Thesis for the degree of Licentiate of Engineering

Fuel-Efficient Truck Platooning using Speed Profile Optimization

SINA TORABI

Department of Mechanics and Maritime Sciences
CHALMERS UNIVERSITY OF TECHNOLOGY
Göteborg, Sweden 2017

Fuel-Efficient Truck Platooning using Speed Profile Optimization
SINA TORABI

THESIS FOR LICENTIATE OF ENGINEERING no 2017:07

Department of Mechanics and Maritime Sciences
Chalmers University of Technology
SE–412 96 Göteborg
Sweden
Telephone: +46 (0)31–772 1000

Cover:
A truck following its optimized speed profile over a road profile.

Chalmers Reproservice
Göteborg, Sweden 2017

Fuel-Efficient Truck Platooning using Speed Profile Optimization

Sina Torabi
Department of Mechanics and Maritime Sciences
Chalmers University of Technology

Abstract

This thesis is concerned with fuel-efficient driving strategies for heavy-duty vehicles driving on highways with varying topography. A method for reducing the fuel consumption of single trucks and platoons consisting of several trucks is described and evaluated both in simulation and in real trucks. The method, referred to as speed profile optimization (SPO), uses a genetic algorithm to find fuel-efficient speed profiles. Using SPO, the fuel consumption of a single truck was reduced by 11.5% (on average) relative to standard cruise control. The method's extension to platooning (P-SPO), reduced the fuel consumption by 15.8% to 17.4% for homogeneous and heterogeneous platoons (with different mass configurations), respectively, relative to the combination of cruise control and adaptive cruise control, when applied to road profiles of 10 km length. Furthermore, it was demonstrated that the results obtained in the simulations are sufficiently accurate to be transferred to real trucks.

The SPO and P-SPO methods also outperform the commonly used MPC-based methods by a few percentage points: For single trucks, SPO outperformed an MPC-based approach by 3 percentage points, in a case with identical roads and similar experimental settings. Similarly, for a platoon of two trucks, P-SPO outperformed an MPC-based approach by around 3 percentage points.

Keywords: fuel efficiency, truck platooning, heavy-duty vehicle platooning, speed profile optimization.

"Taking a new step, uttering a new word, is what people fear the most."

Fyodor Dostoyevsky

Acknowledgements

First and foremost, my sincere gratitude goes to my supervisor, and friend, Prof. Mattias Wahde for giving me a chance, believing in me, and inspiring me to begin my journey towards my PhD degree. Without you, and your endless guidance, I would not be the person I am today.

I would also like to thank the SERET project for funding the research presented in this thesis. I am also thankful for all the help and valuable inputs from the project partners in setting up the experiments.

I would also thank my colleagues in the VEAS division, and the Adaptive Systems research group in particular, for providing an excellent and friendly working environment, especially my fellow PhD students and friends in VEAS. A special thanks goes to Tushar, Pär, Dr. Leon Henderson, Toheed, and the RVAD PhD students for all the interesting discussions and the fun social activities we have had together so far.

And finally, thank you mom and dad, Siyavash, Nasim, Maryam, and Saman for your support. I would not be where I am today without you.

List of included papers

This thesis consists of the following papers. References to the papers will be made using Roman numerals.

I. Caltagirone, L., Torabi, S., and Wahde, M., *Truck Platooning based on Lead Vehicle Speed Profile Optimization and Artificial Physics.* In: Proceedings of the 18th International Conference on Intelligent Transportation Systems, Las Palmas, Spain, pp. 394-399, 2015.

II. Torabi, S. and Wahde, M., *Fuel Consumption Optimization of Heavy-Duty Vehicles using Genetic Algorithms.* In: Proceedings of the 2017 Congress on Evolutionary Computation, San Sebastián, Spain, pp. 29-36, 2017.

III. Torabi, S. and Wahde, M., *Heavy-Duty Vehicle Platooning based on Speed Profile Optimization.* Submitted to IEEE Transactions on Intelligent Transportation Systems, August 2017.

Technical terms used in the thesis

Table of Contents

Chapter 1

Introduction and motivation

The central topic of this thesis is fuel-efficient driving strategies for trucks operating on highways. Fuel accounts for approximately one third of the total cost of owning and operating a truck. Given that hauling companies own many vehicles that typically travel around 200000 km per year, reducing the fuel consumption even by a few per cent can translate to significant savings for these companies. Moreover, trucks are responsible for around 5% of total EU greenhouse gas emissions [9]. With the expected increase in demand for transportation of goods in the coming years, hauling companies and vehicle manufacturers are under pressure to take appropriate measures.

Vehicle **platooning**, a configuration in which a group of vehicles drive at small longitudinal inter-vehicle distances, has been investigated, both in academia and in the vehicle industry, as a means to reduce fuel consumption. Driving at small inter-vehicle distances may reduce the fuel consumption significantly, by reducing the overall air drag resistance. For example, if a platoon of two trucks drive at 80 km/h and at a constant spacing of 10 m, the aerodynamic drag experienced by the second vehicle is reduced by around 40%. Moreover, vehicle platooning can reduce traffic congestion by making better use of current road infrastructure, and also increase safety by automatically controlling the longitudinal motion of vehicles, something that may reduce the risk of rear-end collisions.

Most of the early work and experimental tests on platooning focused on the aerodynamic effects of driving at close inter-vehicle distances, while the **Lead vehicle** in the platoon maintained a constant speed, using **standard cruise control system** (CC), see, for example [10, 12, 8, 3]. This approach is optimal if there is no slope variation. However, on roads with varying slope, maintaining a constant inter-vehicle distance, or allowing only small variations, leads to excessive acceleration and braking, and consequently increases the fuel consumption. Therefore, the impact of road topography must be considered when developing fuel-efficient driving strategies.

1.1 Fuel-efficient driving

Vehicle platooning is one of the earliest proposed methods for reducing the fuel consumption of heavy-duty vehicles. However, it is not always possible to form or join a platoon. Thus, it is necessary to develop fuel-efficient driving strategies that work well both for single vehicles and for platoons.

Similar to maintaining a constant inter-vehicle distance in a platoon, driving at a constant speed on roads with varying topography is not fuel efficient. Consequently, the **speed profile** of a truck, i.e. its reference speed as a function of its longitudinal position along the road, must be allowed to vary. Speed variation methods could be used both in the case of single trucks as well as in truck platoons.

Fuel-efficient speed variation can be achieved using different methods, which will be reviewed in this thesis (see also Chapter 2). Many of the proposed methods are based on the optimal control framework, and model predictive control (MPC) in particular. In these methods, solutions are generated using the dynamic programming optimization method to obtain an optimal speed trajectory for the vehicle, which is tracked using iteratively calculated inputs from the MPC-based controller.

An alternative approach, however, is to use the **speed profile optimization** (SPO) method (first introduced in Paper I) to generate fuel optimal speed profiles, over a long road segment, for the trucks to follow (see Chapter 3). As shown in Papers II and III, SPO leads to larger savings than MPC-based methods, while being computationally less expensive. Moreover, using the SPO approach, there is no need for frequent updates of the speed profile.

This thesis consists of three papers which are concerned with the problem of fuel-efficient driving for heavy-duty vehicles over roads with varying topography. The main research idea presented and tested (both in simulations and experiments) in these papers is based on speed profile optimization, both in single trucks (Paper II) and in truck platoons (Paper I and III). In Paper I, a simple stochastic optimization method was used to design an optimized speed profile for the lead vehicle in the platoon while the rest of the vehicles follow the lead vehicle using various methods from the artificial physics framework. In Paper III, however, it was proposed that each vehicle should follow its own speed profile instead. This was achieved by optimizing the speed profiles of all vehicles together, using the optimization method introduced in Paper II, while considering the safety of the platoon.

The performance of the SPO method was improved by (i) modifying the road and speed profile representation from simple point lists to composite Bézier curves in Paper II, and (ii) using a more conventional genetic algorithm for optimization. Furthermore, in Paper II, it was shown that the results of using SPO for fuel-efficient driving can be transferred from simulations to real trucks despite the simple model used in the simulations.

1.2 Scope and author's contributions

As was mentioned above, the fuel-efficient driving strategy considered in this thesis is based on the concept of speed profile optimization, for single trucks and platoons (consisting of only trucks) driving on highways with varying topography. In this thesis, only the longitudinal dynamics of trucks is considered, which is customary in the field. Moreover, all trucks considered in this work have the same engine model, but with different masses in the case of heterogeneous platoons. Furthermore, trucks are equipped with cruise control systems in order to follow the optimized speed profiles. In general, during the testing of the methods presented in this thesis, it was assumed that other traffic does not interfere with the platoon (or the single truck, where applicable), e.g. by cutting in. However, this problem affects all platooning methods, and its effect on fuel savings will be considered later in this thesis.

The author was the main contributor to Papers II and III, and one of the main contributors to Paper I.

Chapter 2

Fuel-efficient driving strategies

Chassis and power train development has improved vehicle efficiency over the past decades, making vehicles safer and more fuel-efficient. In addition, driving behavior itself gives another opportunity of an even higher fuel-efficiency. There are several systems, which are usually referred to as **driving assistance systems**, that give recommendations to truck drivers through human-machine interaction systems so that they can reduce their fuel consumption, see, for example [16, 4, 28]. However, with the increasing level of autonomy in vehicles, one could develop fuel-efficient driving strategies that actively control the longitudinal motion of a vehicle, or a platoon of vehicles. In this chapter such strategies will be reviewed, and then the concept of speed profile optimization will be introduced.

The earliest works on fuel-efficient driving were focused on roads with constant slope. Schwarzkopf and Leipnik [25] formulated a fuel-optimal problem for a non-linear passenger vehicle model, and an analytical solution was proposed, based on the maximum principle, for highway driving on road segments with constant slopes. A similar approach was taken in [11] where the analytical solutions to a set of constructed simple road segments were derived for heavy-duty vehicles.

In order to solve the problem of fuel-efficient driving on general roads (with varying topography), different methods based on **optimal control** have been proposed, in which the optimization problem is solved numerically. In [20], Lattemann *et al.* proposed an upgrade on the standard cruise control system, called **predictive cruise control** (PCC), that allows the vehicle to drive at a lower speed as it drives through the uphill segments and to speed up when it traverses the downhill segments. PCC systems often operate at a narrow pre-defined speed range, typically around ±5 km/h.

