-consuming, this method proved to yield the best solutions to the optimisation problem usually requiring only two or three simulations to be performed as the low noise factor data rarely diverges.

The use of the optimised inverted topographic initial velocity vector has great benefit to finding an approximate solution in a significantly reduced timeframe. When using the optimal inverted topographic

profile as the velocity vector for the vehicle, a route energy of 6.67 MJ is found. When optimising the velocity profile for a model fidelity of 10, using the optimal inverted topographic as the initial guess, a route energy of 6.40 MJ is found, which means the initial guess was within 4% of the complete optimisation solution, taking less than a minute to run, instead of 8 minutes for one complete *fmincon* solution for a fidelity of 10. When a model fidelity of 50 is used, the best solution yielded an energy value of 5.73 MJ, taking 2 to 4 hours to yield one solution. The improvement there is $(6.67-5.73)/6.67 = 14\%$ from the inverted topographic initial guess. The drawback in reality is that the optimisation needs to be run a couple of times before the answer can be trusted with confidence. The worst route energy of the 10 simulations performed with a fidelity of 50 and noise factor of 0.01 was 5.86 MJ, $(6.67-5.86)/6.67 = 12\%$ better than the initial guess value on its own.

A major advantage of using the optimised inverted topographic velocity profile as the optimal is that it has the same resolution as the topographic profile specified by the available database. For longer routes, the inverted topographic will still have good resolution, whereas the complete optimisation approach is limited to the practical upper limit for the model fidelity, rather than the length of the route. This makes the inverted topographic approach insensitive to route topographic complexity as well. Another advantage is that the route length need not be a convenient length, like the one of route 4.1, which can be divided equally and easily. Distance blocks of varying lengths can be applied to reduce this problem but add complexity to finding a solution.

The velocity profile for the optimised inverted topographic profile and that of the complete optimisation solution are presented in Figure 4.26a. Figure 4.26b shows the route energy consumption for the two strategies. One limitation of the inverted topographic profile approach is that it is not able to deal with local speed limits elegantly, which may reduce its applicability in some situations. The completion time for the two velocity profiles is almost identical at 922s for the inverted topographic and 919s for the complete optimisation solution respectively.

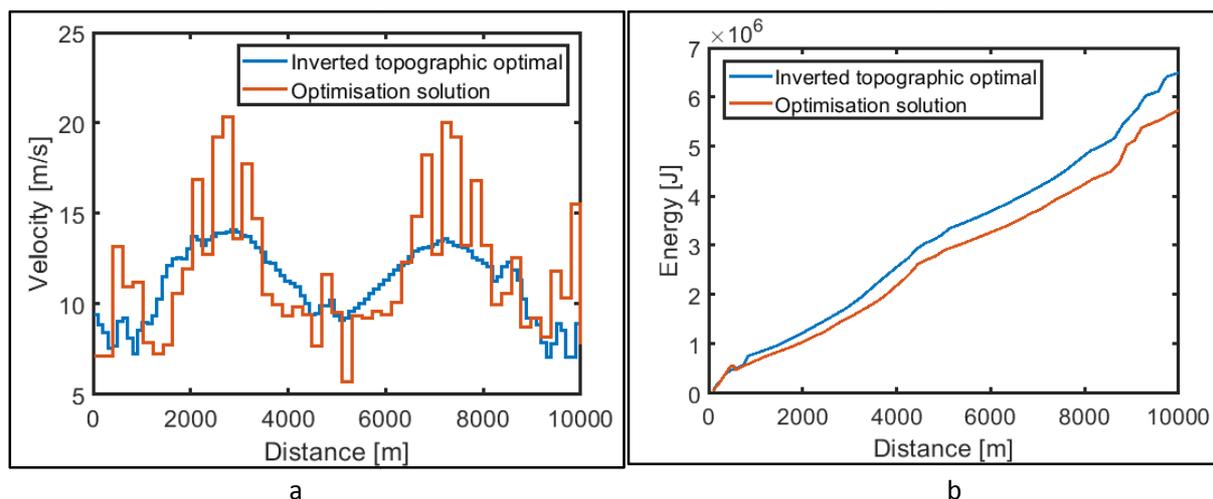


Figure 4.26: Inverted topographic vs. complete optimal solution. a) Velocity profiles, b) Energy

4.4.6. The effect of regeneration

The ability of an electric or hybrid vehicle to store energy in a usable form can greatly benefit the reduction of overall energy consumption. Hybrid electric vehicles have been developed with this as their main focus. They can store energy when excess energy is available, as in the case when going downhill or when reducing velocity. The efficiency at which the vehicle can convert the potential and kinetic energy to stored energy can be termed as the regeneration efficiency. This is typically a function of the electric motor, motor driver, power converter and battery efficiencies. Many researchers assume a constant value for the regeneration efficiency, though some report that it diminishes at very low speeds as the motors have very low efficiencies in that range. Fully electric vehicles, if equipped with the functionality, can also store energy in their onboard batteries for later use. The typical combined efficiency of all of the factors listed above is stated to be around 40% [42].

With a model available that can find an optimal solution to the velocity profile, we can use this model to perform velocity profile optimisations which consider the effect of regeneration in a vehicle. The test routes will be simulated first, for zero and 40% regeneration efficiencies, to determine whether the program yields sensible results, after which a thorough regeneration study can be performed on route 4.1. Figure 4.27 a shows the optimal velocity profiles for all of the constant incline test routes for zero regeneration, and Figure 4.27 b for 40% regeneration. Figure 4.27 c shows the optimal velocity profiles for the 5% Up-Down and 5% Down-Up routes with 0% and 40% regeneration efficiency. Figure 4.27 d shows the optimal velocity profiles for the 7% Up-Down-Up-Down and 7% Down-Up-Down-Up routes for 0% and 40% regeneration efficiencies.

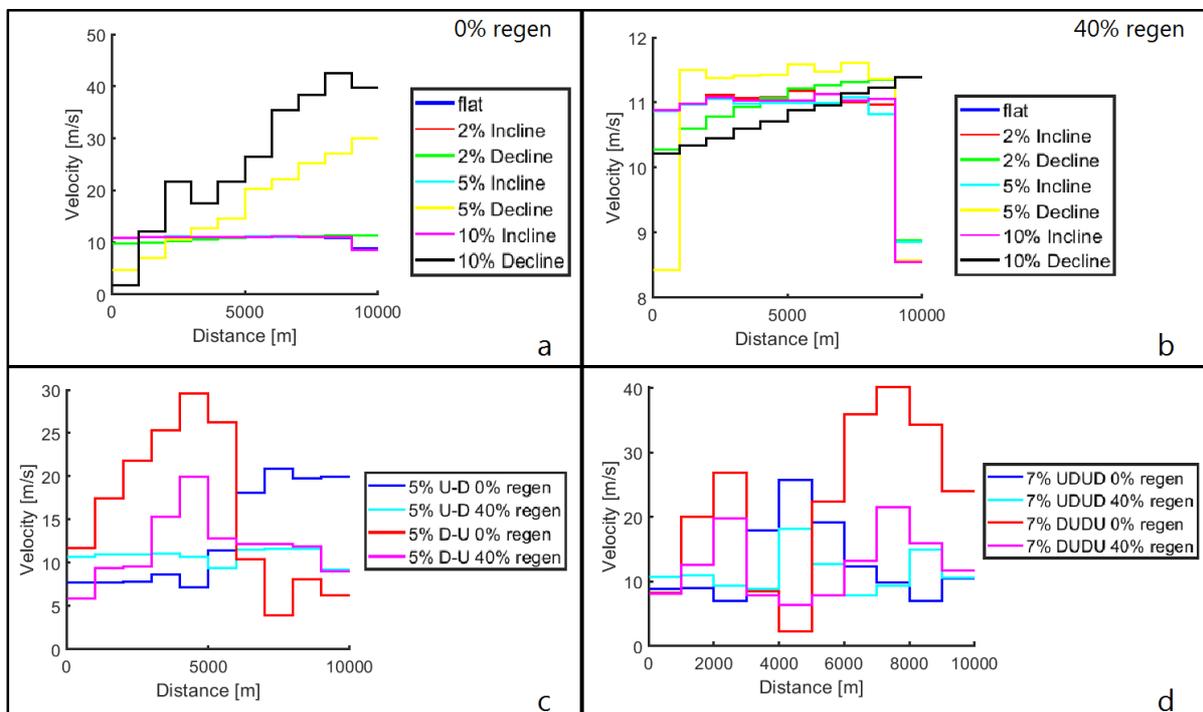


Figure 4.27: Optimal velocity profiles for the test routes with 0% or 40% regeneration efficiency

It is seen from Figure 4.27 that regeneration is not very helpful when the route has a flat or constant incline profile. For routes with constant declines, the ability to store energy in the form of regeneration is beneficial in that the vehicle can complete the route with a more constant velocity profile and can even end up with a final energy that is negative, meaning that the storage system contains more energy after the route than before. This applied the 5% and 10% downhill routes, where the final energy for zero regeneration was close to zero, and was -2.47 MJ and -6.38 MJ for the 5% and 10% decline routes at a regeneration efficiency of 40%. From this simulation, it is seen that an increased regeneration ability leads to more constant velocity profiles, not requiring the vehicle to try and store energy by driving the route like a rollercoaster.

A sweep was performed using constant regeneration efficiencies for a model fidelity of 50 when travelling route 4.1. The results are presented in Figure 4.28. The final route energy amount for each regeneration efficiency is shown in the legend.

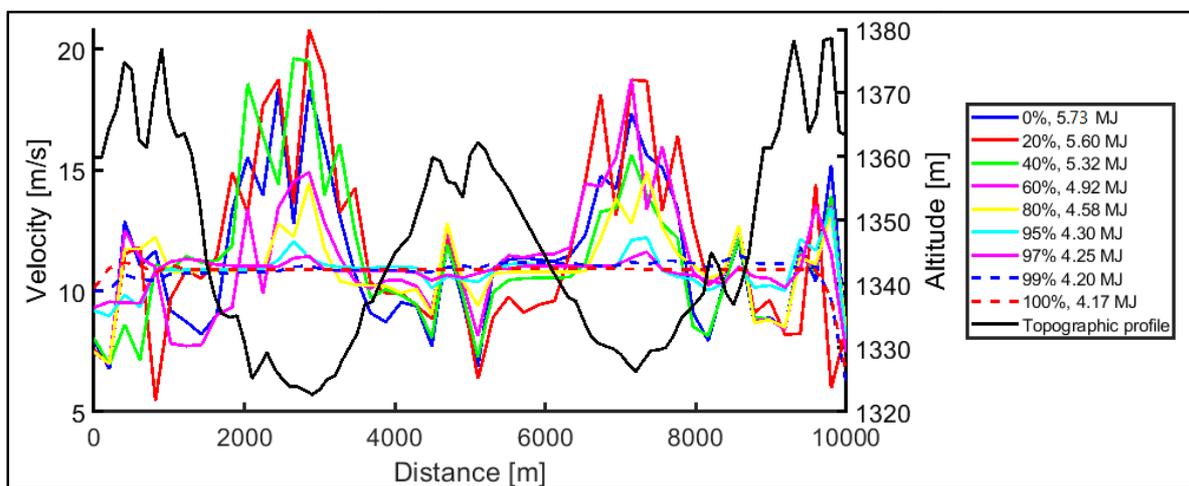


Figure 4.28: Optimal velocity profiles for various regeneration efficiencies, model fidelity of 50, for Route 4.1

From the results presented in Figure 4.28, it is seen that regeneration can greatly reduce the overall energy consumption of a vehicle. If energy can be stored with 100% efficiency it also reduces the need to speed up when going downhill and reduce speed when going uphill, as there is no need to store energy between kinetic/potential form instead of sending it to the onboard storage system. It is seen that at efficiencies higher than 80% the varying of speed is barely required and the vehicle may simply travel at the lowest possible speed to complete the route in the maximum allowable time. With current technology regenerative efficiencies are only around 40%, which means that velocity profile optimisation still poses a significant benefit to energy saving in a vehicle. Hybrid vehicles have the ability to not only store energy but also allow running the IC motor in a more efficient range of operation. This effect was not investigated in this study, as it depends on many vehicle and engine-specific characteristics, which lie outside the scope of this study.

