CHALMERS
UNIVERSITY OF TECHNOLOGY

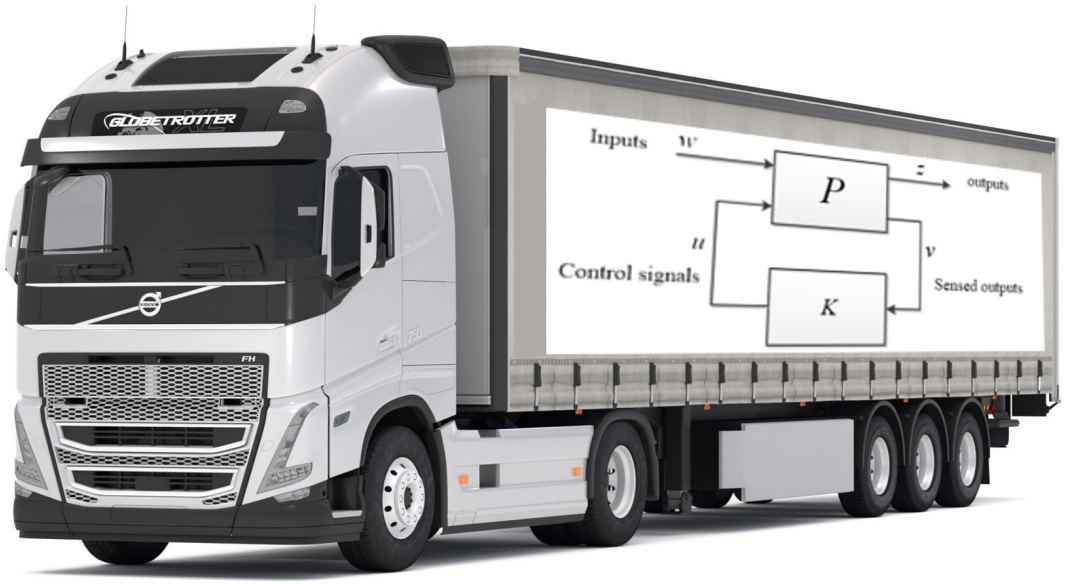

Advanced Thermal Regulation of Muffler

Using MPC and Preview Data

Master's thesis in Systems, Control and Mechatronics (MPSYS)

JOHAN EKENVI & NICLAS PERSSON

DEPARTMENT OF ELECTRICAL ENGINEERING

CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2025
www.chalmers.se

Master's thesis 2025

Advanced Thermal Regulation of Muffler

Using MPC and Preview Data

JOHAN EKENVI
NICLAS PERSSON

CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Electrical Engineering
Division of Systems and control
Chalmers University of Technology
Gothenburg, Sweden 2025

Advanced thermal regulation of muffler
Using MPC and preview data
JOHAN EKENVI, NICLAS PERSSON

Supervisor: Maxim Kristalny PhD, Johan Dahl PhD, Volvo AB
Examiner: Prof. Nikolce Murgovski, Electrical Engineering

Master's Thesis 2025
Department of Electrical Engineering
Systems and Control
Chalmers University of Technology
SE-412 96 Gothenburg
Telephone +46 31 772 1000

Cover: A picture of a truck and control schematic

Typeset in LaTeX
Printed by Chalmers Reproservice
Gothenburg, Sweden 2025

Advanced thermal regulation of muffler using MPC and preview data
JOHAN EKENVI, NICLAS PERSSON
Department of Electrical Engineering
Chalmers University of Technology

Abstract

Tighter legislation on emissions, such as the upcoming EURO-7, puts increased pressure on manufacturers to further decrease emissions from their vehicles. This study address Nitrogen Oxides (NO_x) emission from heavy-duty trucks. For the selective catalytic reduction catalyst, an important component of the exhaust after-treatment system to properly react NO_x into pure nitrogen and water, is that the temperature needs to be sufficiently high in the muffler. During extended periods of low engine load, such as downhills or idling, the muffler cools down as colder exhaust enters the muffler, as a result NO_x emissions increases.

To address this, the study investigates a thermal property of the muffler's dynamics that allows the use of time normalization, which simplifies the dynamics. Time normalization was then utilized in the design of the resulting model predictive controller. Furthermore, the impact of incorporating preview information was investigated to determine what benefits could be gained by using a model predictive controller with preview, compared to one without.