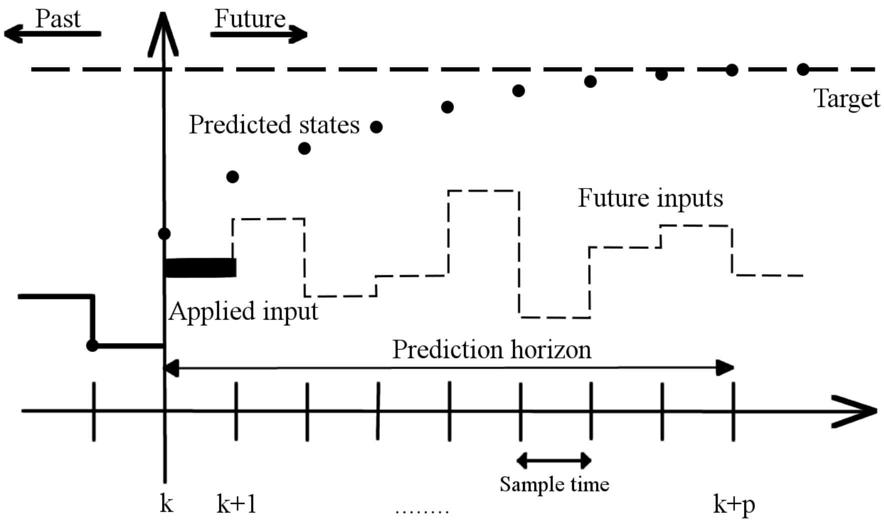

Figure 2.1: *Illustration of the basic idea in the model predictive control framework. At each time instant k, an optimal control problem is solved (using dynamic programming, for example) based on the predicted state of the model, which returns an optimal control sequence (future inputs) for the entire prediction horizon p. Then, only the first of the future control inputs (Applied input) is applied to the system. At the next time instant k+1, the same procedure is repeated with updated states.*

2.1 Model predictive control

PCC systems have been improved by the use of the **model predictive control** (MPC) framework in designing and tracking speed trajectories. In approaches that use this framework, an optimal speed trajectory is generated through online optimizations (discussed below), with respect to fuel consumption, typically over a 2 – 4 km horizon. The trajectory is then tracked by a specially designed controller in which the instantaneous control inputs, at each step, are computed based on the predicted state of the system.

In the MPC framework, **dynamic programming** (DP) [6] has been used extensively to solve the optimization problem numerically at each iteration, thus generating an optimal speed trajectory for the vehicle to follow.

MPC, in general, is a control framework that relies on iterative solutions, for which the update frequency depends on the discretization step of the optimal control problems (here minimizing the fuel consumption of trucks considering road topography). As illustrated in Figure 2.1, at each time instant k, an optimal control problem is formulated based on the predicted state of the model, and it is then solved numerically, a procedure that returns a sequence of control inputs. Of the computed control inputs, only the first element is applied to the system. This procedure is then repeated for each time instant. For the case of fuel-efficient driving strategies, the discretization step is typically in the range of 25 to 200 meters [14, 15, 17, 27].

2.2 Speed profile optimization

In this thesis, an alternative approach is proposed where an optimal speed profile is generated for a longer section of road, without having to solve the optimization problem iteratively at every position. In this approach, which is referred to as the **speed profile optimization** (SPO) method, the vehicle simply follows the optimized speed profile using a standard PID controller (for a detailed description of the method, see Chapter 3). In SPO, the optimization is carried out using a genetic algorithm [18] (see Chapter 3).

One of the main advantages of using SPO, as opposed to methods based on the MPC framework, is that SPO does not require any iterative online calculations: Once the speed profiles have been generated over the entire horizon, the vehicle can follow them without any further optimization. In this thesis, a horizon of 10 km length has been considered, but longer horizons are certainly possible in principle. However, it is also possible to use a shorter horizon to generate optimized speed profiles, and then gradually build a speed profile for an entire road while driving; since a given profile applies to a specific truck with a given load, the speed profile might not be known a priori, and in those cases it has to be optimized while driving. In the papers, the choice of a 10 km horizon was motivated partly by the fact that the time required to find suitable speed profiles (see Chapter 3) over such a horizon is typically a few minutes making it possible, in principle, to optimize speed profiles for the *following* 10 km section, while driving over the current section.

Additionally, as mentioned above, in the SPO method there is no need to use a specifically designed controller: Unlike the case of MPC-based methods where it is required that the vehicle should follow the speed trajectory exactly (or at least with very small error), the speed profiles in SPO act as recipes for the (varying) reference speed used in the vehicle's simple PID controller.

2.3 Applications

The above-mentioned **fuel-efficient driving strategies** can be used both in single vehicles and in vehicle platooning. In this section, different implementations of these strategies, along with their respective performance characteristics, will be reviewed.

2.3.1 Single vehicles

Fuel-efficient strategies for single vehicles have been implemented and tested rather extensively in the literature, both in simulations and experiments. The fuel savings obtained by methods based on PCC and MPC approaches typically fall in the range 3 to 7% (relative to standard cruise control with constant set speed, for highway driving).

Lattemann *et al.* [20] showed, in simulations, that the PCC system obtained fuel savings of about 3% when driving over a road segment of 25 km. In [13], Hellström *et al.* proposed an algorithm based on the MPC framework, referred to as look-ahead control (LAC), where dynamic programming was used to generate optimal speed trajectory for a truck to follow. Look-ahead control was tested, in simulation, over 120 km of highway road, and the fuel consumption was reduced by 3.5% on average.

Dynamic programming, however, suffers from the *curse of dimensionality*, that is, its computation time grows exponentially with the number of states. Much work has been done to overcome this problem. For instance, in [17], Henzler *et al.* proposed a method in which the problem of fuel-efficient driving was reduced to a convex MPC formulation that can be solved efficiently. This approach was tested in simulations only and obtained fuel savings of about 7% relative to standard cruise control.

In Paper I, the concept of SPO was introduced and tested, in simulations, for generating optimal speed profiles for the lead vehicle of a platoon. SPO lowered the fuel consumption of the lead vehicle by 15% (on average), relative to standard cruise control over 10 km highway road profiles with varying topography. In Paper II, the SPO method was implemented and tested, for a single truck, both in simulations and experiments. An important result in Paper II was the demonstration that the results obtained in simulations, even by using a rather simple vehicle model (see Subsect. 3.1.1), can be transferred to real trucks. That is, in general, the fuel savings obtained in real trucks are similar to those obtained in simulations.

Moreover, it was shown that SPO's fuel savings compares favorably to the savings obtained by other methods mentioned in Sect. 2.1. For example, in a close comparison between SPO and a standard PCC system, SPO improved the fuel savings by 3 percentage points, using the exact same road, with the same experiment settings. In more relaxed settings where the truck's speed was allowed to vary between 60 and 90 km/h, SPO obtained fuel savings of 10.2% (on average) when driving over 10 road profiles of 10 km length, generally outperforming the MPC-based methods by a few percentage points.

2.3.2 Platooning

Most of the early work on platooning was concerned with the concept of string stability, i.e. the ability of the controlled vehicle string to attenuate disturbances as they propagate through the platoon. Therefore, at the time, the main focus was on proposing control strategies and **spacing policies** that ensured the platoon's string stability [29, 30, 23, 26]. With the early experiments on platooning, and their relative success, the more practical aspects of platooning (other than just truck automation), such as increased fuel-efficiency and safety gained interest in the research community.

Most of the fuel-efficient driving strategies for platooning are focused on maintaining a close inter-vehicle distance to benefit from reduced air drag resistance. This was first achieved by using the combination of standard cruise control (CC), for the lead vehicle to maintain a constant speed, and **adaptive cruise control** (ACC) for the rest of the platoon to maintain their distance to the vehicle in front.

ACC is an upgrade on standard CC that allows a vehicle to control both the speed and the distance to the preceding vehicle by calculating the desired acceleration $a_R(t) = k_1(d_{i,i-1} - d_0) + k_2(v_{i-1} - v_i)$, where $x_i = x_i(t)$ is the longitudinal position of the i^{th} vehicle along the road, $d_{i,i-1} = x_{i-1} - x_i$ is the inter-vehicle distance, $v_i(t)$ is i^{th} vehicle's speed, k_1 and k_2 are hand-tuned gains, and d_0 is the desired distance to the preceding vehicle. The

ACC function has been used for controlling the vehicles of a platoon by allowing them to maintain a desired spacing policy by controlling both the inter-vehicle distance and the speed of a vehicle.

This approach, however, is not suitable on highways with varying topography, as was mentioned in previous chapter. On flat highways, truck platooning showed fuel savings of up to 10% [31, 7]. However, tests on highways with varying topography showed that the positive impact of reduced air drag can be neutralized by slope variations [1].

There are only quite few studies that consider the impact of road topography in order to improve the fuel-efficiency potential of platooning. In [2], Alam *et al.* proposed a new platooning strategy, look-ahead controller for platooning, that is based on the LAC method (see Subsect. 2.3.1) and was tested on synthetic road profiles (simple uphill and downhill segments). In this approach, which is based on the MPC framework, the LAC method is used to generate a fuel-optimal speed trajectory for every vehicle in the platoon first, and then the profile that requires the largest adjustment in velocity (compared to driving at constant velocity) is set as the common speed trajectory for all vehicles. This controller was tested with a platoon of two vehicles on a synthetic road profile of 4 km length (including an uphill and a downhill segment) and the fuel savings obtained were up to 14% on the downhill segment, and 0.7% during the uphill segments, relative to the CC and ACC combination. In [22], a similar approach was used to generate a common speed trajectory for all vehicles by combining each vehicle's fuel-optimal speed trajectory. The combination of speed trajectories was carried out by minimizing the deviations of each vehicle's speed reference from the common trajectory. This approach reduced the fuel consumption of four light (3500 kg) vehicles by around 6% when driving on a 90 km highway in Germany. Thus, to summarize, these methods generally provide fuel savings of around 6-7%

As discussed before, the SPO method for fuel-efficient driving has (several) advantages over MPC-based methods and was used (in Paper I) for generating fuel-optimal speed profiles for the lead vehicle of a platoon on road profiles of 10 km length on a Swedish highway. In Paper I, the rest of the platoon followed the lead vehicle using various formation control methods from the framework of **artificial physics** (AP), such as non-linear spring-damper model, modified artificial gravity model, etc., to keep a safe inter-vehicle distance and to follow the same speed profiles (see Paper I). This approach, i.e. SPO + AP, reduced the fuel consumption of the entire platoon by 15% (on average) compared to the case of CC + ACC, a result that compares favorably to the methods presented above. However, the use of the AP framework to control the **follower** vehicles of a platoon gave no significant improvement over and above those obtained with standard ACC systems (that employ a *linear* spring-damper approach as introduced above). Thus, the entire improvement was a result of the speed profile optimization for the lead vehicle.

In a more recent approach in [27], Turri *et al.* proposed a new platooning strategy where, instead of combining each vehicle's fuel-optimal speed trajectory, a common feasible trajectory is generated using DP considering all vehicles' characteristics. Then, this common trajectory is tracked by all vehicles, using MPC-based controllers. This method, which is referred to as cooperative look-ahead control (CLAC), was tested in simulations, both for homogeneous and heterogeneous platoons, on a typical Swedish highway of 45 km length. The performance of the platooning strategy was compared with the case in which each vehicle uses CC (constant speed) driving over the same road. CLAC obtained fuel savings of 10.8% in case of heavier second vehicle (relative to CC), and 5.4% in case of lighter second vehicle.

In Paper III, a rather different approach for controlling a platoon of vehicles was considered. In this approach, which is based on the SPO method, each vehicle received its own optimized speed profile, and then followed it independently of other vehicles. In other words, in this approach, which is referred to as **platooning SPO** (P-SPO), vehicles are not required to follow a specific spacing policy. This approach was tested in simulations, resulting in fuel savings of 15.8% for a homogeneous platoon and 16.7-17.4% for heterogeneous platoons of different mass configurations. The results obtained in this paper compares favorably to those obtained by MPC-based approaches.

Chapter 3

Speed profile optimization

The core idea behind the fuel-efficient driving strategy proposed in this thesis is the optimization, before driving (i.e. offline), of a vehicle's speed profile, using a **genetic algorithm** [18].