From the simulation performed it is seen that the regeneration efficiency acts to simply reduce the peaks and does not change the overall characteristic of the optimal velocity profile. This is of benefit to

the approach of simply using the optimised inverted topographic profile as it will simply reduce the magnitude of the scale factor, m , from Chapter 4.5, with all other aspects remaining the same.

An interesting aspect to note is that regeneration efficiency has a small effect if the topographic characteristics of a route are constant. When travelling on a route with a constant incline at a constant velocity there is no excess energy to store for later use, with the only case being a decline steep enough to overcome the frictional losses and that causes the vehicle to speed up past the speed limit by gravity alone, in which case there exists excess energy that may be stored. If no regeneration is available all of this excess energy will normally have to be converted to heat in the form of friction brakes. If the vehicle can store energy, the excess energy can be stored for later use with the only penalty being the regeneration efficiency.

An interesting consideration is to compare the optimal route energy for a fidelity of 10 to that of 50, as presented in Table 4.6 in Chapter 4.4.5. The best route energy for a fidelity of 10 was 6.50 MJ, with the best for a fidelity of 50 being 5.73 MJ. The improvement here is approximately 12%, going from an optimal route velocity profile to one with better resolution. The results for the various regeneration efficiencies presented above showed that the route energy for a realistic regeneration efficiency of 40% is only $(5.73-5.32)/5.73 * 100 = 7.2\%$ lower than that of a vehicle without regeneration. This emphasizes that driver behaviour and following better velocity profiles can be of greater benefit than regeneration on its own. The ideal of course is to have both.

Table 4.7 lists the route energies for the optimal inverted velocity profile and that of the complete optimisation solution, both for a sweep of regeneration efficiencies from zero to 1 (with 1 being 100%).

Model fidelity	Regen eff	Optimal Inverted topographic final energy (MJ)	Optimal Route final energy	Average simulation time (s)	% Error: Opt/inv
10	0	6.7728	6.5902	376	2.7
	0.2	6.2894	6.0617		3.6
	0.4	5.7524	5.7378		0.3
	0.6	5.2707	5.2399		0.6
	0.8	4.7428	4.731		0.2
	1	4.2149	4.1954		0.5
20	0	6.7728	6.5783	1287	2.9
	0.2	6.2894	5.9041		6.1
	0.4	5.7524	5.5927		2.8
	0.6	5.2707	5.1522		2.3
	0.8	4.7428	4.6789		1.3
	1	4.2149	4.1838		0.7
25	0	6.7728	5.9704	2320	11.8
	0.2	6.2894	5.7855		8.0
	0.4	5.7524	5.3041		7.8
	0.6	5.2707	4.9757		5.6
	0.8	4.7428	4.6042		2.9
	1	4.2149	4.1791		0.8
50	0	6.4993	5.7305	7250	13.4
	0.2	6.2894	5.6033		10.9
	0.4	5.7524	5.3239		7.4
	0.6	5.2707	4.9234		6.6
	0.8	4.7428	4.5753		3.5
	1	4.2149	4.169		1.1

Table 4.7: The effect of regeneration on route 4.1 energy consumption

It was seen that, as the regeneration efficiency increases, the optimal inverted topographic velocity profile used for the initial guess would approach the same route final energy as the fully optimised solution. This shows that, if the vehicle can implement regeneration with a realistic 40% efficiency, that the optimised inverted topographic velocity profile may be used directly without any significant penalty. This is of great benefit to the computational burden to find an optimal velocity profile as it runs more than 300 times faster than the full optimiser for higher model fidelities, as was shown in the simulation times of Table 4.7.

4.4.7. The effect of mass changes

To investigate the effect that vehicle mass has on the velocity profile a realistic sweep of the expected mass range is performed and the optimal velocity profiles are compared. Three optimisation solutions are performed for each mass tested to ensure that a good minimum is consistently found. The route

energies for the three passes of the route with the same mass were all within 1% of each other for all three runs of the same mass. Due to the linear relationship between vehicle mass and route energy, as was shown in Chapter 3.3.1. the velocity profile itself was not affected by the vehicle mass, only the final route energy. The sweep of mass was performed from 2000 kg up to 4000 kg, in increments of 500 kg and the optimal velocity profiles for all of those are shown in Figure 4.29a. The route energy consumption is shown in Figure 4.29b. It can be seen that, even though the vehicle mass was varied significantly, the resultant optimal velocity profile remained fairly consistent for all masses for the route simulated. It is expected that the mass should not greatly influence the optimal velocity profile, seeing as the mass-dependent terms in the energy equation are not greatly influenced by the velocity, as was presented in Chapter 4.3.3.1. The energy profiles scale linearly with mass, as was shown in chapter 4.3.3.1.

The significance of this is that the optimal velocity profile is insensitive to vehicle mass, meaning that real-time updating of the velocity profile need not be carried out due to changes in vehicle mass that occur along the route. The mass still has a significant effect on the route energy consumption and should not be discarded from the study, which is the whole purpose of the greater study in which this velocity profile optimiser plays a part. Keeping track of mass changes will allow better estimation of energy consumption, which will aid in increasing the accuracy of remaining range estimation in vehicles.

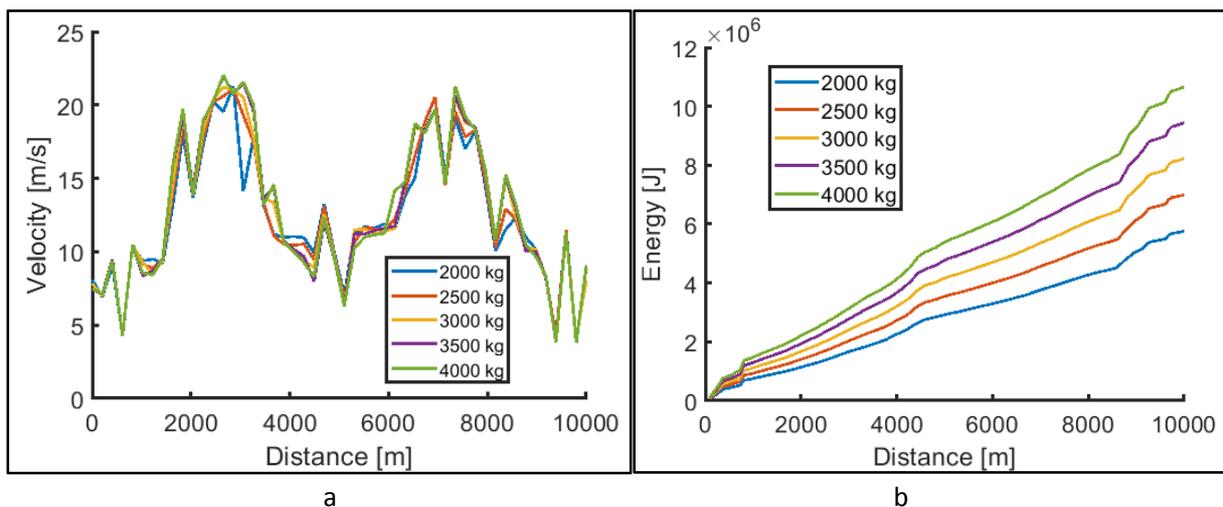


Figure 4.29: Varying mass on route 4.1. a) Optimal velocity profiles, b) Route energy

Seeing that the optimal velocity profile is unaffected by vehicle mass it means that the optimal inverted topographic method still has application here, as it is not necessary to change the profile if the mass changes. This has significant simulation time benefits to making this approach feasible for real-time execution.

4.5. Conclusions

In this thesis, a method is proposed by which the initial guess to a velocity optimisation can be improved to increase the likelihood of finding a global minimum in route energy consumption. A first-principles model was constructed that can estimate the energy consumption as a function of the route topography

and length, the vehicle type, the load and the velocity profile. A *fmincon* local solver was used to try to find the minimum of the cost function for the route energy based on finding a velocity profile that will meet the logistical requirements to complete a route in the prescribed time, but using the least amount of energy when doing so. A summary of the steps performed in this study is presented in Figure 4.30.

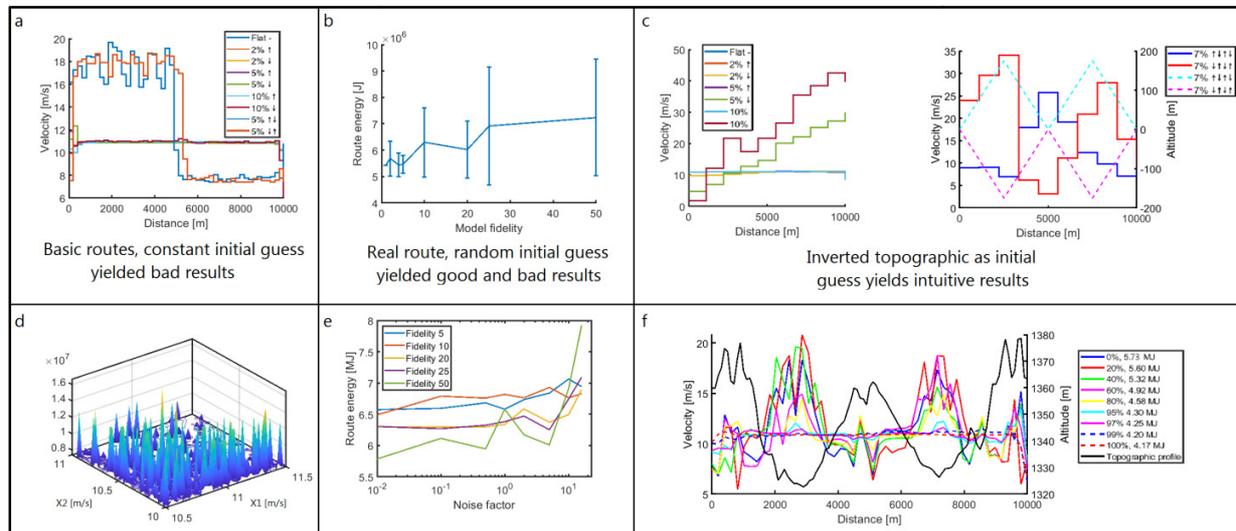


Figure 4.30: Summary of steps to find the optimal velocity profile of a vehicle on a known route

An in-depth study was performed to find a correlation between the initial guess and the final optimisation solution, as having an initial guess closer to the solution allows for finding a better solution than what can be done with bad initial guesses. A constant initial guess proved to force the solution to have constant regions as well, which did not make sense for the test routes evaluated, as shown in Figure 4.30a. Random initial guesses proved to sometimes offer good solutions, but often bad solutions as well, making the use of random starts impractical out of a time constraint point of view, as shown by Figure 4.30b.