The results demonstrated that it was possible to utilize the structure of the dynamics to construct a controller in a linear time-invariant environment, to control the linear time-varying system. Additionally, incorporating preview information improved the performance in cases where the exhaust temperature drops fell outside the region of which the controller could compensate. In such cases the controller pre-heated the muffler to maintain a sufficient temperature, thus extending the duration of constraint satisfaction but at the cost of increased energy consumption. However, when the temperature drop was within the region of what the controller could supply, the controller without preview performed equally well.

Keywords: Diesel engine, Model predictive control, Time normalization, Exhaust after-treatment system , muffler, Thermal regulation, EURO-7, SCR,

Acknowledgements

First and foremost, we are deeply thankful to our supervisors, Maxim Kristalny and Johan Dahl, for their guidance, insightful feedback, and support throughout the thesis. Without their knowledge and encouragement, this work would not have been possible.

We would also like to thank Prof. Nikolce Murgovski, who served both as examiner and academic supervisor. It was Nikolce who introduced and taught us the topic of model predictive control at Chalmers which awakened our interest in the subject and ultimately led to this thesis.

A special thanks to our colleague Mojtaba Amini-Omam, for his technical support and valuable input throughout the project, particularly during the development, debugging and tuning of the numerical solver.

We are also grateful to Greeshma Koothumparambil, our group manager at Volvo AB, for facilitating this thesis opportunity and for her support in the industrial and practical aspects of the project.

Finally, we extend our appreciation to the entire team at Volvo AB and Chalmers University of Technology for providing the environment, resources, and support that made this thesis possible.

Johan Ekenvi & Niclas Persson

Gothenburg, June 2025

List of Acronyms

Below is the list of acronyms that have been used throughout this thesis listed in alphabetical order:

DOC	Diesel Oxydation Catalyst
DPF	Diesel Particle Filter
EATS	Exhaust Aftertreatment System
EMF	Exhaust Mass Flow
HDV	Heavy-Duty Vehicle
LTI	Linear Time Invariant
LTV	Linear Time Varying
LQ	Linear Quadratic
MPC	Model Predictive Control
NOx	Nitrogen Oxides
SCR	Selective Catalytic Reduction

Nomenclature

Below is the nomenclature of indices, sets, parameters, and variables that have been used throughout this thesis.

Signals and Constraints

$\mathbf{x}$	State vector x
$\dot{\mathbf{x}}(t)$	Derivative of x with respect to t
$x_{\min}, x_{\max}$	Lower and upper limit of x
$\mathbf{u}$	Control input vector
$u_{\min}, u_{\max}$	Lower and upper limit of u
$\mathbf{y}$	Output vector
$y_{\min}, y_{\max}$	Lower and upper limit of y
emf(t)	Exhaust mass flow [Kg/s]
$\mathrm{T_{exh}}(t)$	Exhaust temperature [$^{\circ}C$]
$\mathrm{T_{in}}$	Inlet temperature to muffler, sum of $T_{exh}(t)$ and $u(t)$
T_{lim}	Lower temperature limit of SCR [$^{\circ}C$]

Matrices & vectors

$\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D}$	State-space matrices
$\mathbf{A_n}, \mathbf{B_n}, \mathbf{C_n}, \mathbf{D_n}$	Normalized state-space matrices
$\mathbf{A_d}, \mathbf{B_d}, \mathbf{C_d}, \mathbf{D_d}$	Discrete time state-space matrices
$\mathbf{A_c}, \mathbf{B_c}, \mathbf{C_c}, \mathbf{D_c}$	Controller state-space matrices

Independent variables

t	Natural time, [s]
τ	emf normalized time, [s']

Optimal control symbols

V_N	Cost function over horizon N
$u^*(0:N-1)$	Optimal sequence of future control
u^*	Optimal control at current time
$x^*(0:N)$	Optimal sequence of future states

Constants

α	0.5, Time constant, [1]
$\mathrm{emf}_{\mathrm{nom}}$	0.1, Nominal emf value, $[\mathrm{kg\,s^{-1}}]$
C_{g}	Specific heat capacity muffler section, $[\mathrm{J\,kg^{-1}\,K^{-1}}]$
C_{p}	Specific heat capacity muffler section, $[\mathrm{J\,kg^{-1}\,K^{-1}}]$