As was discussed in Chapter 2, a different approach, based on dynamic programming has been used extensively in the literature for generating fuel-efficient speed profiles during driving (i.e. online). However, dynamic programming suffers from the curse of dimensionality meaning that, in this approach, one must somehow limit the complexity of the problem. This is normally done by reducing the number of states considered, something that, in turn, implies a limited speed range. Moreover, the required discretization (in order to keep the problem computationally manageable) changes the problem in a way that may prevent one from finding the optimal solution, since the control inputs are assumed to be constant over the discretization step, which is typically around 80 m or more [14, 15, 17, 27]. Moreover, at least as it has been used in the literature, the dynamic programming problem is solved online (while driving) every few seconds (once for every discretization step) whereas, in our approach the entire speed profile is generated offline, *a priori*.

In this chapter, the vehicle model (used in Papers I – III) along with the PID controller used in the simulations will be described first. Next, the speed profile and road profile representation as well as the safety constraints considered during platooning are presented. Then, finally, the speed profile optimization and the evaluation method are described in detail.

3.1 Modeling

In order to evaluate the performance of the SPO method, a dedicated simulation environment was written, in the C# .NET programming language, implementing the **vehicle model** and the optimization method. In this simulation environment, the SPO method can be used both for a single vehicle and for a platoon of vehicles.

3.1.1 Vehicle model

Trucks are complicated systems with a large number of interacting parts requiring sophisticated mathematical models. However, in this work, only the **longitudinal motion** of a vehicle is considered, both in the case of single vehicles and in platoons of vehicles. Therefore, the dynamics of a vehicle can be expressed in the following form:

$$m(G)\dot{v} = F_{\mathrm{e}} - F_{\mathrm{b}} - F_{\mathrm{d}} - F_{\mathrm{r}} - F_{\mathrm{g}}, \tag{3.1}$$

where the terms on the right-hand side of the equation correspond to the forces experienced by the vehicle, namely, the engine force F_{e}, the braking force F_{b}, the air drag resistance force F_{d}, the rolling resistance force F_{r}, and the gravity force F_{g}. Furthermore, $m(G)$ is the total inertial mass of the HDV and it is computed as follow:

$$m(G) = m + \frac{J_w + \gamma_G^2 \gamma_f^2 \eta_G \eta_f J_e}{r_w^2}, \tag{3.2}$$

where G is the active gear, m is the mass of the vehicle, J_w and J_e represent the engine and wheel inertia, respectively, γ_G and γ_f are the gearbox and final-drive ratios, η_G and η_f denote the gearbox and final-drive ratio efficiencies, and r_w is the wheel radius. The various forces acting on the vehicle are described below.

Engine force:

The generated torque from a truck's engine is transferred to the wheels through the drive-line, i.e. clutch, gearbox, etc., and it is related to the engine force (F_{e}) acting on the vehicle in the following way

$$F_{\mathrm{e}} = \frac{\gamma_G \gamma_f \eta_G \eta_f}{r_w} T_e \equiv k_e T_e \tag{3.3}$$

where T_e is the generated torque, γ_G and γ_f again are the gearbox and final-drive ratios, η_G and η_f denote the gearbox and final-drive ratio efficiencies, and r_w is the wheel radius. In this work, in order to determine T_e, the inverse dynamics of the truck is considered: The requested acceleration, a_R, is computed from the speed profile, and then the required engine force is calculated by considering all the external forces $S = F_\mathrm{d} + F_\mathrm{r} + F_\mathrm{g}$ and the requested acceleration. Therefore, by rearranging terms in Eq. (3.1), one gets the requested engine torque T_e^R as

$$T_e^R = \frac{m(G)a_R + S}{k_e} \tag{3.4}$$

The effective acceleration, a_E, generated by the engine is then calculated as:

$$a_E = \begin{cases} a_R & \text{if } T_e^R < T_e^{\max} \\ \frac{k_e T_e^{\max} - S}{m(G)} & \text{if } T_e^R \geq T_e^{\max} \end{cases} \tag{3.5}$$

where k_e is the torque coefficient, and $T_e^{\max}$ is the maximum torque that can be generated by the engine. The instantaneous fuel consumption is determined by interpolation of the torque-RPM-fuel map for the modeled engine. In the case of vehicle platooning, it is here assumed that all the vehicles are equipped with the same engine.

Braking force:

A modern truck is equipped with several braking systems, such as foundation (disc) brakes, engine brakes, and retarders. The braking torque (and, consequently, the braking force) is often difficult to model since its characteristics vary significantly with the vehicle configuration and the braking logic. Therefore, in this thesis, the braking system is not modeled in detail. Instead, it is assumed that the brakes can generate any requested deceleration down to a limit of -2.5 m/s^2 [24].

Air drag resistance:

The air drag resistance experienced by a single vehicle is described as:

$$F_\mathrm{d} = \frac{1}{2} c_D A \rho_a v^2 \tag{3.6}$$

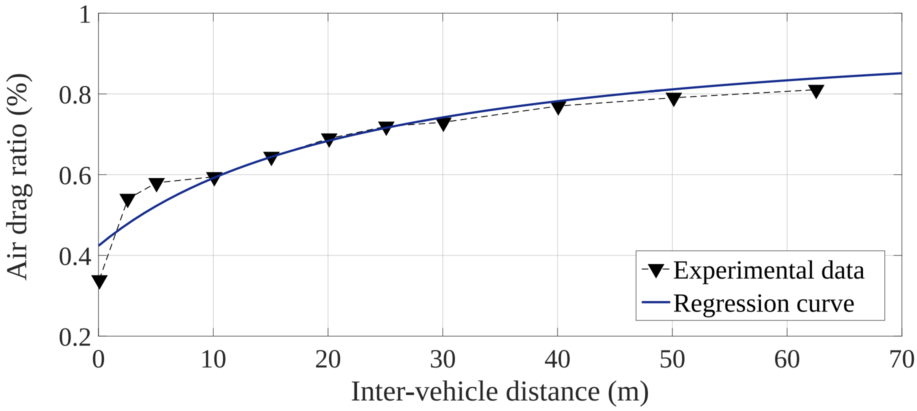

Figure 3.1: *Mapping of air drag reduction as a function of inter-vehicle distance. After Turri et al. [27].*

where c_D is the air drag coefficient, ρ_a is the air density, A is the frontal area of the vehicle, and v is the vehicle's speed. As was mentioned in previous chapters, driving at close inter-vehicle distances reduces the air drag resistance. This reduction is taken into account by considering a non-linear air drag ratio, $\Phi(d)$

$$F_{\mathrm{d}} = \frac{1}{2} c_D A \rho_a \Phi(d) v^2 \tag{3.7}$$

where $\Phi(d)$ is a coefficient that quantifies the **air drag reduction** as a function of the inter-vehicle distance, d, when driving behind another vehicle. $\Phi(d)$ is typically modeled based on empirical results from wind-tunnel experiments as [27, 21]

$$\Phi(d_{i,i-1}) = \left(1 - \frac{C_{D,1}}{C_{D,2} + d_{i,i-1}}\right) \tag{3.8}$$

where $d_{i,i-1}$ is the i^{th} vehicle's distance to its preceding vehicle, and $C_{D,1}$ and $C_{D,2}$ are constants, obtained through regression on the experimental data in [19]. The experimental data and the fitted curve are shown in Figure 3.1. Note that, in this model, the air drag reduction on the *preceding* vehicle is neglected (since it is much smaller than the reduction experienced by the follower vehicle).

Rolling resistance:

The rolling resistance, which is caused by the friction between the tires and the road surface, is a resistive force and is expressed as

$$F_{\mathrm{r}} = mgc_r \cos\alpha \tag{3.9}$$

where c_r denotes the rolling resistance coefficient, m is the vehicle's mass, g is the gravitational acceleration, and α is the road slope. Note that in this thesis, α is positive when driving through uphill segments of a road and negative for downhill segments.

Gravitational force:

Due to the large masses of trucks, the gravitational force plays an important role in the longitudinal dynamics and, consequently, the fuel consumption of trucks when driving on roads with varying topography. The gravitational force is expressed as

$$F_{\mathrm{g}} = mg\sin\alpha \tag{3.10}$$

where, again, m is the vehicle's mass, g is the gravitational acceleration, and α is the road slope. The gravitational force can either be propulsive (during downhill driving) or resistive (during uphill driving).

3.1.2 PID controller

In order for a truck to follow a desired speed profile, a simple **PID controller** has been used, with the error signal $e(t) = v_s(t) - v(t)$ where $v_s(t)$ is the reference speed (set to constant in standard CC), and $v(t)$ is the vehicle's instantaneous speed. The control output is calculated as

$$u(t) = K_p e(t) + K_i \int_0^t e(\tau)\mathrm{d}\tau + K_d \dot{e}(t) \tag{3.11}$$

where K_p, K_i, and K_d are hand-tuned proportional gain, integral gain, and derivative gain respectively. The requested acceleration, a_R, which is sent to the vehicle's power-train, is calculated by dividing the control output by the vehicle mass.

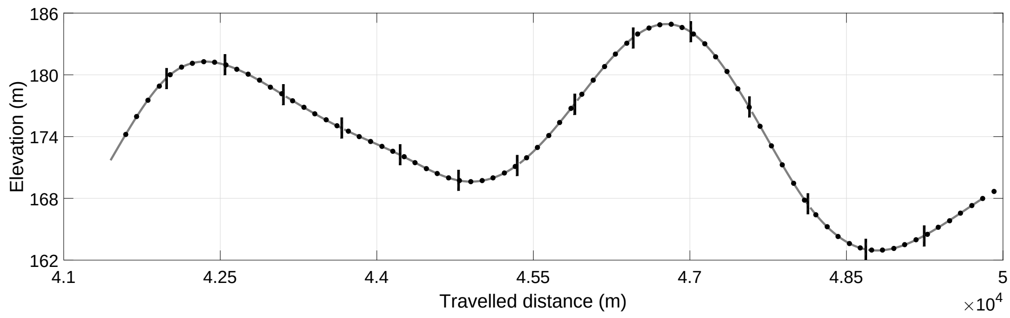

Figure 3.2: *Comparison between the fitted composite Bézier curve and the original data, over a part of the road between Göteborg and Borås. The dots represent the original data and the gray curve represents the fitted splines. The vertical lines separate individual splines in the composite Bézier curve.*

3.1.3 Road and speed profile representation

In order to evaluate the performance of the SPO method, both speed profiles and road profiles need to be modeled in the simulation environment. In the papers forming this thesis, two approaches were used to represent the road and speed profiles. In Paper I, simple lists of two-dimensional points were used to represent the profiles, giving elevation and speed values every 10 m. By contrast, in Papers II and III, a more compact representation of speed and road profiles was considered. In these papers, both profiles were represented using **composite Bézier curves**, i.e. sequences of Bézier splines; see e.g. [5]. In Papers II and III, two-dimensional cubic Bézier splines were used, of the general form

$$\begin{aligned}\mathbf{x}(u) \quad &\equiv (x(u), y(u))^T = \quad \mathbf{P}_0(1-u)^3 + 3\mathbf{P}_1 u(1-u)^2 \\ &\qquad + 3\mathbf{P}_2 u^2(1-u) + \mathbf{P}_3 u^3, \end{aligned} \tag{3.12}$$

where the vectors $\mathbf{P}_j$ are two-dimensional control points and u is a parameter ranging from 0 to 1. Using this representation reduces the search space during optimization (see Sect. 3.2) relative to the case involving simple lists.