It was found that there exists a correlation between the route topography and that of the optimal velocity profile in that the optimal velocity profile looks to be closely related to an inverted version of the route topography. A simple to implement low fidelity cost function was optimised that can deliver a good initial guess to the optimisation within a timeframe of less than a minute. This initial guess allowed the *fmincon* local solver to consistently find a good minimum solution to the optimisation problem, as shown in Figure 4.30c. It was found that this inverted topographic method sometimes gave solutions that were not representative of the global minimum, and that slight random variations in the element values of the initial guess proved to find a good solution 72% of the time, as is shown by Figure 4.30d. This strategy proved capable of consistently yielding good solutions to the optimisation problem, as any deviation from this method usually resulted in worse solutions, as shown in Figure 4.30e. With a method that has proven itself capable of finding good solutions the effect of regeneration on the optimal velocity profile can be studied and it was found that regeneration reduces the need to exploit the topography to save energy, in effect riding the route like a rollercoaster.

By studying the effect of model fidelity on the solution to the optimisation problem it was found that higher fidelities resulted in better solutions, at the cost of significantly increased simulation times. A trade-off might be required when simulation time has practical limitations.

It was noted that the initial guess method that was developed in this study can give an initial guess that may be sufficient to be used as a rough estimate of the route energy on its own without further optimisations required. The advantage of this method is that it reduces simulation time from a couple of hours down to less than a minute, making real-time application of this method more feasible. The excess energy consumed by using the initial guess on its own was found to be less than 15% when compared to the highest fidelity complete optimal solution in this study, for no regeneration ability available in the system. When regeneration is a possibility the error between the initial guess and the final solution is reduced to 7% for a realistic regeneration efficiency of 40%.

It was found that mass has a direct influence on the total energy consumption, but not on the optimal velocity profile for a given route, which means that the optimal velocity profile need not be adapted when the vehicle mass changes along its route. The only exception to this statement is that there exists a practical and safe limitation to the acceleration that a heavily loaded vehicle may undergo and care should be taken to not exceed these limits, especially when unrealistic time constraints are being imposed on the route. The results showed that the optimal velocity profile is simply a function of the route topography and the regeneration efficiency if the same vehicle is used.

An optimal route velocity profile proved capable of significant savings in energy consumed when traversing a route, even when the logistical constraint on time is constant. For route 4.1 the vehicle consumed a total energy of 8.5 MJ in a real-world driving style. When optimising the velocity profile for that route the total energy could be reduced down to 5.7 MJ while still completing the route in a similar time. This is a 33% reduction in energy, which proves that driving style can greatly improve the energy consumption and cost of operation of a vehicle, be it electric, hybrid or IC as this optimisation is based on the route and not the vehicle.

4.6. Future work

The optimisation strategy proposed in this work proved to be successful in finding a velocity profile for a route that consumes the lowest energy. There are however elements that could be improved upon in future studies. Presented here is a list of aspects that can be studied further in future.

It would be good to investigate the effect of speed limits on the optimal velocity profile. In some applications, there might be specific zones where the maximum and minimum allowable speeds need to be constrained. These might include normal speed limits on public roads and even minimum speeds for freeways etc. To apply speed limits to the model presented in this study would be easily achieved by simply adding constraints to the optimisation problem. Applying speed limits to the inverted topographic method as the final answer is not feasible due to the low model fidelity not being designed to gracefully account for these.

The strategy was developed and tested on some standardised test routes and one real-world route. It would be beneficial to perform more simulations on real-world routes.

For heavily loaded commercial vehicles the ability to quickly change velocity is limited by braking and engine power. It would thus make sense to be able to limit the rate of velocity change to ensure the vehicle is capable of safely following the optimal velocity profile. It should be noted that the model fidelity plays a role here, as a low fidelity for a long route can cause very slow acceleration profiles.

To speed up the process of finding an optimal inverted topographic profile the system can be characterised and replaced by a look-up table, which will allow real-time application of the model for optimisation of velocity profiles as the vehicle is travelling. Alternatively, a surrogate model of the entire first principles system can be created to reduce the computational burdens.

5. Real-time route velocity profile optimisation

This chapter summarises the adaptations to the optimisation strategy that facilitates operation in real-time.

5.1. Background

In Chapter 4 a method was proposed that can find a global minimum for the route energy as a function of the velocity profile with good confidence. The drawback of that method is that a complete *fmincon* simulation run can take several hours to complete, making real-time application of this method impossible. A method was developed to enable finding a good initial guess vector to be used in the complete optimisation problem. It was found that this initial guess method can be executed in a significantly reduced timeframe due to the significantly lower fidelity that the model possesses. The accuracy of the low fidelity model is not as good as the full optimisation, but in most cases reached final route energy consumption values within 10% of the complete optimal solution. If regeneration is available in the vehicle this model proved to improve in accuracy as well, with typical errors being to the order of 7% as compared to the maximum fidelity models that were optimised for a realistic regenerative efficiency of 40%. Higher regenerative efficiencies are not widely attainable at this stage.

5.2. Method

Due to the low fidelity model only having two variables, the scale and shift parameters, it makes it possible to easily perform a sweep of the input variables such that these two parameters may provide a lookup table that can be used to quickly locate an estimate to the optimal velocity profile in a real-time application. From Chapter 4.4 it was seen that the only variables that influence the scale and shift parameters are the regeneration efficiency and the maximum completion time. These two variables are swept over a broad range of values that are expected to exceed the realistic values that will be encountered in real-time applications, and the resultant *m* and *c* values are stored in a lookup table for fast execution. Running the lookup table takes around 0.7 ms to yield an estimate of the optimal velocity vector. Figure 5.1 a and b show surface plots for the scale and shift parameters *m* and *c* for route 4.1 respectively. Though not extremely smooth, these can effectively be used to yield estimates of the optimal velocity profiles. The peaks in one parameter correspond to peaks in the other, so interpolation should still yield usable results. It is now possible to perform real-time velocity optimisation for a vehicle travelling a known route with a known mass vector with a good estimation of the energy required to complete the route with a combination of the real-time and proposed optimal velocity profile. A schematic diagram showing the procedure for the real-time optimisation is given in Figure 5.2.

In Figure 5.1 a and b there is an apparent discontinuity in the *m* and *c* parameters in the zero regeneration efficiency line at the back. This discontinuity lies close to the realistic values for the maximum time for the route and will lead to noise in the resulting optimal profile. The matrices for the scale and shift parameters were filtered using a built-in Matlab smooth filter with the 'rloess' setting applied. This setting has the advantage that it is insensitive to local outliers. Figure 5.2 shows the zero regeneration line of the *m* and *c* parameters with the original and the smoothed data for route 4.1. From this figure, it can be seen that the 'rloess' filter setting is able to ignore outliers and follow the

trend of the data almost exactly. The smoothed surface plots for the scale and shift parameters are shown in Figure 5.3 for comparison purposes.

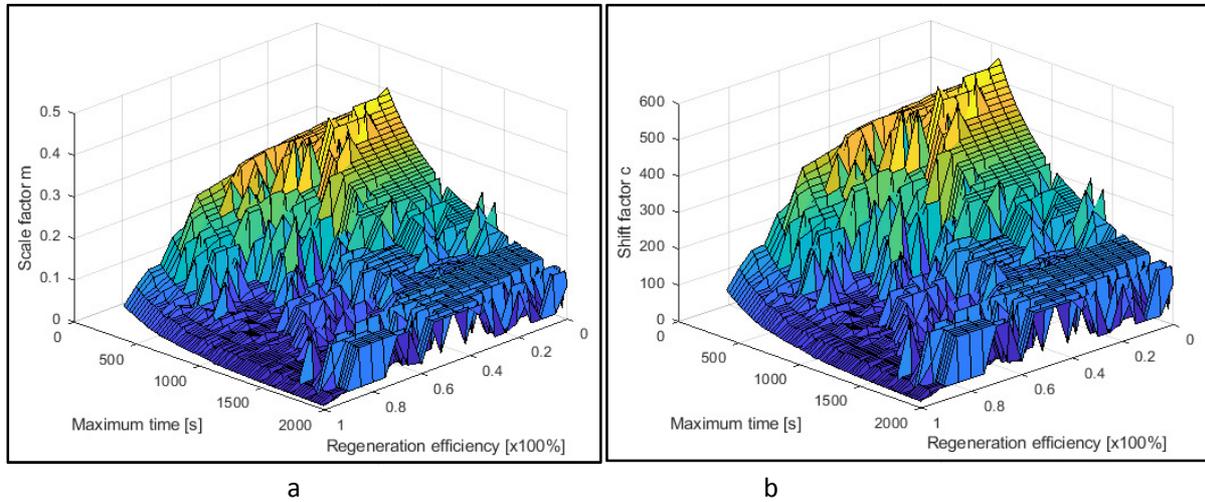


Figure 5.1: Scale (a) and shift (b) parameter surface plot vs. regeneration efficiency & maximum completion time for route 4.1

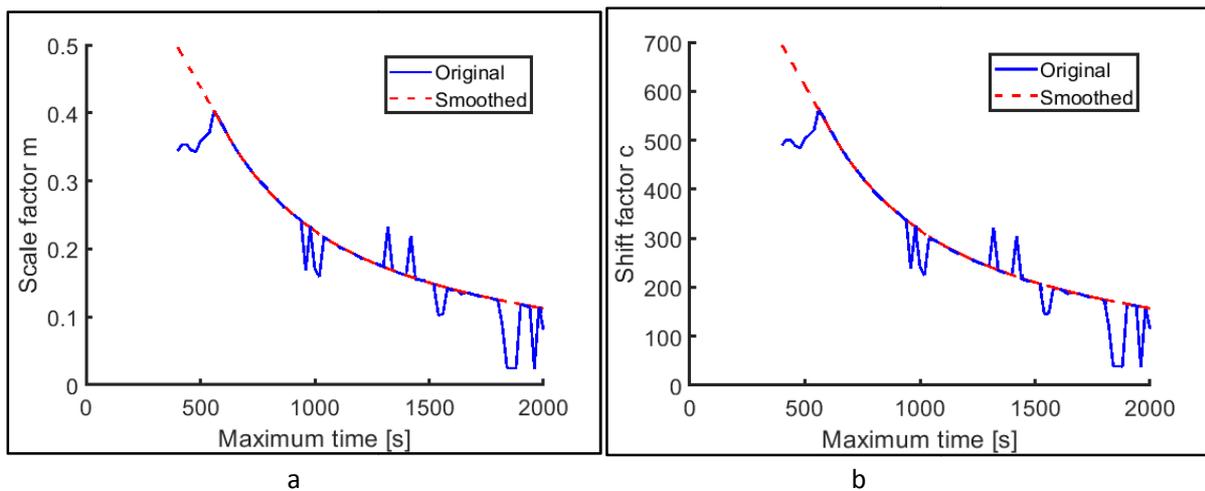


Figure 5.2: Smoothed scale a) and shift b) parameters for the zero regeneration line

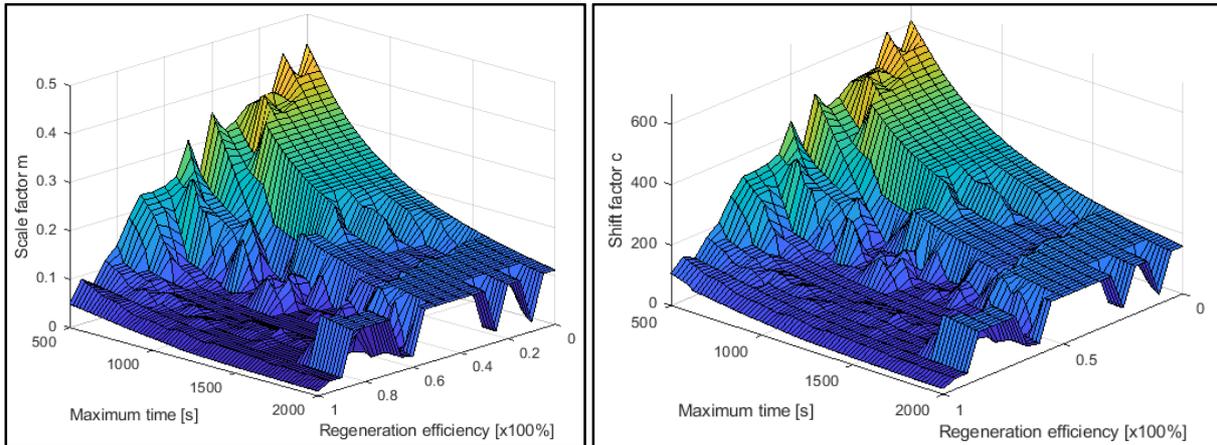


Figure 5.3: Smoothed a)scale and b)shift parameter surface plots for a sweep of regeneration efficiency and maximum time for route 4.1