MPC variables and parameters

N, M	Prediction and control horizon
ϵ	Slack variable
W_ϵ	Linear slack weight
W_ϵ^2	Quadratic slack weight
W_{u}	Linear control weight
W_{u}^2	Quadratic control weight

Contents

List of Figures

List of Tables

1
Introduction

In this section, we provide an overview of key concepts and challenges related to the thermal control of core components in the Exhaust Aftertreatment (EATS) systems of diesel engines. We begin by discussing the role of the EATS in reducing harmful emissions and the importance of maintaining optimal operating temperatures for effective nitrogen oxide conversion. Next, we introduce the challenges associated with thermal regulation and motivate the use of Model Predictive Control (MPC) as a promising approach, highlighting the potential benefits of utilizing preview information to improve temperature management. We also explore the concept of time-scaling, which simplifies the handling of system dynamics. Finally, the purpose and objectives of this work are presented.

1.1 Background

1.1.1 Exhaust Aftertreatment System (EATS)

The commercial transportation industry is a vital part of society's infrastructure, but the use of fossil fuels leads to undesirable emissions. In addition, regulations, such as the EURO-6 [1] standard and the upcoming EURO-7, impose increasing pressure on manufacturers to adopt alternative fuels or further reduce emissions. The need to comply with constantly tightening emission regulations drives the development of advanced strategies for emission control.

Today's EATS is a complex assembly comprising several components, such as the Diesel Oxidation Catalyst (DOC), the Diesel Particulate Filter (DPF), and the Selective Catalytic Reduction (SCR) catalyst. These components must be carefully coordinated to effectively reduce harmful emissions such as carbon monoxide, particulate matter, and nitrogen oxides (NOx), mainly NO_2 and NO. This work focuses in particular on the SCR, which is responsible for converting NOx. The SCR uses a urea solution (NH_3), commonly known as AdBlue, and heat to convert NOx to pure nitrogen and water.

$$\mathrm{NO_x\ Conversion} : \mathrm{NO_x} + 2\mathrm{NH_3} + \mathrm{HEAT} \Rightarrow 2\mathrm{N_2} + 3\mathrm{H_2O}$$

During normal operating conditions, the conversion is nearly complete, and the muffler is kept warm by a steady supply of heat from the engine. However, certain situations, such as prolonged low engine load while descending hills or idling, can

cause the muffler to cool down. These periods are often followed by a sudden increase in engine load and exhaust flow, but the EATS is not sufficiently warm to convert the NOx. As a result, even with urea injection, conversion is incomplete due to inadequate temperature. These transient conditions are responsible for a significant share of NOx emissions. Low temperatures can also lead to crystallization, which reduces the effective surface area of the SCR and increases back pressure, negatively affecting engine performance [2].
To maintain optimal temperatures, additional heating is therefore sometimes required, either by adjusting combustion parameters or using an electrical heater to increase the gas temperature entering the muffler. However, adding heat increases fuel consumption, making accurate control essential for balancing emission reduction with fuel efficiency.

1.1.2 Thermal Control of EATS

Due to the complex dynamics of both engine combustion and the exhaust aftertreatment system, thermal control presents a significant challenge. Finding an effective control strategy is not trivial and analytical solutions, if they exist, are difficult to derive, and most related work relies on numerical methods, which are more tractable. However, under certain assumptions, analytical solutions can be derived as shown by Holmer [3].

One challenge in muffler thermoregulation arises from the asymmetry in actuation: while heat can be added to the system, there is no mechanism for actively removing it. This leads to a control problem with a hard constraint on the control input, which must remain positive. That said, excessively high temperature is rarely an issue, as component damage typically occurs only at very high temperatures.

Another difficulty lies in the muffler's slow and time-varying thermal dynamics, caused by its large physical dimensions and thermal inertia. By the time the SCR temperature sensor detects a drop in temperature, it is often too late to take corrective action. The significant lag between heat input at the muffler entrance and the resulting temperature change at the SCR motivates the use of preview information. Given the hard constraints on both control and temperature, combined with the desire to incorporate preview capabilities, Model Predictive Control (MPC) emerges as a promising control strategy.