Road profiles

Since only the longitudinal motion of vehicles is considered here, a **road profile** can be modeled by using a two-dimensional composite cubic Bézier curve, as described above. With this representation, the two dimensions are the longitudinal position along the road and the elevation, in the following form:

$$(s, z) \equiv (s_i(u), z_i(u)), \quad i = 0, \ldots, n-1 \tag{3.13}$$

where n is the total number of splines used to model a road section. With this representation, it is possible to write the elevation as $z = z(s)$, since for any given position s along the road, the corresponding spline index and $u-$value can be found.

The number of splines needed to fit a composite Bézier curve to a list of position-elevation pairs can be selected in various ways. For instance, it is possible to fit a composite Bézier curve to a data set such that the curve passes through all the data points. This approach, however, is not so suitable. First of all, if the number of splines approaches the number of data points, no reduction in the search space size is achieved (having the speed profile optimization procedure in mind) compared to the case where a simple point list is used. Second, fitting the curve to all data points is unsuitable due to the presence of noise in the data. In fact, in Paper II, it was shown that by using a more compact representation of the road (and speed) profiles, the overall performance of the optimization method improved. For the data set used here, the number of splines ranges from 18 to 22 for each road profile. An example of a fitted curve is shown in Figure 3.2.

Speed profiles

Similar to the road profiles described above, the **speed profiles** are modeled using two-dimensional composite cubic Bézier curves as:

$$(s, v) \equiv (s_i(u), v_i(u)), \quad i = 0, \ldots, n-1 \tag{3.14}$$

where v is the vehicle's longitudinal speed and s again is the longitudinal position of the vehicle along the road. Thus, the speed of a vehicle can be written as $v = v(s)$. With this representation, when generating speed profiles in order to calculate the desired speed based on the vehicle's current longitudinal position, one must use the same splines for s as in the road profiles.

3.1.4 Safety constraint

Since in the proposed approach for platooning, P-SPO, the inter-vehicle distances are not controlled directly, it is necessary to have **safety constraints** during the optimization. In this thesis, the safety of a platoon is guaranteed by preventing the inter-vehicle distance from going below a safe distance at any time. At each time step during the optimization, the **safe distance** is calculated as [30]:

$$d_{i,i-1}^{\text{safe}}(t) = d_0 + h(t)v_i(t) \tag{3.15}$$

where $d_{i,i-1}^{\text{safe}}(t)$ is the minimum allowed distance between the i^{th} and $(i-1)^{\text{th}}$ vehicles at time t, d_0 is the absolute allowed minimum distance, $h(t)$ is the variable time headway, and $v_i(t)$ is the i^{th} vehicle's speed at time t. The variable time headway $h(t)$ is expressed as:

$$h(t) = h_0 - c_h v_r(t) \tag{3.16}$$

where $h_0 > 0$ is the (constant) minimum time headway, $c_h > 0$ is a constant, and $v_r(t) = v_{i-1}(t) - v_i(t)$ is the relative speed. For safety reasons, the variable time headway $h(t)$ is not allowed to become negative, while very large headways are undesirable as they may increase the inter-vehicle distances beyond the limit where the platoon can be considered coherent. Here, the variable time headway value has been limited to the interval $[0, 1]$ (s) and the values of h_0 and c_h have been set to 0.1 (s) and 0.2 (s^2/m), respectively, as was proposed in [30].

3.2 Optimization method

Consider the problem of moving a truck from a given starting point to a given finishing point, in a case where the road profile, or at least a part of it (for example 10 km), is known *a priori*. Then, the motion of the truck can be formulated as a **desired speed profile**, $v_d(s)$, defined over the entire road profile, where v_d is the desired speed at any given longitudinal position s along the road. Of course the vehicle's speed must vary based on the road topography in order to reduce its fuel consumption, as discussed in Chapter 1. Therefore, the problem of minimizing the fuel consumption of a truck can be formulated as finding an optimal speed profile first, and then driving accordingly over the road. In this problem formulation, the effect of the external traffic on the vehicle's motion is not considered; that is, the vehicle is assumed not to be disrupted from following its optimized speed profile during driving. However, the effect of external traffic has been considered in Paper II, and it will be further discussed in Chapter 4.

For a platoon of vehicles, two approaches have been considered. In Paper I, the **leader-follower** approach has been used, in which the motion of the lead vehicle is formulated as an SPO problem for a single vehicle, while the follower vehicles control their distance to the preceding vehicle. In Paper III, however, the P-SPO method (see Subsect. 2.3.2) has been used, in which the fuel consumption minimization problem for all vehicles is formulated using SPO. Thus, in this case, the problem of minimizing the fuel consumption of a platoon of trucks is reduced to finding an optimal speed profile for each vehicle.

3.2.1 Evaluation method

For the purpose of evaluating a speed profile with respect to fuel consumption, a dedicated simulation environment was written (in C# .NET) where the truck model described in Subsect. 3.1.1 and the controllers described in Subsect. 2.3.1 were implemented, as well as the road and the speed profiles introduced in Subsect. 3.1.3. Now, assuming that a speed profile is available, the truck can be made to follow it so that its fuel consumption can be measured.

The **speed profile evaluation** for a single vehicle proceeds as follows: At each time step, the current longitudinal position of the truck is used to calculate the desired speed from the speed profile. Since the speed profile is defined in advance and is thus available during driving over the given road profile, the desired speed can be calculated easily, either by linear interpolation in the case of a point list (Paper I), or from the splines, as described in Subsect. 3.1.3 (Papers II and III). Thus, the speed profile is used as a lookup table from which the desired speed is extracted at any given longitudinal position along the road. The obtained desired speed is then fed to the truck's PID controller as its reference point. When the vehicle passes the finishing point of the road section a single number, namely the fuel consumption, is returned.

For a platoon of trucks, if the P-SPO method is used, all the vehicles follow their individual speed profiles and the evaluation procedure is thus the same (for each vehicle) as in the case of a single vehicle; see Paper III. Unlike other methods, with P-SPO each vehicle receives and follows its

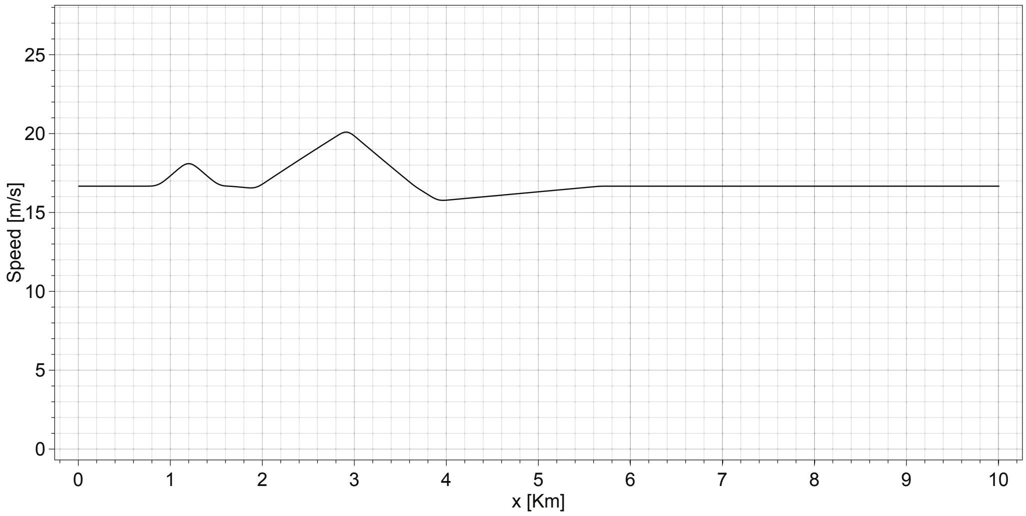

Figure 3.3: *An example of speed profile tweaking in the RMHC method. In the snapshot shown, taken from the early stages of an optimization run, the initial (flat) speed profile has undergone three tweakings.*

own speed profile *independently* of other vehicles. In the case of the leader-follower approach (Paper I), only the lead vehicle has a speed profile define *a priori*. In this case, the evaluation proceeds as follows: At any time step, the distance between the current vehicle and its preceding vehicle is calculated and is then used to compute the desired acceleration based on the adopted model, for example ACC (see also Subsect. 2.3.2).

3.2.2 Optimization algorithm

In this thesis, the optimization of the speed profiles is carried out using evolutionary algorithms with respect to fuel consumption. In Paper I, a simple genetic algorithm with one individual that encodes a speed profile using a simple list, specifying the desired speed at a number of discrete points (here every 10 m) was used. This method which can also be referred to as **random mutation hill climbing** (RMHC) proceeds as follows: First, it starts from a flat speed profile, i.e. identical to the case where the vehicle uses the CC function. Then, the truck's fuel consumption f_0 is computed using the evaluation procedure described in Subsect. 3.2.1. Since only one individual is considered here, f_0 is also the minimum fuel consumption f_{min} found so far. Next, the speed profile is tweaked by randomly selecting a point x_{t} at a random location. The speed profile at that point is then changed by a fraction β (either positive or negative) and, in addition, a randomly selected range r is used during tweaking such that the speed values in the interval $[x_{\mathrm{t}} - r, x_{\mathrm{t}} + r]$ are changed linearly. The change $\Delta v_{\mathrm{d}}(x)$ in the speed profile is thus computed as

$$\Delta v_{\mathrm{d}}(x) = \begin{cases} \left(1 - \left|\frac{x - x_{\mathrm{t}}}{r}\right|\right) \beta v_{\mathrm{d}}(x) & \text{if } |x - x_{\mathrm{t}}| < r \\ 0 & \text{otherwise} \end{cases} \tag{3.17}$$

An example of the tweaking procedure is shown in Figure 3.3. Finally, a smoothing step is applied to the speed profile using a simple, centered moving average of length $L = 3$. Once the new speed profile is generated, the fuel consumption of the truck is measured when following this profile. If the resulting fuel consumption, f_{new}, is lower than the current minimum fuel consumption, the new speed profile is kept and the value of the minimum fuel consumption is updated accordingly. Otherwise, the new speed profile is discarded and a new tweaking is applied to the previous speed profile as described above. One should note that the generated speed profile must fulfil certain constraints, namely (i) the instantaneous maximum speed v_{max} must never exceed the road's speed limit and (ii) the average speed should always be above a certain threshold $\bar{v}_{\mathrm{min}}$. In Paper I, these constraints were applied as hard constraints, i.e. if any of the constraints were violated during evaluation, the corresponding speed profile was discarded immediately.