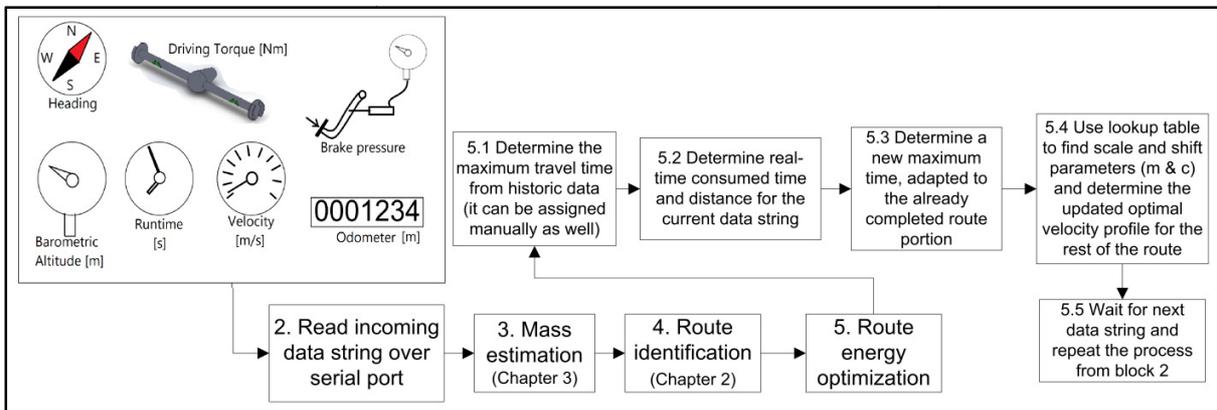


Figure 5.4: Real-time program schematic

In the schematic diagram of Figure 5.4 the mass estimation part, block 3, was explained in detail in Chapter 3 and route identification part, block 4, was explained in Chapter 2. One of the main contributing factors in route energy consumption is the time constraint, with the optimisation being dependent on this value. The total travel time may be obtained from historic data or can be input manually based on logistical constraints. From here the completed distance and time for the route may be obtained from the real-time data that is being streamed in. If the amount of distance and time consumed are known it is possible to calculate the amount of time left by subtracting the consumed time from the maximum time, with a similar process followed for the distance. An average velocity for the completion of the route can be determined by taking the remaining distance and dividing it by the remaining time. This average velocity is now used to determine a hypothetical value for the maximum time used to determine the optimal velocity profile that should be followed for the rest of the route. This process is repeated for each new 100-metre distance increment streamed in until the route is completed.

5.3. Real-time optimal velocity profiles

To verify whether the real-time optimal program code runs successfully, the optimal result will be input as the "real-time" data and it is then expected that the real-time optimal should be the same as the input data. The results of this simulation are shown in Figure 5.5 a. The black line represents the original optimal velocity profile, which is used as the input to the real-time optimiser, and the blue lines represent the optimal route profiles for the route in real-time for every 100 m increment as the route is being completed. Figure 5.5a used the m and c parameter lookup data from Figure 5.1 a and b. It is noted that the approximation method does introduce some noise, as is seen by the blue lines for every one of the 100 data points not lying on top of the original optimal one. The smoothed m and c parameters were used to yield the real-time optimal tracing shown in Figure 5.5 b. It is seen that the amount of noise generated by using the smoothed data is significantly less than the original data. The result presented in Figure 5.5 confirms that the program can deliver a sensible result, following the optimal profile when the optimal profile was input as the real-time data, though it may be imprecise at some locations due to approximations made. This approach will now be used to perform real-time optimisations on a route with real-world data being streamed in.

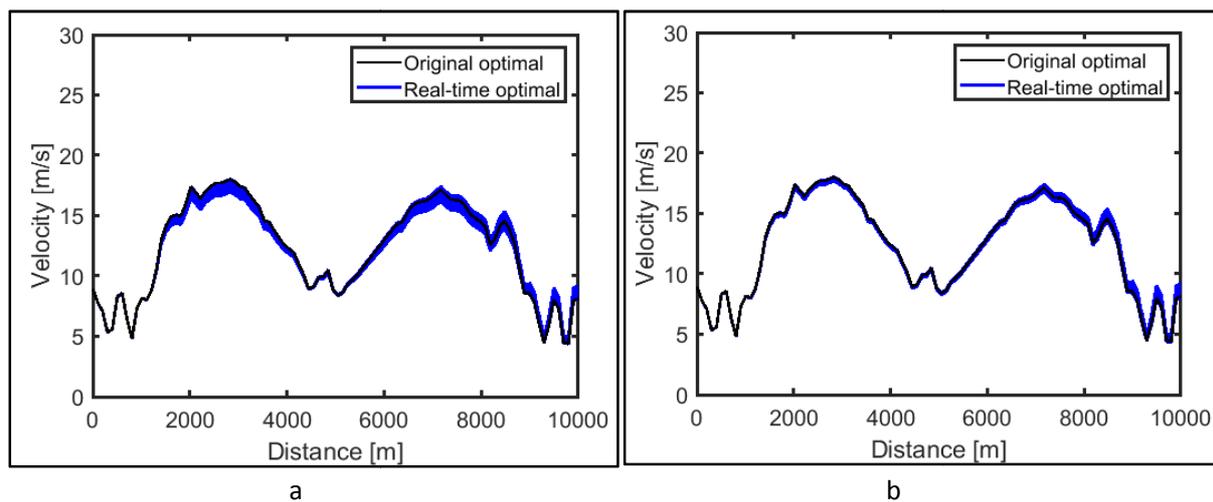


Figure 5.5: Optimal route velocity profile for route 4.1 using the optimal for that route as input. a) original m and c , b) smoothed m and c parameters

The real-time updating of the velocity profile for a vehicle travelling on route 4.1 is shown in Figure 5.6. This route was travelled in an uninformed manner, with velocity optimisation not influencing the real-time data during the completion of the route. The plot thus shows the difference between a real-world real-time uninformed travel, indicated by the red line, and the real-time updated optimal for the rest of the route, which is indicated by the blue line. The faint black line indicates the original approximation to the optimal for the whole route, as a reference. With the route being identified early, the system can propose an improved velocity profile to be taken for the rest of the route and update this velocity profile based on what has happened in the real-time travel data. It is seen that, as the vehicle drives slower than the optimal, the optimal for the rest of the route has an increased form to compensate (Figure 5.6 a and b). As the vehicle drives faster than required, the optimal profile for the remainder is reduced to compensate (Figure 5.6 c and d). This updating occurs in real-time. The way this method is

set up allows for graceful accommodation of moderate levels of unplanned events, like traffic, on a route, though speed limits and vehicle power limitations may cause the final completion time to become infeasible if there is not enough distance to make up for lost time. Similarly, the comparative profiles for routes 4.2 through 4.4 are indicated in Figure 5.7 a-c respectively.

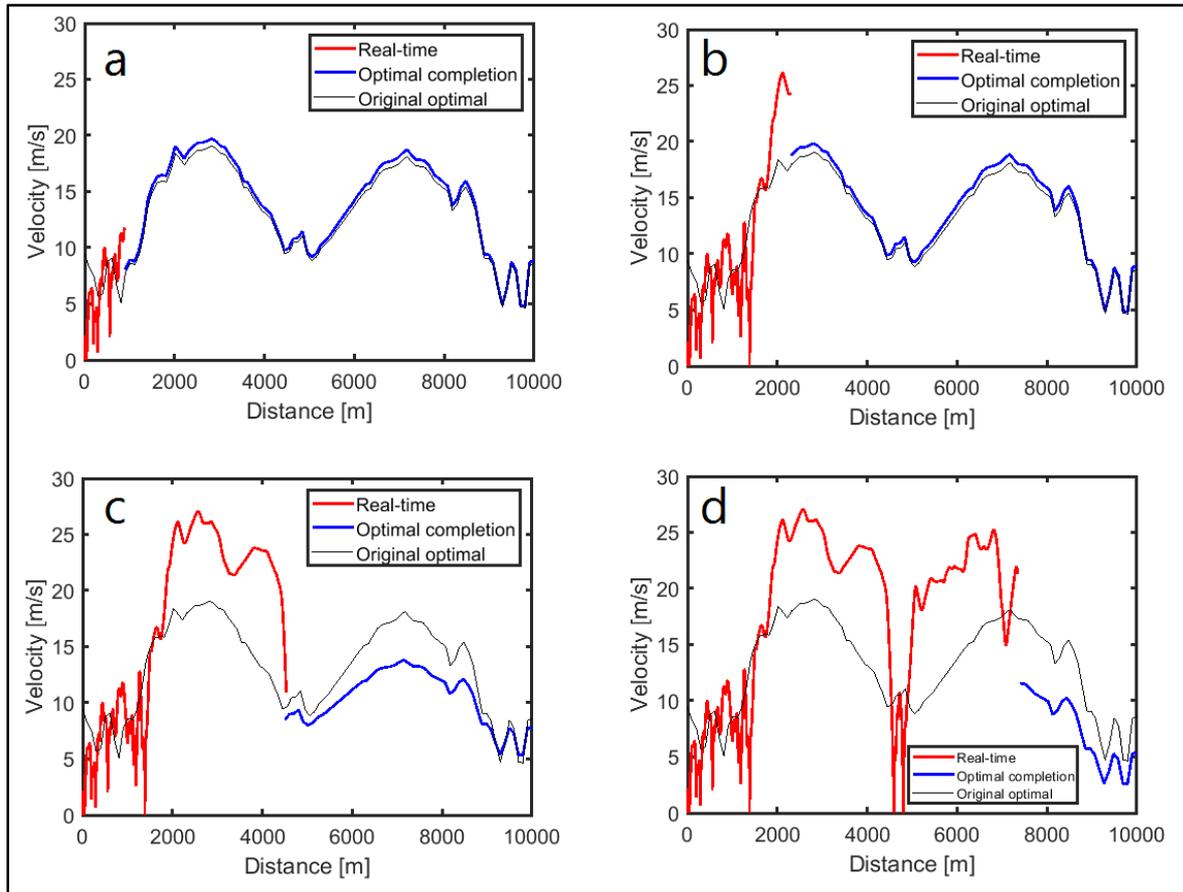


Figure 5.6: Real-time updating of the optimal velocity profile of route 4.1

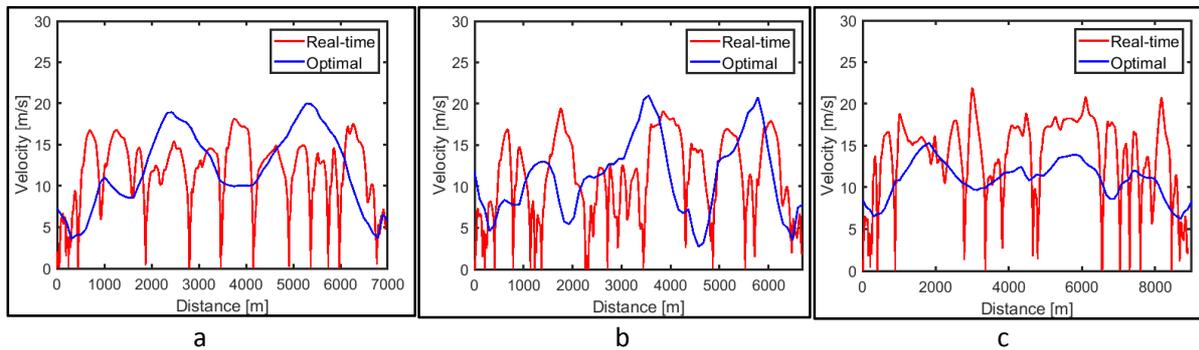


Figure 5.7: Real-time and final optimal velocity profile. a) Route 4.2, b) Route 4.3 and c) Route 4.4

The plots of energy consumed vs. distance are indicated in Figure 5.8 for the four routes evaluated. It is seen that the real-time optimal route velocity profiles are able to yield significant improvements in the

total energy consumed for a route. The complete optimal solution is able to yield even better savings, but is not able to run in real-time. The savings are summarised in Table 5.1.