1.1.3 Model Predictive Control (MPC)

Extensive research has already been conducted to minimize fuel consumption and NOx emissions [4, 5, 6], and further improvements continue to emerge. In the previously mentioned papers, various MPCs were implemented, which will also be the focus of this work. MPC is an advanced control method widely used in various industries, such as process control and path planning. More recently, it has gained traction in the automotive sector due to the increase in computational power available in modern vehicles, enabling the real-time optimization which MPC requires.

Furthermore, the thermal dynamics in the EATS are very slow, which creates further incentive for predictive control.

The main advantage of MPC lies in its built-in ability to handle both actuator and state constraints, critical for our application, where the control input and system temperature must remain within strict limits. In addition, MPC naturally allows for the integration of preview information, enabling the controller to anticipate and respond to future disturbances. We will develop and evaluate the performance of an MPC that incorporates preview information for both the exhaust temperature and exhaust mass flow. The use of preview data has demonstrated improved performance in related domains, such as vehicle suspension [7, 8] and transmission control [9]. As real-time computation is a requirement in the automotive and heavy-duty vehicle (HDV) industry, it is desired to develop an efficient formulation, which we in this work achieve by time-scaling.

1.1.4 Time-Scaling

A method used in this thesis is *time-scaling*. Time-scaling refers to the method of redefining the time axis of which a system evolves. Normally, the axis is the physical time, but the system dynamics can be reformulated with respect to another, artificial independent variable, commonly denoted "τ". This method was originally introduced to simplify non-linear dynamics [10]. The main benefit of time-scaling is that complex behaviour appears more tractable or even linear when viewed through this artificial axis, making the system easier to both analyse and control. For instance, in certain batch crystallization processes, time can be normalized with respect to the crystal growth length rather than physical time, leading to more intuitive and manageable system models [11]. Other works where this also has been used is for example in [12] and [13]. Although commonly applied to simplify non-linear dynamics, the method can also be applied to linear time-varying systems to yield a time-invariant representation, which is the approach taken in this work.

1.2 Aim

This work aims to contribute to the ongoing research of minimizing harmful emissions from the transportation industry by the following:

- Find an **efficient** way to formulate the problem of muffler thermoregulation as MPC.
- Develop a framework for muffler thermoregulation that allows utilization of preview information within a moving preview window.
- Analysing the impact of preview information for the thermoregulation problem.

1.3 Limitations and Assumptions

This section aims to bring up the limitations of this work and give a transparent and critical assessment of its content. The limitations include different aspects, such as the chosen method, data selection, theoretical assumptions, and resource limitations.

- **Heat source considerations:** While potential sources of added heat are discussed, the practical constraints and feasibility of these sources in a real-world setting are not analysed in detail.

- **Assumed inputs:** The model assumes that both the exhaust mass flow and exhaust gas temperature are known quantities. The process of estimating or measuring future values of these variables is not within the scope of this work.

- **Model simplifications:** The proposed framework is based on a model that includes several assumptions: ideal heat exchange, no conduction in the substrate, no exothermic or endothermic reactions, and perfectly insulated surfaces. These assumptions may limit the accuracy of the simulation results in more complex or real-world scenarios. However, the model is widely used in industry and validated to sufficiently well reflect a real system.

- **Simulation scope:** The controller is tested exclusively in a simulated environment. As such, its real-world performance may differ, and the results should be viewed as indicative of the potential of the proposed methodology, rather than definitive proof of its effectiveness.

- **Evaluation of emissions:** Although the motivation for controlling the muffler temperature is rooted in improving emission conversion, the controller is evaluated solely based on its thermal performance. Emission-level effects are not assessed, as this would require either a significantly more complex emission model or implementation and testing on a real engine setup with emission measurement capabilities. Furthermore, the trade-off between increased fuel consumption and improved NOx conversion efficiency is not evaluated. Therefore, the overall emission impact of the proposed control strategy remains an open question.

- **Resource constraints:** The work has been conducted under limited time (20 weeks) and computational resources, which influenced the scope of testing and the depth of the analysis.

1.4 Data disclaimer

In this work, several simulations were performed. Due to confidentiality requirements, certain data used in the simulations, such as exhaust temperature, exhaust mass flow, and system constraints have been modified. These modifications do not

affect the validity of the methodology or the conclusions drawn from this work.