In Paper II and Paper III, the optimization was carried out using a fairly standard GA. The optimization algorithm keeps a population of M individuals (typically around 100) where each individual defines N speed profiles (in case of a platoon of N trucks, Paper III), or one profile (i.e. $N = 1$) in case of a single truck (Paper II). In these papers, the speed profiles were represented by composite Bézier curves in order to improve the performance of the optimization process; see Subsect. 4.4. The individuals (chromosomes) are encoded using floating-point numbers, where each number (gene) represents the second component (i.e. the speed, v) of a spline control point $\mathbf{P}_{i,j}$ (where $i = 0, \ldots, n-1$ denotes the spline index and $j = 0, \ldots, 3$ denotes the control point index for the spline in question); see also Eq. (3.12). In order to make sure that the decoded individual results in smooth speed profiles, two additional requirements are considered during **encoding**. Since a speed profile (i.e. a composite Bézier curve) consists of several splines, the **positional continuity** (C0) of the profile should be taken into account. In other words, the encoding must be such that a decoded individual forms a speed profile that is continuous over the entire stretch of the road. The second requirement is to ensure that the generated speed profile is smooth. Therefore, the **derivative continuity** (C1) of the speed profile must be preserved in the encoding. In order to make sure that these requirements are met, the following conditions must hold:

$$P_{i,3} = P_{i+1,0} \quad \text{for} \quad i = 0, \ldots, n-1 \tag{3.18}$$

and

$$P_{i,3} - P_{i,2} = P_{i+1,1} - P_{i+1,0} \quad \text{for} \quad i = 0, \ldots, n-1 \tag{3.19}$$

Considering these requirements, the number of parameters (i.e. the length of the individuals) will be $L = N\,(4 + 2(n-1)) = N(2n+2)$ where n is the number of splines. The decoding procedure results in a set of N speed profiles which is then evaluated as described in Subsect. 3.2.1. The **fitness measure** of an individual is then taken as the inverse of the fuel consumption of the truck while following the decoded individual.

Similar to the optimization method used in Paper I, the generated speed profiles (in Papers II and III) must fulfil certain requirements, namely (i) the instantaneous maximum speed $v_{\max}$ must never exceed the road's speed limit, (ii) the average speed should always be above a certain threshold $\bar{v}_{\min}$, and (iii) the instantaneous minimum speed $v_{\min}$ should be above a user-defined threshold to ensure that the vehicle does not affect the traffic negatively (a condition that was not used in Paper I). Moreover, in Paper III, two additional constraints were considered to ensure the safety and cohesion of the platoon. For the platoon to remain coherent, the inter-vehicle distance was bounded from above by a threshold (here 40 m) at all times. Moreover, the inter-vehicle distance was required always to be larger than the safe distance defined in Eq. (3.15). If any of the **optimization constraints** described above are violated, the fitness value is multiplied by a penalty term smaller than 1. For instance, the **penalty term** for a case in which the instantaneous inter-vehicle distance exceeds its maximum allowed value is calculated as follows:

$$p_{d_{i,i-1}}(d) = \mathrm{e}^{-c_d\left(\frac{d_{i,i-1}}{d_{\max}}-1\right)^2}, \tag{3.20}$$

where (in this case) $d_{i,i-1}$ is the *maximum* distance between the i^{th} and the $(i-1)^{\text{th}}$ vehicle during the evaluation, c_d is a constant, and $d_{\max}$ is the maximum allowed inter-vehicle distance. Similar penalty terms are used for the other constraints.

Selection is carried out using standard tournament selection with the tournament size S_t (from 2 to 5) and a tournament selection probability p_t (around 0.7 to 0.8). Single-point crossover is used with the probability p_c (around 0.7 to 0.9). Once a new generation is formed, mutation is applied with probability p_m (defined as $1/L$ where L is the length of the individual). The mutation is carried out either full-range (with probability of 0.5) or as a real-number creep mutation (also with probability of 0.5) where the creep rate is set to 10% of the full-range mutation. Finally, to ensure that the current best solution is preserved, elitism is used to pass the best current individual to the next generation. An example of speed profile optimization using the GA is shown in Figure 3.4.

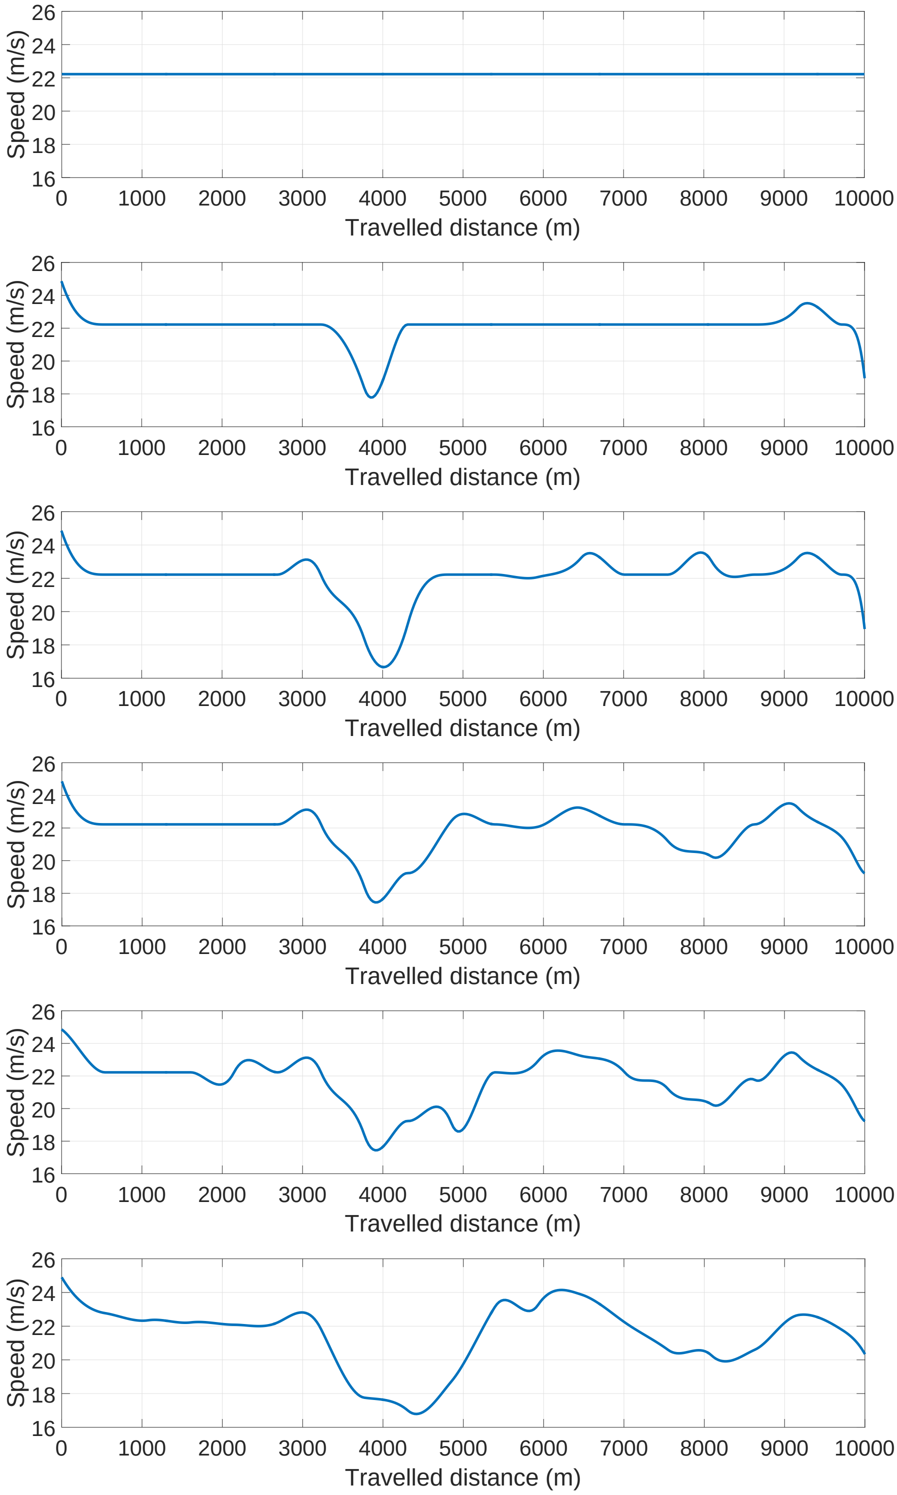

Figure 3.4: *An example of speed profile evolution using the GA applied to road profile 9 from Paper II. In this figure, the evolution of the speed profile from its initial flat shape (top panel) is shown at different generations. The bottom panel shows the final optimized speed profile used for driving over this road profile.*

Chapter 4

Discussion

This chapter begins with a discussion on fuel consumption reduction (relative to CC and ACC) of single trucks and platoons, with emphasis on the SPO and P-SPO methods introduced in Chapter 3. Next, the advantages regarding fuel consumption of using these methods over recent MPC-based approaches are discussed, emphasizing the control architecture, the vehicle model, and horizon length. The optimization method and its components are then discussed in further detail. Then, the effects of external traffic are considered. The final section discusses the performance of the SPO method when applied to platoons in which the trucks have different power trains.

4.1 Single vehicles

In Paper II, the SPO method was applied to 10 road profiles of 10 km length for a single truck. The fuel consumption was reduced by 11.5% in simulations and 10.2% in experiments (on average), relative to standard CC (see Tables I and II in Paper II). The results reported in Paper II show that the SPO method does indeed work well in real trucks, and that the results obtained in the simulations are transferable to real vehicles despite the simplicity of the adopted vehicle model. While the differences in the fuel savings obtained in the simulations and in the experiments (see Table I in Paper II) generally are very small, there are some exceptions. For example, one should note that the motion of the truck can be disrupted during the experiments, an effect that is not implemented in the simulations. Thus, on occasion, the real truck may not be able to follow the speed profile; hence the exclusion of road profiles 1 and 9 from Table I in Paper II. Moreover, during the experiments, the truck can experience the drag-reducing effect of driving behind other trucks, an effect that is not modeled in the single-vehicle simulations (but of course included in the platooning simulations; see below). This influences the fuel consumption in the experiments and causes differences between the results obtained in the simulations and the experiments. Furthermore, another important factor is the (rare) inability of the real truck to follow the speed profile at a few instances due to the imposed limits on the requested acceleration from the truck's PID controller which, in Paper II, were slightly different from the corresponding limits in the simulations.

4.2 Platoons

As presented in Subsect. 2.3.2, using the SPO method for platooning, in combination with common controllers such as ACC, can significantly reduce the fuel consumption of a platoon of trucks. In Paper I, it was reported that the SPO+ACC method reduced the fuel consumption of a homogeneous platoon of three trucks by around 15% on average, relative to the case of CC+ACC over ten road profiles (see Table II in Paper I). Since all the vehicles within a platoon will follow the same speed profile using the SPO+ACC method, it can be argued that this method would achieve smaller fuel savings in a *heterogeneous platoon*[1]. This was the main idea for proposing the P-SPO method (Paper III) for platooning; see also Subsect. 2.3.2.

The P-SPO method was introduced and applied to both homogeneous and heterogeneous platoons of trucks over 10 road profiles. For a homogeneous platoon, P-SPO reduced the fuel consumption by 15.8% relative to CC+ACC. Compared to SPO+ACC, P-SPO improved the fuel savings by just 0.2 percentage points (see Table I in Paper III). However, the P-SPO method has several advantages over SPO+ACC; First of all, P-SPO does not require any direct communication between vehicles. The speed profiles must, however, be uploaded to the trucks when the platoon is formed. In the cases where no set of optimized speed profiles is available for a given road, the optimization must be completed before forming the platoon. However, an offline

[1]In Paper III, the heterogeneity was limited to differences in the *masses* of the trucks. However, heterogeneity can of course also refer to differences in the *power trains* of the trucks, something that has been considered in Sect. 4.6.