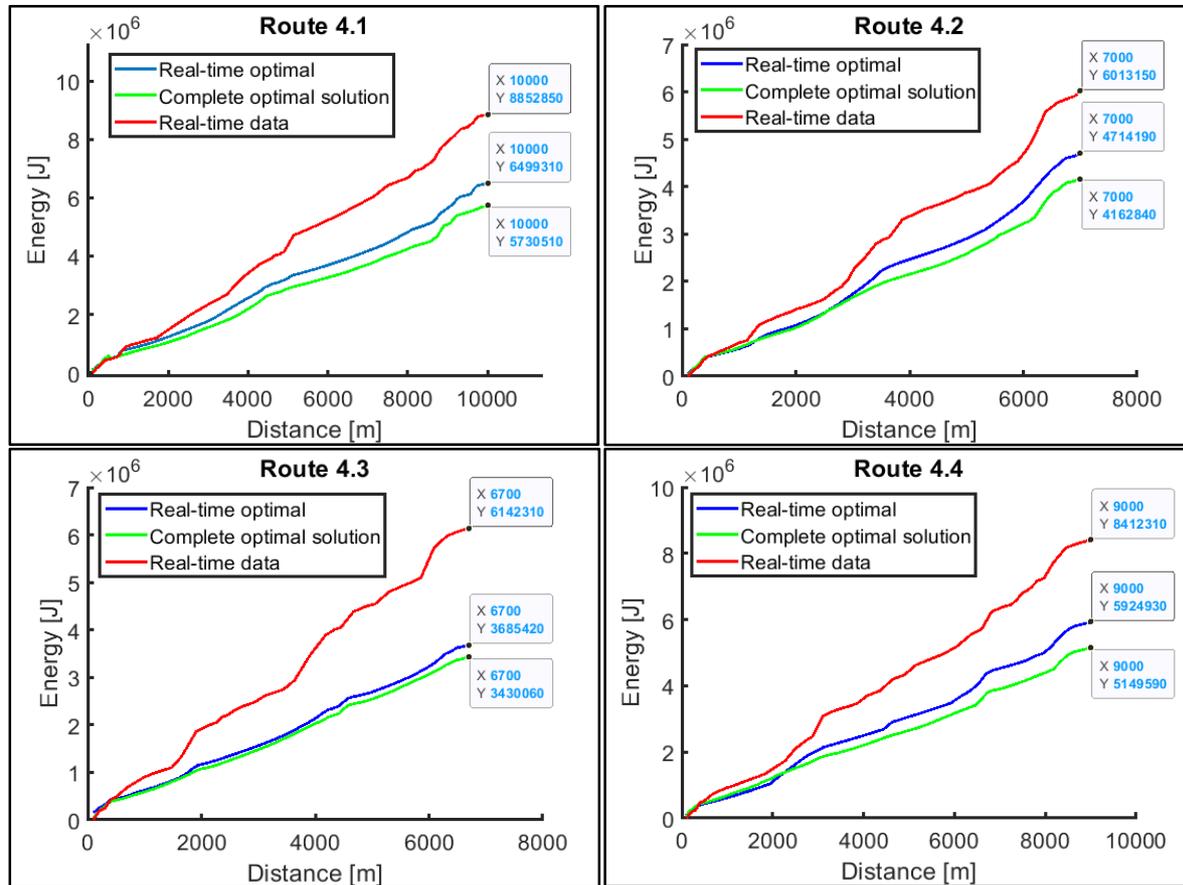


Figure 5.8: Route energy profiles

Route	Real-time proposed optimal saving%	Complete optimal solution saving %
Route 4.1	26.6	35.3
Route 4.2	21.6	30.8
Route 4.3	39.9	44.1
Route 4.4	29.6	38.8

Table 5.1: Energy saving comparison

From the results presented in Figure 5.8 and summarised in Table 5.1 it is seen that the complete optimal solution shows significant improvements over the original energy consumed. The real-time optimal is not as effective as the complete solution, but still shows the ability to significantly improve the total amount of energy consumed, compared to the uninformed travel data. If the vehicle has the ability to store energy onboard in the form of regeneration, this error reduces, as was shown in Chapter 4.4.6.

To illustrate the advantage of following a better velocity profile, the cumulative energy difference between the real-time data and the optimal velocity profile may be studied. For test route 4.1, this energy difference is plotted for varying regeneration efficiencies, shown in Figure 5.9.

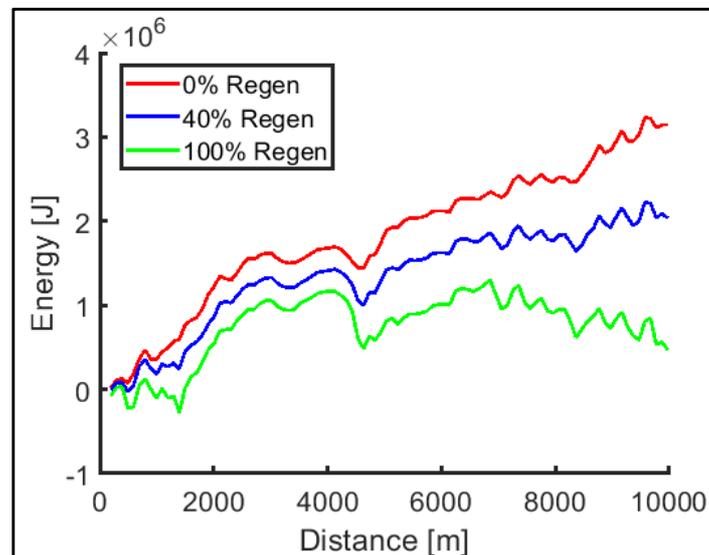


Figure 5.9: Total energy wasted by not following the optimal velocity profile

It was seen in Figure 5.6 that the optimal profile and that of the real-world data are not the same. In some locations, the optimal profile is significantly slower than the real-world, and in other cases, the real-world profile is slower than the optimal. The steep dip in the wasted energy of Figure 5.9 just after the 4000-metre point is because the optimal velocity here is significantly higher than the real-world data. If the optimal profile were to follow the real-world data instead, it would lead to the route not being completed in time. A small loss locally must be considered acceptable as it allows a bigger gain globally. As the regenerative efficiency increases the losses in improper velocity profile also reduce, with the losses mainly being due to the increased aerodynamic drag. It was seen in Chapter 4.4.6. that an increased regeneration efficiency leads to the consumed energy being less sensitive to the travel velocity profile. It should be noted that the current technology is still only capable of 40% regenerative efficiencies, with the 100% case being shown for interest only.

The original route had an energy consumption of around 8.8 MJ, and the optimal profile for this route had a final consumption of 5.7 MJ (from 4.4.6), which means that around 3 MJ was wasted in total by travelling at a velocity profile that is less than ideal, with similar total travel times. This equates to approximately 35% excess energy consumed. If the vehicle is capable of storing energy in a regenerative manner the amount of wasted energy can be reduced significantly, as the need for careful velocity profile management reduces.

5.4. The effect of incorrect route identification

It was mentioned in Chapter 4.3 that the initial incorrect identification of a route may be detrimental to finding the lowest energy for the true route being travelled. The optimal velocity profiles for the 4 test routes, Route 4.1 through 4.4, are plotted in Figure 5.10. It can be seen that the optimal velocity profiles

at the start of the routes, where surety on the route being travelled is lower, are very different. To simulate the effect of an incorrect route identification the optimal velocity profile for the wrong route will be used to determine the energy consumed up to the first 2.5 km. From Figure 5.10, the velocity profiles that differ the most are those of route 4.3 and route 4.4. The "optimal" velocity profile for route 4 will now be used as the travel velocity for route 3, and the energies compared.

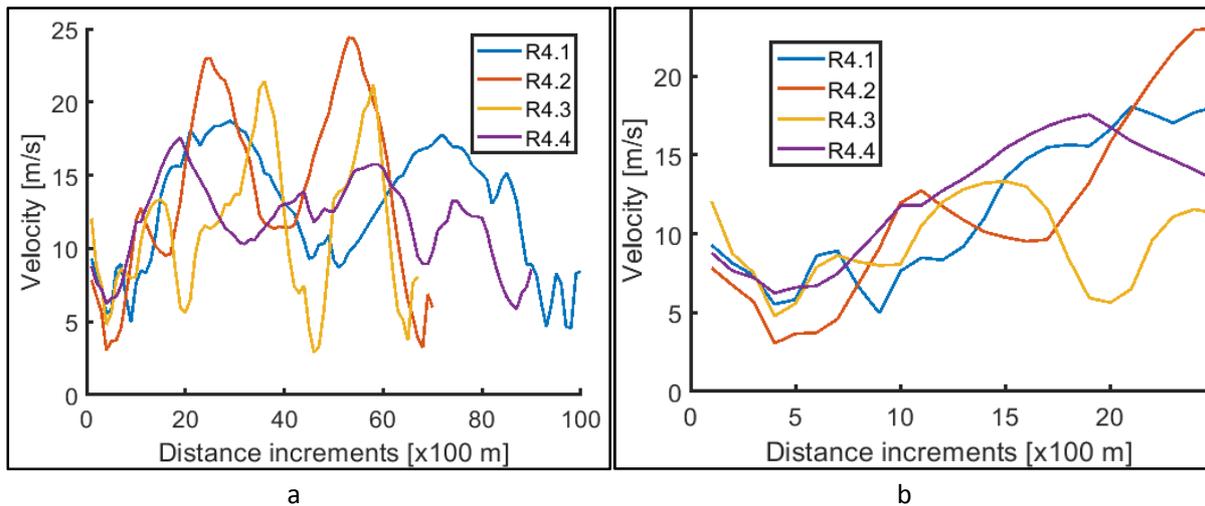


Figure 5.10: Optimal velocity profiles for the 4 real-world routes. a) complete, b) zoomed in initial

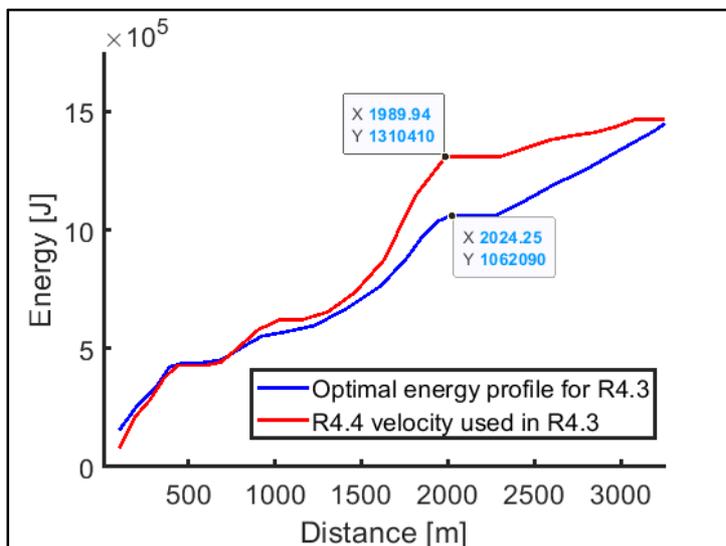


Figure 5.11: Route energy with incorrect route identification. R4.3 travelled with the velocity optimum of R4.4

It is seen from Figure 5.11 that an energy difference of approximately 0.25 MJ is found when route 4.3 was travelled at the optimal profile for route 4.4. The total optimised energy consumption for route 4.3 is 3.3 MJ, meaning that a wastage of 0.25 MJ amounts to an overspending of 8% on energy up to the 2.5

km point. A strategy needs to be implemented to ensure that, even though incorrect route identification may take place, the least amount of energy is wasted.