1.5 Ethical and sustainability aspects

The commercial transportation industry is one of the main sources of pollution today. A reduced NOx emission caused by commercial trucks would benefit the well-being of all humans since these emission caused several problematic environmental and health effects. For example, NOx in the atmosphere reacts with water and creates acidic rain which harms both animals and humans [14]. When the ground becomes more acidic, vital nutrients are lost which hinders crop and forest growth [15]. NOx is considered poisonous and may irritate the repository tract and mucous membranes [16]. It has also been shown that long exposure to NOx lead to increased symptoms of bronchitis in asthmatic children [17].

Alignment with United Nations Sustainability Goals

This work helps improve the following sustainability goals, as defined by the United Nations:

- **Goal 3 Good Health and Well-Being:**
  Working to decrease NOx emissions will directly benefit this goal, as both human and animal health are negatively impacted by these pollutants. By reducing NOx emissions, this project aims to lower the incidence of respiratory issues and promote a healthier living environment.

- **Goal 6: Clean water and sanitation**
  NOx directly acidifies our sources of water, compromising water quality. This project addresses this goal by working to limit NOx emissions, thereby helping to preserve clean water and protect aquatic ecosystems.

- **Goal 9: Industry, Innovation, and Infrastructure**
  Commercial transportation is a crucial part of global infrastructure and substantial efforts are put into making it more sustainable. While the long-term goal is to replace fossil fuels with sustainable energy sources (such as electricity and hydrogen), in the meantime the harmful emission needs to be minimized to make the current fleet of commercial vehicles more sustainable. This study will be another step in this direction, making commercial vehicles more sustainable by aiming to decrease the emissions of NOx.

Ethical risks

The authors of this work found no ethical risks in this work. However, in future work or when implementing preview information into trucks there arises a privacy concern that require handling of sensitive customer data such as GPS location, which probably will be needed to accurately estimate the exhaust temperature and exhaust mass flow.

2

Modelling and Problem formulation

In this section, the underlying dynamics of the thermodynamic model will be described in detail, which is the foundation for this paper. Furthermore, the problem formulation is presented mathematically to describe the nature of the problem and its constraints.

2.1 Thermal Model

The thermal model of the muffler is inspired by [3]. The system is divided into n identical sections (denoted as "stones"), with the assumption that each section has the same mass and material properties. Each stone propagates the inlet gas further in the system until the last stone (the output). Furthermore, we assume that the temperature of each stone and the temperature of the exiting gas are the same (ideal heat exchange), with no conduction in the substrate, no exothermic or endothermic reactions, and perfectly insulated surfaces. A single-stone's equation is presented below.

$$\mathrm{C_p} m \dot{T}_i = \mathrm{C_g} emf(t) \left(T_{i-1} - T_i\right) \tag{2.1}$$

where $\mathrm{C_p}$ is the specific heat capacity for the ceramic, m is the mass, $\mathrm{C_g}$ specific heat capacity for the gas and the emf(t) is the exhaust mass flow which varies over time making the system time dependent. T_{i-1} is the previous stone's output and temperature and T_i is the current stone's output and temperature. Figure 2.1 gives an overview of the system where each stone depicts a lumped thermal mass.

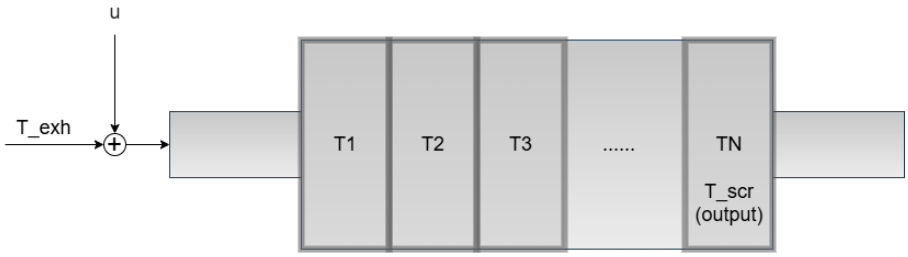

Figure 2.1: Overview of the modelled muffler, showing the stones and T_{exh} as the exhaust temperature, u the control input and T_{scr} the temperature at the SCR (the output).