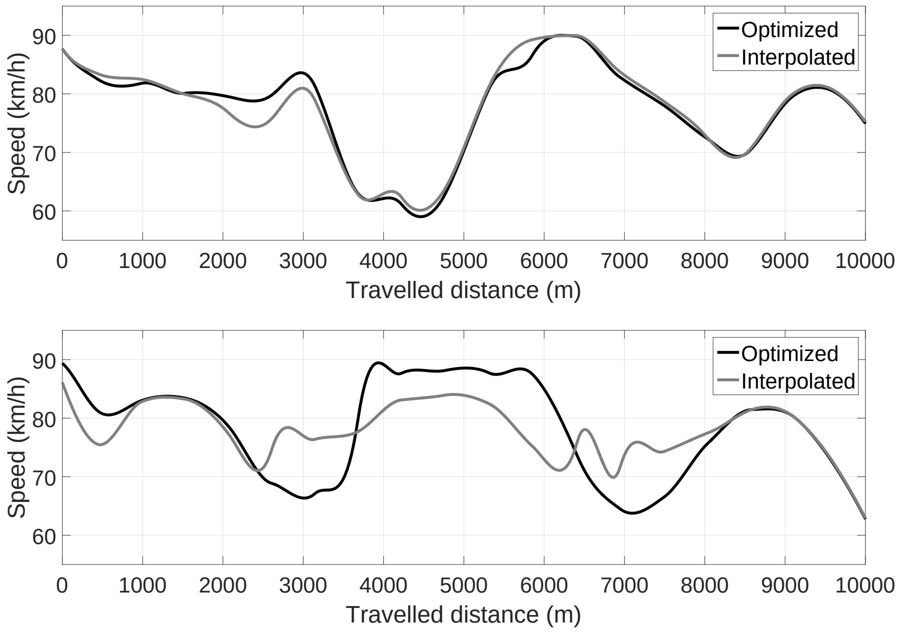

Figure 4.1: *In this figure, two examples of speed profiles generated by interpolating (for a 35-tonne truck) already optimized speed profiles (for a 30-tonne and a 40-tonne truck) are shown. Top panel: An example of successful speed profile interpolation. In this panel, the interpolated speed profile (gray curve) is similar to the optimized profile generated by SPO (black curve) for the 35-tonne truck. Bottom panel: An example of a case where the interpolated speed profile (gray curve) differs quite strongly from the optimized speed profile generated by SPO (black curve).*

database of speed profiles for common vehicle masses could be generated so that the lead vehicle's speed profile can be made available directly. Then, one can use either SPO+ACC which already provides large fuel savings, or run P-SPO starting from the available speed profile to generate a set of speed profiles for the entire platoon. In cases where an optimized speed profile is not available for a specific mass, one could generate a speed profile by interpolating between the available speed profiles. For instance, if the speed profiles for a 30-tonne and a 40-tonne truck are available, a speed profile for a 35-tonne truck could be generated by linear interpolation. However, while this procedure sometimes generates a suitable speed profile, see the upper panel of Figure 4.1, this is not always the case, as shown in the bottom panel of the same figure. Still, the interpolation can provide a good starting point for the optimizer in SPO (and P-SPO) and thus speed up the optimization process.

A second advantage of P-SPO is that it does not require the vehicle to follow a particular spacing policy and therefore has an advantage when driving over steep uphill or downhill segments of a road, where maintaining a constant distance between vehicles will lead to unnecessarily large accelerations or decelerations. Moreover, one should note that during the simulations, the desired inter-vehicle distance in the ACC function was set to the absolute minimum value allowed according to Eq. (3.15), a value that is not practical due to potential failures in the electronic system and sensors. With a more realistic desired inter-vehicle distance (e.g. 15 meters) in the ACC function, the fuel savings of the SPO+ACC drops by 1 percentage point; see the Discussion section of Paper III.

For heterogeneous platoons, the P-SPO method obtained fuel savings of 16.8% to 17.4% (relative to CC+ACC) for different mass configurations. The improvement made by the P-SPO over SPO+ACC is more evident in the case of heterogeneous platoons, see Table II in Paper III, and this can be attributed to two factors; (i) better exploitation of the inter-vehicle distances by not requiring the trucks to follow a specific spacing policy, and (ii) optimization of each truck's speed profile *separately* while considering the safety of the platoon.

4.3 SPO and P-SPO vs. MPC

As discussed in Sect. 2.3 and in Papers II and III, the SPO method's fuel savings (for single trucks) as well as the SPO+ACC and the P-SPO methods' fuel savings (for platoons), exceed those obtained by the methods based on the MPC framework. Apart from the larger fuel savings, using the SPO-based methods has other advantages over MPC: With these methods, there is no need for online iterative re-calculation of the speed profiles; once the optimized speed profiles have been generated, the vehicles can drive accordingly over the entire stretch of the road section.

In recent MPC-based approaches for fuel-efficient driving, two-layer architectures for controllers have been considered where in the first layer the problem of finding a fuel-optimal speed trajectory is solved, assuming that the vehicle is able to track its speed trajectory precisely, and in the second layer the problem of actually tracking the generated speed trajectory is han-

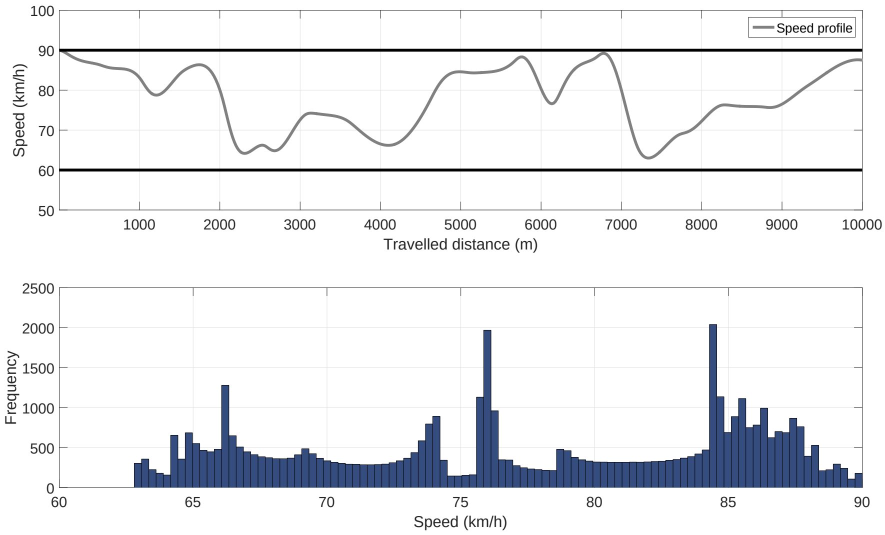

Figure 4.2: *Top panel: Optimized speed profile for a single truck over road profile 5 from Paper II, where the gray curve shows the speed profile and the black lines show the speed limits set during the optimization. Bottom panel: The frequency histogram of the truck's speed measured every 10 ms (in simulation).*

dled. However, in SPO the truck is not required to follow the speed profiles precisely, since the method accounts for the deviation between the desired and actual speed profile when using a simple PID controller. Thus, there is no need for implementing a more advanced controller for tracking the speed profiles and therefore also the need of having a two-layer control architecture similar to the controllers introduced in [27, 21] is eliminated.

In the approaches that are based on the MPC framework, it is common to linearize the model (during the optimization process) to reduce the problem's complexity, something that forces one to use a narrower speed range. In the SPO method, however, no such linearizations are required for solving the optimization problem. This enables SPO to generate speed profiles with larger speed range compared to other methods. In general, for driving on highways, truck drivers prefer to avoid driving at very low speeds (e.g. lower than 60 km/h). As mentioned in Subsect. 3.2.2, the requirements on the average speed, as well as the constraint on the allowed minimum speed, both of which are considered during optimization (see the end of Subsect. 3.2.2),

Table 4.1: *The fuel savings obtained in the simulations for a single truck for different horizon lengths. The second row shows the fuel savings obtained with SPO. The normalization was carried out relative to CC with a set speed of 80 km/h.*

Horizon length (km)	2	5	10	20	40	50
Fuel savings (%)	4.88	9.20	9.24	9.59	9.72	9.91

prevent the truck from reaching low speeds. In fact, in the experiments carried out in Paper II, only around 5% of the time the vehicle's speed was below 65 km/h. An example of a speed profile along with its speed histogram is shown in Figure 4.2. As can be seen from the figure, the majority of the time the truck is driving with a speed higher than 70 km/h.

In general, the fuel savings obtained with the SPO-based methods exceed those obtained with MPC which typically are between 3% and 7% (for single trucks) and between 10% and 13% (for platoons). It should be mentioned that in recent work based on MPC, where a convex approximation of the problem has been considered, e.g. [17, 21], larger speed ranges and, consequently, larger fuel savings (compared to e.g. [15]) were obtained than with earlier MPC-based approaches. However, the reported savings are still below the savings obtained by the SPO-based methods; see also the Discussion sections in Papers II and III.

In the three papers that form this thesis, a horizon length of 10 km has been considered during the optimization. However, one should note that the choice of horizon length is somewhat arbitrary for SPO, as opposed to the MPC-based methods for which the increase in horizon length makes the problem computationally challenging, even in offline computations [21]. Regardless of the method used, the horizon must be long enough to guarantee that the required average speed can be maintained, while leaving enough space for speed variations to reduce the fuel consumption. Here, the choice of a 10 km horizon was motivated partly by the fact that the time required to find suitable speed profiles (see Chapter 3) over such a horizon is typically a few minutes, making it possible, in principle, to optimize speed profiles for the *following* 10 km section, while driving over the current section. Nevertheless, with the SPO method, speed profiles of any length can be generated *a priori* (provided that neither the road nor the truck's characteristics change) such that, during driving, the sole task of the controller is to follow the profile. In fact, by considering longer horizons, or even the entire road, the fuel savings obtained by SPO can be (slightly) improved; see Table 4.1. The results presented in Table 4.1 were obtained by generating fuel-optimal speed profiles over road profiles of different lengths with the same optimization parameters.

4.4 Optimization method

In this thesis, the optimization of speed profiles has been carried out with RMHC (Paper I) and standard GA (Paper II and III). Although it appears that the performance of the standard GA used here is similar to that of the RMHC, in terms of fuel savings, the GA has the advantage of being parallelizable, using GPUs or multiple CPUs for instance, something that can be used for speeding up the optimization process. Nevertheless, both optimization methods are able to obtain a sufficiently good speed profile for the next road section during the time required to drive over the current section.

Two different representations for road profiles and speed profiles were used in this thesis, namely, a simple point list representation and a spline representation. In Paper II, it was shown that by using a more compact representation of the road and speed profile (i.e. the spline representation), the optimization performance can be improved by up to 4 percentage points relative to a case where the simple point list is used. For a 10 km road profile, with the spline representation, each chromosome length dropped from 1000 to the range of 38 to 46 in the GA, thus significantly reducing the size of the search space.