The route identification strategy presented in Chapter 2 yields a vector of route scores, which is a real-time tracking of the surety of each of the routes being the actual route being travelled. A proposed solution to the choice of optimal velocity profile to follow would be to use a normalised score value for the routes multiplied by their optimal velocity profile values. This approach is shown by Eq.5.1, with n being the number of routes. This will allow all routes that are identified with good surety to dictate the optimal profile, but not disregard the fact that other routes are still a possibility, as was shown in Figure 2.9. As more surety of the route being travelled is gained, the other routes will quickly drop down in surety score and lose their influence on the optimal result.

$$\overline{V}_{opt} = \frac{Score(1) \cdot \overline{V}_{opt}(1) + Score(2) \cdot \overline{V}_{opt}(2) + \dots + Score(n) \cdot \overline{V}_{opt}(n)}{\sum_1^n (Score)} \quad (\text{Eq.5.1})$$

5.4. Conclusion

The method presented to optimise route velocity profiles in real-time relies heavily on information from the previous chapters, such as the route identified and the vehicle mass. Variations (noise) in these parameters will introduce errors in the optimisation strategy as well, which will influence the lookup table's accuracy and thus influence the ability of the optimisation system to perform well. It is important to ensure that these previous systems deliver good quality data so that the optimisation approximation will end up with the lowest errors possible. If these are improved it will yield faster identification of routes, which will lead to less energy being wasted.

The real-time velocity profile optimisation method presented in this chapter proved capable of continuously updating the target travel velocity profile for the rest of a route to reduce energy wastage. For the four routes evaluated, the real-time optimal target velocity profile is able to reduce the overall energy consumption by 20-40%, depending on the route being travelled and how it was travelled originally. The complete optimal solution still provides even better results than the real-time approximation, typically up to 9% better, but without the ability to be run in real-time.

5.5. Recommendations

To try to improve the accuracy of the optimal velocity profile estimations it may be beneficial to look into making a surrogate model of the complete optimisations of Chapter 4 instead of the basic lookup table used in this study. This will potentially allow the real-time model to be able to account for specific constraints on the velocity that the current model is not able to do.

6. Conclusion of work

Reducing the consumption of energy is of paramount importance to society if we want to move closer to a sustainable future. A system capable of reducing the energy consumption of a vehicle in real-time was the focus of this study. The study was performed using a generalised approach, to increase the potential applications where the systems developed may be beneficial. The vehicle type that the study will be applied to is an underground electric/hybrid vehicle used in the mining sector in South Africa. It was found that the energy consumption of a vehicle to complete its route is mainly a function of the vehicle type, the route being travelled, the vehicle mass and the velocity profile that is followed to complete the route in the allowable timeframe. To enable successful optimisation of a system it is beneficial to have a good idea of the role players in the goal of the system. The vehicle type was assumed constant for the entire project and is easily updated for variations in the vehicle type and its associated parameters. A system capable of identifying the route being travelled was developed. A strategy to estimate the mass of the vehicle was developed. A strategy to find a solution to the optimal route velocity profiles was developed. A real-time executable approximation to the optimisation problem was developed. Individual conclusions are summarised for each phase of the project before an overall conclusion will be presented.

6.1. Route Identification

Using low-cost sensors an easily applied method of identifying a route has been developed. The advantage of the system developed is that it is capable of identifying routes above and below ground without the use of external transmitters or beacons. The system is capable of operating in real-time due to its computationally simple method of looking for patterns in heading and altitude changes for data as compared to a library of saved routes. The altitude data is beneficial to the rest of the project as well, as it is required in mass estimation and energy consumption calculations. If the route being travelled is not identified, the system stores this route as a new route in the database and can use it in route identification in future traversing of that route. The system can identify a route within a short distance, usually only requiring around 30 data points, depending on the noise in the data. A journal article was published on the content of this chapter [6].

6.2. Mass estimation

A method of estimating a vehicle's gross combined mass is developed and patented [7]. The system can estimate the combined mass of a vehicle and trailer (if fitted) to within 5% of the actual value. The system can run in real-time, making it usable in a real-time route energy estimation strategy as well. The system makes use of a single load cell and is insensitive to the weight distribution in the vehicle and trailer, but rather makes use of the propulsion force to move the entire system. The topography estimation from the route identification section is crucial to estimate inclines accurately, which plays a significant role in the force balance equation used to determine the vehicle mass. The system proved to be sensitive to calibration values for each of its parameters and caution should be taken when setting it up on a vehicle, thus requiring a decent calibration process. It is not reliable on soft terrain like gravel and sand roads but will work well on tarmac and railway applications.

6.3. Velocity Profile Optimisation

A robust method was developed that can provide a good initial guess to a local minimum solver such as *fmincon*, in Matlab. A first-principles model was constructed that can estimate the energy consumption for a vehicle as a function of the route topography and length, the vehicle type, the load and the velocity profile. A *fmincon* local solver was used to try to find the minimum of the cost function for the route energy based on finding a velocity profile that will meet the logistical requirements to complete a route in the prescribed time, but using the least amount of energy when doing so.

It was found that the optimal velocity profile for a given route had features resembling the inverse of the topographic profile of the route. An easily executable low-fidelity optimisation was performed to find the appropriate scale and shift factors that could be linearly applied to the topographic profile to yield a good initial guess for the optimisation problem. This method proved capable of yielding consistent local minima that were close to the global minimum for a given route. Simulation time was a concern though as high model fidelities resulted in simulations that would take hours to complete, making real-time application of the optimiser impossible.

It was found that mass has a direct influence on the total energy consumption, but not on the optimal velocity profile for a given route, which means that the optimal velocity profile need not be adapted when the vehicle mass changes along its route. This is of great help as it reduces the amount of system variables to account when trying to find a solution to the optimising problem. The results showed that the optimal velocity profile is simply a function of the route topography and the regeneration efficiency if the same vehicle is used.

An optimal route velocity profile proved capable of significant savings in energy consumed when traversing a route, even when the logistical constraint on time is constant. For test route 4.1 the vehicle consumed a total energy of 8.5 MJ in a real-world driving style. When optimising the velocity profile for that route the total energy could be reduced down to 5.7 MJ while still completing the route in a similar time. This is a 33% reduction in energy, which proves that driving style can greatly improve the energy consumption and cost of operation of a vehicle, be it electric, hybrid or IC as this optimisation is based on the route and not the vehicle. Constraints like speed limits, traffic and compulsory stops will influence the ability of the vehicle to travel at these optimal velocities, which is why a real-time implementation of this system is required.

The effect of regeneration was also studied and it was found that it is possible to reduce the total amount of energy required to complete a route even further by storing energy in times of excess. There exists a practical limit to the efficiency of this, which is around 40% for the majority of today's vehicle types. A 40% regeneration efficiency can lead to a 7% reduction in total energy usage. This shows that, when considered on their own, better driving styles may have a more significant effect on overall energy consumption than what hybridization of an IC vehicle or regeneration of an electric vehicle may have, and optimal energy conservation can only be successful when every aspect of energy consumption is optimised.

6.4. Real-time velocity optimisation

A method was developed to find an optimal route velocity profile for a vehicle on a known route, with known mass as well. This method showed potential in improving the energy consumption of a vehicle but had the drawback of being too slow to work in a real-time application. An approximate method of finding an optimal velocity profile was proposed and it proved capable of providing the vehicle controller, be it human or computer, with an improved target velocity profile to follow such that the overall energy consumption of the vehicle throughout the route may be reduced in real-time. This strategy works on the principle of scaling and shifting the topographic profile of the identified route, allowing the vehicle to use the topography to its advantage. The scale and shift parameters are easily found by performing a low-fidelity optimisation for a vehicle on its route. By performing a sweep of the key parameters; regeneration efficiency and maximum route completion time, the scale and shift parameters can be stored in a lookup table. This lookup table is smoothed to remove local outliers and provide a more predictable result. The smoothed lookup table data is used for very fast estimating of the optimal route velocity profile with errors of up to 9% introduced due to the model's inaccuracy, as compared to the complete optimal solution. The real-time method is still capable of achieving significant reductions in energy consumption in excess of 20% over the uninformed driver. This method allows graceful adaptation of the optimal velocity profile in real-time for a vehicle covering its route when unforeseen changes to the velocity plan occur. These unforeseen changes can be caused by unknown influences such as traffic.

6.5. Overall project

It is only possible to optimise a system which is understood well. Chapters 2 and 3 were required to provide knowledge on the route being travelled and the total mass of the vehicle such that the route energy could be accurately estimated and optimised for in Chapter 4. The optimisations of Chapter 4 proved successful, but are deemed impractically slow for any real-time application. A lookup-table-based approximation to the velocity profile is developed and applied in real-time. This allows the system to supply the vehicle control system with real-time updates on the best velocity to travel at whilst completing the route to conserve the most amount of energy while still satisfying the logistical constraints.

This thesis has contributed to the engineering audience by presenting a novel method for identifying a route being travelled, without the use of external sensors or beacons. A novel method was developed that is capable of estimating the gross combined mass of an entire vehicle rig in real-time. A method was developed that allows robust optimisation of the route velocity profile to take place. A method was developed that allows real-time approximations of the optimisation problem of the velocity profile that a vehicle should follow if it wants to reduce its overall energy consumption. Figure 6.1 summarises the programme strategy.

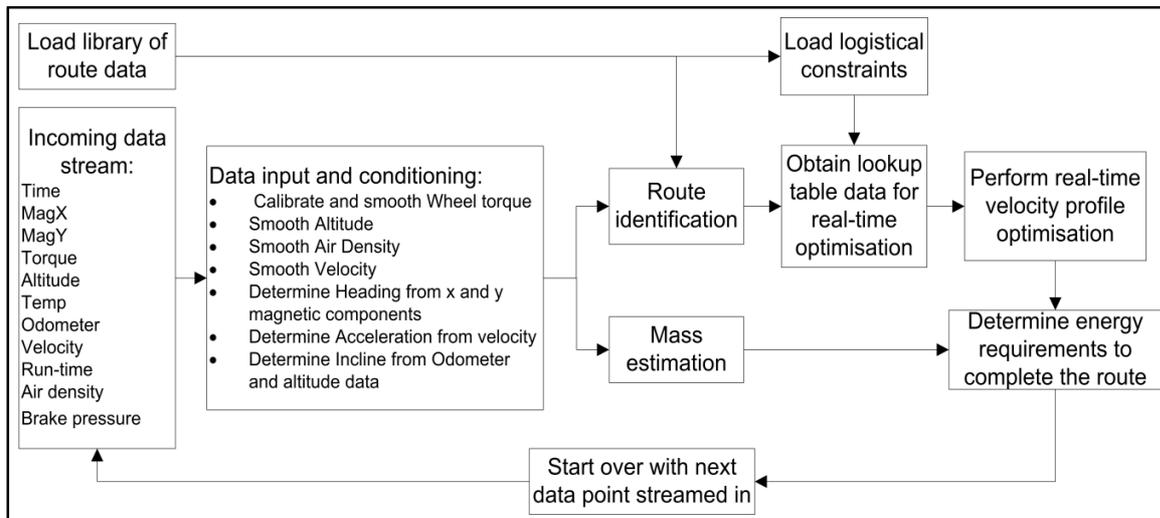


Figure 6. 1: Schematic of real-time velocity profile optimising system functioning

That being said, some recommendations for improving the various studies are presented in Chapter 7.

7. Recommendations and future work

7.1. Route Identification

For future work, the standard deviation value for each point of the route can be updated based on the stability of the streamed data for that portion of the route, and updated by averaging the data for many passes of the same section. This has the potential to improve route identification accuracy and responsiveness. If the system is implemented in an actual deep underground mine the absolute pressure obtained by the barometric pressure sensor can also be used to rule out certain routes from the stored set for that mine, as mine levels are usually many meters apart making the effect of weather changes to the absolute pressure readings fall away, which will lead to reduced calculation times.

7.2. Mass estimation

Build a statistical method that can be used to estimate the reliability of a mass estimation value in certain circumstances and increase or reduce its contribution to the estimated mass equation, which will lead to increased accuracy and responsiveness of the strategy.

7.3. Optimising of velocity profiles

The optimisation strategy proved successful in repeatably finding minimums to the optimisation problem. One drawback is the simulation time, especially when the model fidelity becomes high. A proposed improvement to the initial guess method used is to study the effect of adding a third parameter to the vertical scale and shift parameters. This parameter could be used to shift the inverted topographic profile horizontally as well, and studies could be done to investigate whether this could be beneficial to the optimisation problem.

It was seen that, if the route starts with an uphill section, that the initial part of the optimal velocity profile corresponding to the length of the uphill will have an approximately constant velocity. It can be studied whether altering the inverted topographic initial guess to have a flat initial part if the route starts with an uphill may be beneficial to finding better solutions to the optimisation problem or reducing the time it takes to find a solution.

The effect of making mass a vector can be studied, for better final energy estimations.

7.4. Real-time optimisation of the velocity profile

To try to improve the accuracy of the optimal velocity profile estimations it may be beneficial to look into making a surrogate model of the complete optimisations of Chapter 4 instead of the basic lookup table used in this study. This will potentially allow the real-time model to be able to account for specific constraints on the velocity that the current model is not able to do.

7.5. Factors not directly relevant to the work performed

This study focussed on developing a method of consistently and predictably finding a solution to the optimal velocity profile that a vehicle should follow to reduce the energy consumption of a vehicle travelling its route. The project specifically focussed on optimising factors outside of the vehicle itself.

For the case of plug-in hybrids, or fully electric vehicles the effect of charging off of mains can be studied, by looking at factors such as mains charging models etc.

For vehicle acceleration, there is a limit to the amount of power that the vehicle can deliver. This is influenced by motor sizing and traction, and in the case of electric and hybrid vehicles, battery sizing and motor driver sizing as well. These parameters may be included in optimisations to yield even better real-world applications of the real-time optimisation model presented in this study.

For this study, the regenerative ability of the vehicle was simulated to study the effect of these abilities. The models made use of a regenerative efficiency assumed to be constant. In reality, these efficiencies are not constant but depend on speed, temperature, power usage etc. For an improved optimisation an improved regeneration model may be of benefit.

To estimate route completion energies more accurately the mass estimation model may be updated to include looking at historic mass changes at specific locations on the route.

A detailed speed limit model can be incorporated into the optimisation strategy. This model can include historic speed profile data as well as legislative peaks.

Determining and optimising for the total cost of ownership of the vehicle can be included in the optimisation study, so not simply reducing the energy usage per route, but looking at reducing the cost to own and operate the vehicle can be studied. This can include accurate battery life models etc.

8. Acknowledgements

- The author would like to thank Gold Reef City Theme Park for their kind assistance in allowing underground testing at their facility.
- The author would like to thank Professor Schalk Kok for assisting in workload distribution and making time available to complete this project.
- The author would like to thank Dr. Helen Inglis for her continuous assistance over the last years in taking part of his workload upon herself and always motivating him to complete this project.
- The author would like to thank Dr. Nico Theron for his role as supervisor in the initial part of this project and for being a co-author to the published article that was done as part of Chapter 2, and co-author to the article that will be submitted for publication on the work covered in Chapter 3.
- The author would like to thank Professor Nico Wilke for his role in being the study leader for the last part of the project.
- The author would like to thank Dr. Danie Burger for the role he played in helping identify the topic for this study.
- The author would like to thank his parents for their continuous motivation and support throughout his life, and especially in completing this project. Their example has made him the man and father he is today.
- The author would like to thank his wife Amorie, for her love and support during the completion of this project.
- The author would like to thank his Lord and Saviour, Jesus Christ, for never letting go of me and guiding and teaching me throughout my journey in life.

9. Bibliography

- [1] A.D. Bosman, 'Investigating Load Shift and Energy Efficiency of New Technology Loco Battery Chargers', North-West University, South Africa, 2006. Accessed: Nov. 29, 2023. [Online]. Available: <https://repository.nwu.ac.za/handle/10394/1057>
- [2] W. G. Shaw, M. Mathews, and J. Marais, 'Holistic analysis of the effect on electricity cost in South Africa's platinum mines when varying shift schedules according to time-of-use tariffs', *J. Energy South. Afr.*, vol. 30, no. 4, pp. 26–40, Dec. 2019, doi: 10.17159/2413-3051/2019/v30i4a5675.
- [3] J. J. Eckert, L. C. A. Silva, E. S. Costa, F. M. Santiciolli, F. G. Dedini, and F. C. Corrêa, 'Electric vehicle drivetrain optimisation', *IET Electr. Syst. Transp.*, vol. 7, no. 1, pp. 32–40, Mar. 2017, doi: 10.1049/iet-est.2016.0022.
- [4] W. Wang, J. Shi, Z. Zhang, and C. Lin, 'Optimization of a dual-motor coupled powertrain energy management strategy for a battery electric bus', *Energy Procedia*, vol. 145, pp. 20–25, Jul. 2018, doi: 10.1016/j.egypro.2018.04.005.
- [5] E. Hellström, J. Åslund, and L. Nielsen, 'Design of an efficient algorithm for fuel-optimal look-ahead control', *Control Eng. Pract.*, vol. 18, no. 11, pp. 1318–1327, Nov. 2010, doi: 10.1016/j.conengprac.2009.12.008.
- [6] R. F. Meeser and N. Theron, 'Real-time underground route identification and route progress using simple on-board sensing and processing', *J. South. Afr. Inst. Min. Metall.*, vol. 123, no. 1, pp. 41–52, Feb. 2023, doi: 10.17159/2411-9717/1893/2023.
- [7] R. Meeser, 'A system and method for estimating a mass of a vehicle', PCT/IB2022/061330, 2022
- [8] A. Taghavipour, M. Vajedi, N. L. Azad, and J. McPhee, 'A Comparative Analysis of Route-Based Energy Management Systems for Phevs: Analysis of Route-Based Energy Management Systems for PHEVs', *Asian J. Control*, vol. 18, no. 1, pp. 29–39, Jan. 2016, doi: 10.1002/asjc.1191.
- [9] Y. Huang, H. Wang, A. Khajepour, H. He, and J. Ji, 'Model predictive control power management strategies for HEVs: A review', *J. Power Sources*, vol. 341, pp. 91–106, Feb. 2017, doi: 10.1016/j.jpowsour.2016.11.106.
- [10] S. F. Tie and C. W. Tan, 'A review of energy sources and energy management system in electric vehicles', *Renew. Sustain. Energy Rev.*, vol. 20, pp. 82–102, Apr. 2013, doi: 10.1016/j.rser.2012.11.077.
- [11] 'LOCO1000 Series'. Accessed: Sep. 07, 2021. [Online]. Available: <http://cmtigroup.co.za/LOCO1000-Series/>
- [12] K. Boriboonsomsin and M. Barth, 'Impacts of Road Grade on Fuel Consumption and Carbon Dioxide Emissions Evidenced by Use of Advanced Navigation Systems', *Transp. Res. Rec. J. Transp. Res. Board*, vol. 2139, no. 1, pp. 21–30, Jan. 2009, doi: 10.3141/2139-03.
- [13] L. Johannesson, N. Murgovski, E. Jonasson, J. Hellgren, and B. Egardt, 'Predictive energy management of hybrid long-haul trucks', *Control Eng. Pract.*, vol. 41, pp. 83–97, Aug. 2015, doi: 10.1016/j.conengprac.2015.04.014.
- [14] A. Lajunen, 'Fuel economy analysis of conventional and hybrid heavy vehicle combinations over real-world operating routes', *Transp. Res. Part Transp. Environ.*, vol. 31, pp. 70–84, Aug. 2014, doi: 10.1016/j.trd.2014.05.023.
- [15] R. Wang and S. M. Lukic, 'Review of driving conditions prediction and driving style recognition based control algorithms for hybrid electric vehicles', in *2011 IEEE Vehicle Power and Propulsion Conference*, Chicago, IL, USA: IEEE, Sep. 2011, pp. 1–7. doi: 10.1109/VPPC.2011.6043061.
- [16] S. Boyd and D. J. Nelson, 'Hybrid Electric Vehicle Control Strategy Based on Power Loss Calculations', presented at the SAE World Congress & Exhibition, Apr. 2008. doi: 10.4271/2008-01-0084.