In Subsect. 3.2.2, it was stated that the initial speed profiles, both in the simple GA used in Paper I (i.e. RMHC) and the standard GA used in Papers II and III, were set to flat profiles (as in CC). On the other hand, it is common in GAs to initialize the individuals randomly (within the range of the search space of course). In order to investigate whether starting from random speed profiles improves the fuel savings, several runs were made. Figure 4.3 shows the performance of the GA in reducing the fuel consumption of a single truck both with random initialization of the individuals (gray curves) and flat initialization of the individuals (black curves). As can be seen from the figure, the initialization does not affect the final fuel savings at all; with both initializations the fuel savings obtained were around 13% on average (for 5 road profiles). However, starting from flat speed profiles has

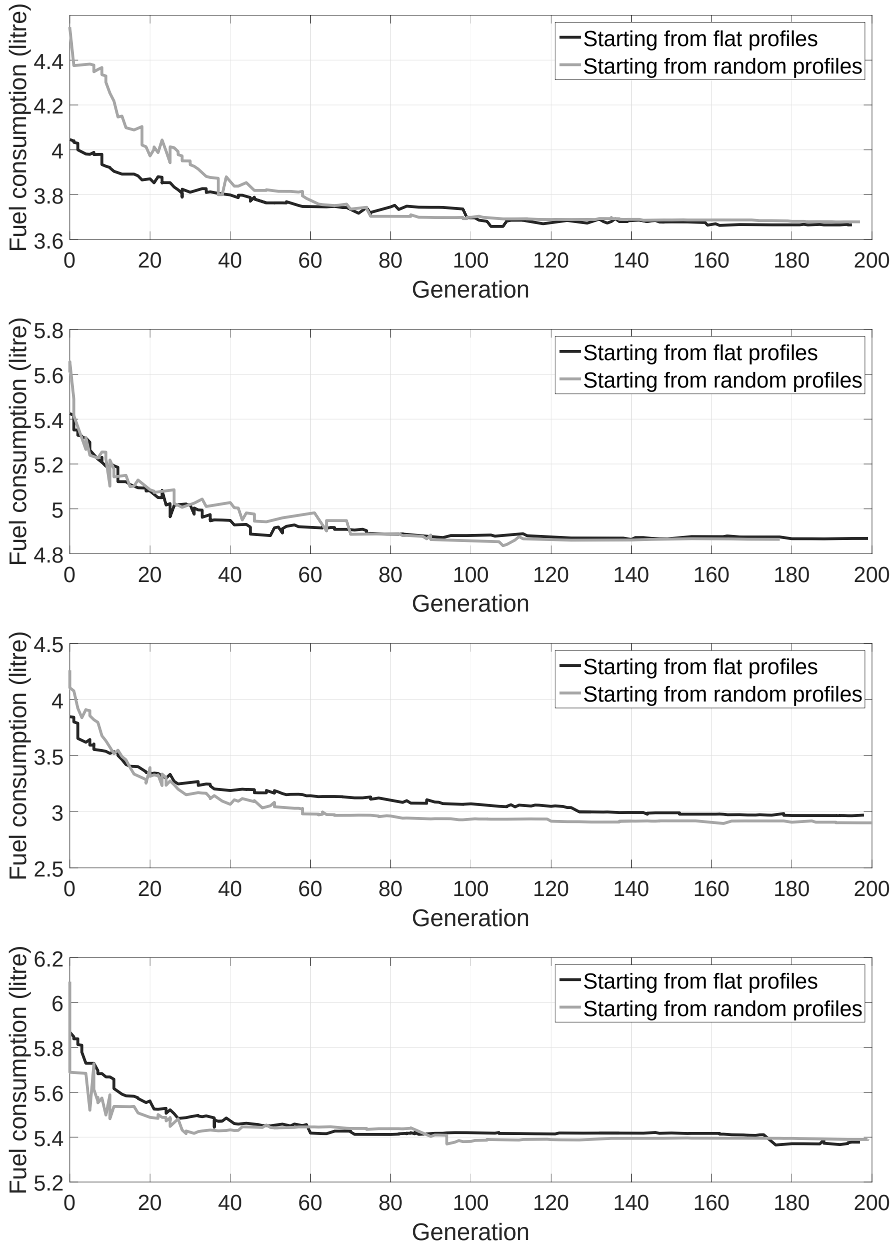

Figure 4.3: *Performance comparison between two initialization strategies for the GA, namely, (i) initialization with flat speed profiles (black curves), and (ii) initialization with random speed profiles (gray curves), for different road profiles.*

a few advantages. For instance, with flat initial speed profiles one can make sure that the current best solution fulfils the requirements and constraints (e.g. speed limit and required average speed) which is useful in cases where, for whatever reason, the GA must be terminated prematurely. Moreover, in the case of platooning (P-SPO), random initialization of the speed profiles could hinder the optimization progress since such profiles are likely to cause collisions within the platoon. However, this problem could be avoided, whether using flat or randomly generated initial profiles, by forcing all trucks to follow the same profile in the early stages of optimization, as indeed has been done in Paper III.

4.5 Handling of external traffic

One of the main assumptions that was made in the SPO method, which is a common assumption in the literature, is the exclusion of the surrounding traffic so that the motion of a truck, or a platoon, is uninterrupted during the simulations. Of course, external traffic can disturb the motion of trucks, regardless of the platooning method used. However, since the SPO method is intended for highway driving and considering that the trucks are typically among the slowest vehicles on a highway, it is unlikely, in a normal traffic situation, that the motion of the trucks would be disturbed frequently by other vehicles. During the experiments with the real truck, carried out in connection with Paper II, it was noted that only in two instances the truck was unable to follow the speed profile due to interference from external traffic, which lasted for 140 seconds in total, over a 100 km drive lasting more than an hour. Moreover, since the SPO method does not require iterative online calculations, and considering that the speed profiles are generated *a priori* and are then available all the time, the truck could simply resume following the speed profile once the disturbance is gone. In Paper II, the effect of external traffic on the fuel savings was investigated (in simulation) by placing a slower vehicle in front of the truck for 30 seconds in every road profile. It turned out the fuel savings of the truck were decreased by around 2 percentage points.

In the platooning scenario where the trucks drive at rather close distances (on average 11 m in the simulations), the external traffic can compromise the safety of the platoon, for example in cut-in situations. In order to handle such situations, the trucks can simply stop following the speed profiles and,

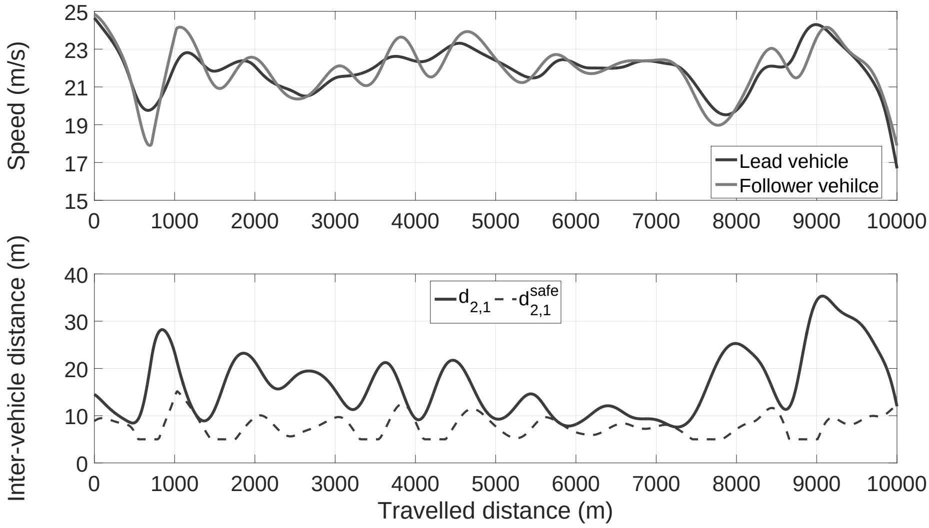

Figure 4.4: *Top panel: One of the optimized set of speed profiles for a heterogeneous platoon of trucks equipped with different power trains. The curves show the lead vehicle's speed profile (black) and the follower vehicle's speed profile (gray). Bottom panel: The inter-vehicle distance (solid line) between the two trucks and the minimum allowed safe distance (dashed line) computed using Eq. (3.15).*

instead, drive according to the ACC function in an attempt to control their distance to the preceding vehicle. Given that the minimum safe distance is around 7 meters on average, according Eq. (3.15), the ACC function has enough time, and space, to control the inter-vehicle distance without braking too harshly. Similar to the case of the single truck, once the disturbance disappears, the trucks can resume following their speed profiles and, possibly, form the platoon again; see also the discussion in Paper III.

4.6 Influence of power train heterogeneity on fuel efficiency

In most of the platooning work to date where heterogeneous platoons have been considered, including Paper III, it is common to assume that the platoon's heterogeneity stems from the difference in the masses of the vehicles. However, another common type of heterogeneous platoon is one in which the vehicles have different *power trains*. In order to investigate the effects of power train differences, the P-SPO method was applied to the 10 road profiles (same as in Papers II and III) using a heterogeneous platoon, where the trucks had different power trains. The P-SPO method obtained fuel savings of 18.2% (on average) over the entire road profiles relative to the baseline case of CC+ACC, showing the ability of the P-SPO method to reduce the fuel consumption also in this case. One of the optimized sets of speed profiles is shown in Figure 4.4. As can be seen from the figure, the speed profiles here differ much more than the case where the platoon members have different masses (see Figures 3 and 4 in Paper III) thus showing the importance of optimizing a speed profile for each vehicle of a platoon rather than assigning identical profiles to all vehicles despite their heterogeneity.

Chapter 5

Conclusion and future work

5.1 Conclusion

The main conclusion that can be drawn from this thesis is that the optimization of speed profiles when developing fuel-efficient driving strategies, and the SPO and P-SPO methods presented in previous chapters, in particular, leads to significant fuel savings both for single trucks and platoons, on typical highways. Specifically, the SPO method obtained fuel savings of 11.5% (on average) for a single truck (see Paper II) relative to the baseline case (CC). Moreover, the P-SPO method achieved fuel savings of 15.8% to 17.4% for a platoon consisting of trucks with different mass configurations (see Papers I and III) relative to the baseline case (CC+ACC).

The results obtained from the SPO method, both for single trucks and platoons, exceed those obtained by the methods described in Chapter 2, and more specifically in Sect. 2.3, by a few percentage points. In fact, in a real truck, the SPO method outperformed a common PCC approach for fuel-efficient driving by 3 percentage points (see Paper II) on the same road with similar settings. Moreover, for platoons of two trucks with different mass configurations, the P-SPO method outperformed an MPC-based approach on the same road, in simulation, by around 3 percentage points for the entire platoon (see Paper III).

Regarding the application of SPO in platooning, it was shown in Paper III that the P-SPO method can improve the SPO+ACC strategy presented in Paper I even though the improvement for homogeneous platoons is not significant. However, the P-SPO method removes the problem of controlling the inter-vehicle distance within a platoon by assigning an optimized speed profile to each truck. Consequently, with the P-SPO method, the trucks are not required to follow any specific spacing policy (in contrast to MPC-based approaches) thus reducing the two-layer architecture commonly used in the MPC framework (where the first layer's task is to generate a fuel-optimal speed profile and the second layer's task is to track that speed profile precisely) to a single problem of finding an optimized speed profile for each vehicle instead.

Another important conclusion from this thesis is the transferability of the SPO method to real trucks. In Paper II, the SPO's performance was evaluated both in simulations and real trucks and it was shown that the SPO method resulted in average fuel savings of 11.5% and 10.2% (relative to CC), in the simulations and the experiments, respectively. The similarity of the obtained results is important and non-trivial since it shows that, despite the simplicity of the simulated truck, the simulations are sufficiently accurate to be transferred to real trucks.