- [17] A. M. Phillips, M. Jankovic, and K. E. Bailey, 'Vehicle system controller design for a hybrid electric vehicle', in *Proceedings of the 2000. IEEE International Conference on Control Applications. Conference Proceedings (Cat. No.00CH37162)*, Anchorage, AK, USA: IEEE, 2000, pp. 297–302. doi: 10.1109/CCA.2000.897440.
- [18] X. Hu, C. M. Martinez, and Y. Yang, 'Charging, power management, and battery degradation mitigation in plug-in hybrid electric vehicles: A unified cost-optimal approach', *Mech. Syst. Signal Process.*, vol. 87, pp. 4–16, Mar. 2017, doi: 10.1016/j.ymsp.2016.03.004.
- [19] M. Back, M. Simons, F. Kirschaum, and V. Krebs, 'PREDICTIVE CONTROL OF DRIVETRAINS', *IFAC Proc. Vol.*, vol. 35, no. 1, pp. 241–246, 2002, doi: 10.3182/20020721-6-ES-1901.01508.
- [20] Y. Yokoi *et al.*, 'Driving pattern prediction for an energy management system of hybrid electric vehicles in a specific driving course', in *30th Annual Conference of IEEE Industrial Electronics Society, 2004. IECON 2004*, Busan, South Korea: IEEE, 2004, pp. 1727–1732. doi: 10.1109/IECON.2004.1431842.
- [21] A. Vahidi, A. Stefanopoulou, and H. Peng, 'Recursive least squares with forgetting for online estimation of vehicle mass and road grade: theory and experiments', *Veh. Syst. Dyn.*, vol. 43, no. 1, pp. 31–55, Jan. 2005, doi: 10.1080/00423110412331290446.
- [22] Z. Chen, L. Li, B. Yan, C. Yang, C. Marina Martinez, and D. Cao, 'Multimode Energy Management for Plug-In Hybrid Electric Buses Based on Driving Cycles Prediction', *IEEE Trans. Intell. Transp. Syst.*, vol. 17, no. 10, pp. 2811–2821, Oct. 2016, doi: 10.1109/TITS.2016.2527244.
- [23] S. Rezaei and R. Sengupta, 'Kalman Filter-Based Integration of DGPS and Vehicle Sensors for Localization', *IEEE Trans. Control Syst. Technol.*, vol. 15, no. 6, pp. 1080–1088, Nov. 2007, doi: 10.1109/TCST.2006.886439.
- [24] S. Jeon, S. Jo, Y. Park, and J. Lee, 'Multi-Mode Driving Control of a Parallel Hybrid Electric Vehicle Using Driving Pattern Recognition', *J. Dyn. Syst. Meas. Control*, vol. 124, no. 1, p. 141, 2002, doi: 10.1115/1.1434264.
- [25] C.-C. Lin, S. Jeon, H. Peng, and J. Moo Lee, 'Driving Pattern Recognition for Control of Hybrid Electric Trucks', *Veh. Syst. Dyn.*, vol. 42, no. 1–2, pp. 41–58, Dec. 2004, doi: 10.1080/00423110412331291553.
- [26] V. Alvarez-Santos, A. Canedo-Rodriguez, R. Iglesias, X. M. Pardo, C. V. Regueiro, and M. Fernandez-Delgado, 'Route learning and reproduction in a tour-guide robot', *Robot. Auton. Syst.*, vol. 63, pp. 206–213, Jan. 2015, doi: 10.1016/j.robot.2014.07.013.
- [27] K. Song *et al.*, 'Multi-mode energy management strategy for fuel cell electric vehicles based on driving pattern identification using learning vector quantization neural network algorithm', *J. Power Sources*, vol. 389, pp. 230–239, Jun. 2018, doi: 10.1016/j.jpowsour.2018.04.024.
- [28] F. A. Bender, M. Kaszynski, and O. Sawodny, 'Drive Cycle Prediction and Energy Management Optimization for Hybrid Hydraulic Vehicles', *IEEE Trans. Veh. Technol.*, vol. 62, no. 8, pp. 3581–3592, Oct. 2013, doi: 10.1109/TVT.2013.2259645.
- [29] K. J. Rogers and R. S. Trayford, 'Grade measurement with an instrumented car', *Transp. Res. Part B Methodol.*, vol. 18, no. 3, pp. 247–254, Jun. 1984, doi: 10.1016/0191-2615(84)90035-3.
- [30] H. Liu, H. Li, M. O. Rodgers, and R. Guensler, 'Development of road grade data using the United States geological survey digital elevation model', *Transp. Res. Part C Emerg. Technol.*, vol. 92, pp. 243–257, Jul. 2018, doi: 10.1016/j.trc.2018.05.004.
- [31] R. Basso, B. Kulcsár, B. Egardt, P. Lindroth, and I. Sanchez-Diaz, 'Energy consumption estimation integrated into the Electric Vehicle Routing Problem', *Transp. Res. Part Transp. Environ.*, vol. 69, pp. 141–167, Apr. 2019, doi: 10.1016/j.trd.2019.01.006.
- [32] B. Y. Boroujeni, H. C. Frey, and G. S. Sandhu, 'Road Grade Measurement Using In-Vehicle, Stand-Alone GPS with Barometric Altimeter', *J. Transp. Eng.*, vol. 139, no. 6, pp. 605–611, Jun. 2013, doi: 10.1061/(ASCE)TE.1943-5436.0000545.

- [33] Bosch Sensortec, 'BMP280 Data sheet'. May 05, 2015. [Online]. Available: <https://www.bosch-sensortec.com/products/environmental-sensors/pressure-sensors/bmp280/>
- [34] Bosch Sensortec, 'BMP085 Digital pressure sensor'. [Online]. Available: <https://www.bosch-sensortec.com/products/environmental-sensors/pressure-sensors/>
- [35] Mining Technology, 'Top ten deepest mines in the world'. Accessed: Nov. 19, 2019. [Online]. Available: <https://www.mining-technology.com/features/feature-top-ten-deepest-mines-world-south-africa/>
- [36] M. J. Gallant and J. A. Marshall, 'The LiDAR compass: Extremely lightweight heading estimation with axis maps', *Robot. Auton. Syst.*, vol. 82, pp. 35–45, Aug. 2016, doi: 10.1016/j.robot.2016.04.005.
- [37] InvenSense, 'MPU-9250'. Jun. 20, 2016. [Online]. Available: www.invensense.com
- [38] H. I. Inc. Honeywell International Inc., 'HMC5883L_3-Axis_Digital_Compass_IC.pdf'. Feb. 2013.
- [39] I. I. D. I. Arduino, 'Arduino IDE'. GitHub, 2005. [Online]. Available: <https://www.arduino.cc/en/software>
- [40] The MathWorks Inc., 'Matlab'. The MathWorks Inc., 2021. [Online]. Available: <https://www.mathworks.com>
- [41] E. Hellström, M. Ivarsson, J. Åslund, and L. Nielsen, 'Look-ahead control for heavy trucks to minimize trip time and fuel consumption', *Control Eng. Pract.*, vol. 17, no. 2, pp. 245–254, Feb. 2009, doi: 10.1016/j.conengprac.2008.07.005.
- [42] A. Boretti, 'Improvements of Vehicle Fuel Economy Using Mechanical Regenerative Braking', SAE International, 2010. doi: 10.4271/2010-01-1683.
- [43] B. Yazdani Boroujeni and H. C. Frey, 'Road grade quantification based on global positioning system data obtained from real-world vehicle fuel use and emissions measurements', *Atmos. Environ.*, vol. 85, pp. 179–186, Mar. 2014, doi: 10.1016/j.atmosenv.2013.12.025.
- [44] K. Zhang and H. C. Frey, 'Road Grade Estimation for On-Road Vehicle Emissions Modeling Using Light Detection and Ranging Data', *J. Air Waste Manag. Assoc.*, vol. 56, no. 6, pp. 777–788, Jun. 2006, doi: 10.1080/10473289.2006.10464500.
- [45] S. Beyhan, F. E. Sarabi, and Z. Lendek, 'Takzgi-Sugeno Fuzzy payload estimation and adaptive control', in *IFAC PapersOnLine*, 2017, p. 6.
- [46] R. Zhang and E. Yao, 'Electric vehicles' energy consumption estimation with real driving condition data', *Transp. Res. Part Transp. Environ.*, vol. 41, pp. 177–187, Dec. 2015, doi: 10.1016/j.trd.2015.10.010.
- [47] E. Wilhelm, L. Rodgers, and R. Bornatico, 'Real-time electric vehicle mass identification', vol. 6, p. 6, 2013.
- [48] C. Fiori and V. Marzano, 'Modelling energy consumption of electric freight vehicles in urban pickup/delivery operations: analysis and estimation on a real-world dataset', *Transp. Res. Part Transp. Environ.*, vol. 65, pp. 658–673, Dec. 2018, doi: 10.1016/j.trd.2018.09.020.
- [49] C. Fiori, K. Ahn, and H. A. Rakha, 'Power-based electric vehicle energy consumption model: Model development and validation', *Appl. Energy*, vol. 168, pp. 257–268, Apr. 2016, doi: 10.1016/j.apenergy.2016.01.097.
- [50] K. Nishitani, 'Electromagnetic-Type Load Weighing Apparatus', 5,243,146, 1993
- [51] S. K. Yang, T. S. Liu, and Y. C. Cheng, 'Automatic measurement of payload for heavy vehicles using strain gages', *Measurement*, vol. 41, no. 5, pp. 491–502, Jun. 2008, doi: 10.1016/j.measurement.2007.07.003.
- [52] Windfinder, 'Pretoria average wind speed annually', Wind & Weather Statistics. Accessed: Nov. 21, 2019. [Online]. Available: https://www.windfinder.com/windstatistics/pretoria_rietvlei_dam
- [53] Thomas D. Gillespie, *fundamentals of vehicle dynamics*. SAE, 1992.

- [54] D. P. H. Kolodziejak, T. H. Pham, T. Hofman, and S. Wilkins, 'an optimization and analysis framework for TCO minimization of plug-in hybrid heavy-duty electric vehicles.pdf', in *IFAC PapersOnLine*, Elsevier, 2019, pp. 484–491. doi: 10.1016/j.ifacol.2019.09.077.
- [55] L. Szklarski, W. Dudek, and J. Machowski, *Underground electric haulage*, 1st ed. Poland, 1969.
- [56] Y. Zhou, A. Ravey, and M.-C. Péra, 'A survey on driving prediction techniques for predictive energy management of plug-in hybrid electric vehicles', *J. Power Sources*, vol. 412, pp. 480–495, Feb. 2019, doi: 10.1016/j.jpowsour.2018.11.085.
- [57] L. Xie, Y. Luo, D. Zhang, R. Chen, and K. Li, 'Intelligent energy-saving control strategy for electric vehicle based on preceding vehicle movement', *Mech. Syst. Signal Process.*, vol. 130, pp. 484–501, Sep. 2019, doi: 10.1016/j.ymsp.2019.05.027.
- [58] K. N. Genikomsakis and G. Mitrentsis, 'A computationally efficient simulation model for estimating energy consumption of electric vehicles in the context of route planning applications', *Transp. Res. Part Transp. Environ.*, vol. 50, pp. 98–118, Jan. 2017, doi: 10.1016/j.trd.2016.10.014.
- [59] S. Zhang and R. Xiong, 'Adaptive energy management of a plug-in hybrid electric vehicle based on driving pattern recognition and dynamic programming', *Appl. Energy*, vol. 155, pp. 68–78, Oct. 2015, doi: 10.1016/j.apenergy.2015.06.003.
- [60] P. Shen, Z. Zhao, X. Zhan, J. Li, and Q. Guo, 'Optimal energy management strategy for a plug-in hybrid electric commercial vehicle based on velocity prediction', *Energy*, vol. 155, pp. 838–852, Jul. 2018, doi: 10.1016/j.energy.2018.05.064.
- [61] W. Shabbir and S. A. Evangelou, 'Real-time control strategy to maximize hybrid electric vehicle powertrain efficiency', *Appl. Energy*, vol. 135, pp. 512–522, Dec. 2014, doi: 10.1016/j.apenergy.2014.08.083.
- [62] Y. Kim, M. Figueroa-Santos, N. Prakash, S. Baek, J. B. Siegel, and D. M. Rizzo, 'Co-optimization of speed trajectory and power management for a fuel-cell/battery electric vehicle', *Appl. Energy*, vol. 260, p. 114254, Feb. 2020, doi: 10.1016/j.apenergy.2019.114254.
- [63] S. S. Rao, *Engineering optimization: theory and practice*, 4th ed. Hoboken, N.J.: John Wiley & Sons, 2009.
- [64] Y. Zhang, H. Liu, Z. Zhang, Y. Luo, Q. Guo, and S. Liao, 'Cloud computing-based real-time global optimization of battery aging and energy consumption for plug-in hybrid electric vehicles', *J. Power Sources*, vol. 479, p. 229069, Dec. 2020, doi: 10.1016/j.jpowsour.2020.229069.
- [65] Snyman, JA, 'Practical mathematical optimization: an introduction to basic optimization theory and classical and new gradient-based algorithms', *Choice Rev. Online*, vol. 44, no. 05, pp. 44-2754-44–2754, Jan. 2007, doi: 10.5860/CHOICE.44-2754.