5.2 Future work

The proposed P-SPO method for platooning needs to be tested thoroughly in experiments with real trucks, to validate the results obtained in simulations both from the fuel efficiency and safety aspects. Moreover, it is crucial to test the requirements on the implementation and integration of the P-SPO (and SPO method in general) in commercial trucks, especially if the real-time implementation of the SPO is considered, where the speed profiles are optimized during driving.

In this thesis, it was generally assumed that the external traffic does not interfere with the motion of the trucks when following the speed profiles. However, on real roads and in the presence of other vehicles, it is possible that the motion of a truck will be disrupted from time to time (though typically rather rarely, as mentioned in Sect. 4.5). Although with the SPO method, given that the speed profiles are available *a priori*, the truck could resume following the speed profile once the disturbance is gone, it should be noted that the duration of the disturbance could be longer than expected. Therefore, a decision-making system could be developed for making a decision on whether to overtake the slower vehicle in front or simply keep following the slow vehicle until it is gone. In any case, driving behind a slower vehicle for a long enough time will destroy the coherence of the platoon. Therefore, a decision has to be made whether or not catching up with the preceding vehicle in the platoon (which possibly involves excessive accelerations) can improve the expected fuel savings relative to the alternative option of following the vehicle's own speed profile alone, something that already yields substantial fuel savings; see also the discussion in Paper III.

Finally, another important direction for future research could be the coordination and formation of truck platoons. More specifically, given that goods have different origins, destinations, and time restrictions, it is not trivial to estimate how trucks will benefit from platooning during their individual transport missions. Therefore, developing coordinating systems that could use the available information, such as the potential fuel savings of driving on a specific road as well as the logistic requirements, is an important topic for future work.

Chapter 6

Summary of included papers

This thesis consists of three papers which are concerned with the problem of fuel-efficient driving for heavy-duty vehicles, using speed profile optimization, over roads with varying topography (see Sect. 3.2). In Paper I, speed profile optimization (SPO) is implemented using a simple stochastic optimization method for the lead vehicle of a platoon, while the rest of the vehicles follow the lead vehicle using various platooning algorithms adopted from an artificial physics framework as well as ACC. Paper II is concerned with testing and validating the speed profile optimization method for a single truck, with an improved optimization algorithm, in real trucks. In Paper III, speed profile optimization is implemented for the entire platoon in such a way that each vehicle receives its own speed profile to follow, a method which is referred to as P-SPO.

6.1 Paper I

In this paper, the concept of SPO was introduced for fuel-efficient truck platooning over roads with varying topography, by implementing a very simple stochastic optimization method. The lead vehicle followed the generated fuel-optimal speed profile while the rest of the vehicles in the platoon followed the lead vehicle using various platooning algorithms, such as a non-linear spring-damper model, a modified artificial gravity, and an artificial Lennard-Jones force model as well as the standard adaptive cruise control (ACC) method. The performance of the SPO method together with the platooning algorithms was tested in simulations over 10 road profiles of 10 km length. These approaches were then compared to the baseline case of driving with the combination of cruise control (CC) for the lead vehicle and ACC for the followers, resulting in average fuel savings of around 15% for the entire platoon. Moreover, the average fuel savings of the lead vehicle was 15.8% (on average) compared to driving at constant speed.

6.2 Paper II

The main purpose of Paper II was to test the performance of the SPO method both in simulations and experiments. In this paper, it was proved that the results obtained in simulations are transferable to real trucks despite the simple model used in simulations. Moreover, the SPO method's performance in fuel-efficient driving of a single truck on roads with varying topography was further investigated in this paper. Furthermore, a more direct comparison was made between the SPO method and MPC-based methods for fuel-efficient driving (see Sect. 2.1). The SPO method obtained 11% fuel savings, compared to typical savings of 3 to 7% for MPC-based approaches, in cases where large variations of speed were allowed. In addition, in a more difficult setting with narrower speed range, the SPO method outperformed a standard predictive cruise control (PCC) by around 3 percentage points, over the exact same road and speed range. Moreover, a more advanced genetic algorithm was used (compared to the simple method used in Paper I), and the representations of the speed and road profiles were improved by using composite cubic Bézier curves, rather than using simple lists of road points.

6.3 Paper III

In Paper III, a method for platooning based on SPO, which is referred to as P-SPO, was introduced and evaluated both for homogeneous and heterogeneous platoons. In the latter, the trucks had different masses. The P-SPO method reduced the fuel consumption of a homogeneous platoon by 15.8% (on average), and of heterogeneous platoons by between 16.8% and 17.4% (on average), relative to the baseline case where the lead vehicle used the CC function and the follower vehicle used the ACC function. Moreover, it was shown that the P-SPO method can further improve the platooning method introduced in Paper I, i.e. SPO+ACC, by up to 1.8 percentage points, by assigning different speed profiles to each truck. In this case, the speed profiles were optimized together, while considering the safety of the platoon as a hard constraint during optimization. Furthermore, the P-SPO method outperformed an MPC-based approach for platooning by around 3 percentage points when applied to identical roads, with similar settings. Moreover, with the P-SPO method, the inter-vehicle distance control problem and the two-layer control architecture used in MPC-based approaches are eliminated from the platooning problem.

Bibliography

[1] A. Al Alam, A. Gattami, and K. H. Johansson, *An experimental study on the fuel reduction potential of heavy duty vehicle platooning*, in Intelligent Transportation Systems (ITSC), 2010 13th International IEEE Conference on, IEEE, 2010, pp. 306–311.

[2] A. Alam, J. MÅrtensson, and K. H. Johansson, *Look-ahead cruise control for heavy duty vehicle platooning*, in 16th International IEEE Conference on Intelligent Transportation Systems (ITSC 2013), Oct 2013, pp. 928–935.

[3] L. Alvarez and R. Horowitz, *Safe platooning in automated highway systems*, California Partners for Advanced Transit and Highways (PATH), (1997).

[4] C. Andrieu and G. S. Pierre, *Using statistical models to characterize eco-driving style with an aggregated indicator*, in 2012 IEEE Intelligent Vehicles Symposium, June 2012, pp. 63–68.

[5] R. H. Bartels, J. C. Beatty, and B. A. Barsky, *An Introduction to Splines for Use in Computer Graphics & Geometric Modeling*, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1987.

[6] R. Bellman, *Dynamic Programming*, Princeton University Press, Princeton, NJ, USA, 1 ed., 1957.

[7] C. Bonnet and H. Fritz, *Fuel consumption reduction in a platoon: Experimental results with two electronically coupled trucks at close spacing*, tech. report, SAE Technical Paper, 2000.

[8] K. S. Chang, J. Karl Hedrick, W.-B. Zhang, P. Varaiya, M. Tomizuka, and S. E. Shladover, *Automated highway system experiments in the path program*, Journal of Intelligent Transportation Systems, 1 (1993), pp. 63–87.

[9] European Commission, *EU transport in figures - statistical pocketbook*, 2016.

[10] H. Fritz, *Longitudinal and lateral control of heavy duty trucks for automated vehicle following in mixed traffic: experimental results from the chauffeur project*, in Proceedings of the 1999 IEEE International Conference on Control Applications (Cat. No.99CH36328), vol. 2, 1999, pp. 1348–1352 vol. 2.

[11] A. Fröberg, E. Hellström, and L. Nielsen, *Explicit fuel optimal speed profiles for heavy trucks on a set of topographic road profiles*, tech. report, SAE Technical Paper, 2006.

[12] J. K. Hedrick, D. McMahon, V. Narendran, and D. Swaroop, *Longitudinal vehicle controller design for ivhs systems*, in 1991 American Control Conference, June 1991, pp. 3107–3112.

[13] E. Hellström, J. Åslund, and L. Nielsen, *Design of an efficient algorithm for fuel-optimal look-ahead control*, Control Engineering Practice, 18 (2010), pp. 1318–1327.

[14] E. Hellström, A. Fröberg, and L. Nielsen, *A real-time fuel-optimal cruise controller for heavy trucks using road topography information*, tech. report, SAE Technical Paper, 2006.

[15] E. Hellström, M. Ivarsson, J. Åslund, and L. Nielsen, *Look-ahead control for heavy trucks to minimize trip time and fuel consumption*, Control Engineering Practice, 17 (2009), pp. 245 – 254.

[16] M. Henzler, M. Buchholz, and K. Dietmayer, *Online velocity trajectory planning for manual energy efficient driving of heavy duty vehicles using model predictive control*, in 17th International IEEE Conference on Intelligent Transportation Systems (ITSC), Oct 2014, pp. 1814–1819.

[17] ——, *Optimal parameter selection of a model predictive control algorithm for energy efficient driving of heavy duty vehicles*, in 2015 IEEE Intelligent Vehicles Symposium (IV), June 2015, pp. 742–748.

[18] J. H. Holland, *Adaptation in natural and artificial systems: an introductory analysis with applications to biology, control, and artificial intelligence.*, U Michigan Press, 1975.

[19] W. Hucho, *Aerodynamics of Road Vehicles: From Fluid Mechanics to Vehicle Engineering*, Elsevier Science, 2013.

[20] F. Lattemann, K. Neiss, S. Terwen, and T. Connolly, *The predictive cruise control a system to reduce fuel consumption of heavy duty trucks*, in SAE Technical Paper, SAE International, 10 2004.

[21] N. Murgovski, B. Egardt, and M. Nilsson, *Cooperative energy management of automated vehicles*, Control Engineering Practice, 57 (2016), pp. 84–98.

[22] B. Nemeth and P. Gaspar, *Optimised speed profile design of a vehicle platoon considering road inclinations*, IET Intelligent Transport Systems, 8 (2014), pp. 200–208.

[23] L. Peppard, *String stability of relative-motion pid vehicle control systems*, IEEE Transactions on Automatic Control, 19 (1974), pp. 579–581.

[24] K. Santhanakrishnan and R. Rajamani, *On spacing policies for highway vehicle automation*, IEEE Transactions on Intelligent Transportation Systems, 4 (2003), pp. 198–204.

[25] A. Schwarzkopf and R. Leipnik, *Control of highway vehicles for minimum fuel consumption over varying terrain*, Transportation Research, 11 (1977), pp. 279–286.

[26] D. Swaroop and J. Hedrick, *String stability of interconnected systems*, IEEE transactions on automatic control, 41 (1996), pp. 349–357.

[27] V. Turri, B. Besselink, and K. H. Johansson, *Cooperative look-ahead control for fuel-efficient and safe heavy-duty vehicle platooning*, IEEE Transactions on Control Systems Technology, PP (2016), pp. 1–17.

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

[29] D. Yanakiev and I. Kanellakopoulos, *Variable time headway for string stability of automated heavy-duty vehicles*, in Decision and Control, 1995., Proceedings of the 34th IEEE Conference on, vol. 4, IEEE, 1995, pp. 4077–4081.

[30] ——, *Nonlinear spacing policies for automated heavy-duty vehicles*, IEEE Transactions on Vehicular Technology, 47 (1998), pp. 1365–1377.

[31] J. Zhang and P. A. Ioannou, *Longitudinal control of heavy trucks in mixed traffic: environmental and fuel economy considerations*, IEEE Transactions on Intelligent Transportation Systems, 7 (2006), pp. 92–104.