If the emf is assumed to be constant, an LTI-system is obtained. Consequently, a transfer function can be formulated which yields a first-order system for a single stone.

$$T_i = \frac{1}{\frac{s}{\alpha emf} + 1} T_{i-1} \tag{2.2}$$

where $\alpha \coloneqq \frac{\mathrm{C_g}}{\mathrm{C_p}m}$ for shorter notation. To model the entire muffler we may then chain a single stone's transfer function with the number of stones that the muffler has been divided into and then describe that system as

$$y(s) = \underbrace{\left(\frac{1}{\frac{s}{\alpha emf} + 1}\right)^{i}}_{G(s)} \underbrace{T_{i-1}}_{u(s)} = G(s)u(s) \tag{2.3}$$

Hence an n-order system is obtained. However, the assumption that the emf should be constant is a simplification, as the emf will vary. Therefore, the system needs to be represented in state space form of order n as below.

$$\mathbf{x}_{n\times 1}(t) = \begin{bmatrix} T_1(t) & T_2(t) & \dots T_n(t) \end{bmatrix}^T \tag{2.4}$$

$$\mathbf{A}_{n\times n}(t) = \begin{bmatrix} -\alpha emf(t) & 0 & \dots & 0 \\ \alpha emf(t) & -\alpha emf(t) & 0 & \dots \\ 0 & \alpha emf(t) & -\alpha emf(t) & 0 \\ \vdots & 0 & \ddots & \ddots \end{bmatrix} \tag{2.5}$$

$$\mathbf{b}_{n\times 1}(t) = \begin{bmatrix} \alpha emf(t) \\ 0 \\ \vdots \\ 0 \end{bmatrix} \tag{2.6}$$

$$\mathbf{c}_{1\times n} = \begin{bmatrix} 0 & 0 & \dots & 1 \end{bmatrix} \tag{2.7}$$

Furthermore, the emf(t) enters on all non-zero elements in the $\mathbf{A}$ matrix and $\mathbf{b}$ vector. $\mathbf{A}(t)$ is a sparse matrix with a negative main diagonal and a positive lower off-diagonal, while $\mathbf{b}(\mathbf{t})$ is a column vector with only a positive non-zero value on the

first element. Additionally, only the last state of the n-order system is considered the output from the system as seen in equation (2.7), as the SCR is located at the end of the muffler.

Thus, the model can be described in state space as a n-order linear time-varying (LTV) system as below.

$$\dot{\mathbf{x}}(t) = \mathbf{A}(t)\mathbf{x}(t) + \mathbf{b}(t)u(t) + \mathbf{b}(t)\mathrm{T}_{\mathrm{exh}}(t) \tag{2.8}$$

$$\mathbf{y}(t) = \mathbf{c}\mathbf{x}(t) \tag{2.9}$$

Where u is the control signal, $\mathrm{T}_{\mathrm{exh}}$ is the exhaust temperature, which will act as a disturbance on the system. All other variables are defined as presented in equations (2.4)-(2.7).

By analysing the structure of the system, it can be seen that the system is stable, all eigenvalues will be negative as the **A** matrix has a negative diagonal and a lower triangular structure and emf(t) is strictly positive.

Furthermore, the system has an inherent lag from the input to the output, depending on the emf, this lag can range from 1 minute up to 10 minutes, which makes it challenging to control the output. If the controller is purely reactive to the current output and only applies heat when below the threshold, it will lead to poor performance due to the possible long lag.

2.2 Problem Formulation

In this section, the thermal regulation problem will be shown formally and mathematically to clarify the nature of the problem. The main objective of this thesis is to find an efficient way to keep the muffler and especially the SCR sufficiently warm to facilitate a good conversion of NOx. This can be expressed as

$$y(t) \geq T_{lim}$$

where the limit, T_{lim} comes from the chemical reaction and the properties of the material within the SCR. To ensure this limit is adhered to, we can apply additional heat in the form of $u(t)$, which is bounded by

$$0 \leq u(t) \leq u_{max}$$

The maximum value depends on the actual actuator and is kept generic for now, and the lower is zero as we cannot remove any heat from the system. Applying additional heat increases fuel consumption, which is undesirable and should be minimized. The total amount of added energy applied to the muffler becomes

$$\int_0^\infty mC_{sp}u(t)emf(t)dt$$

Meanwhile, the system dynamics must be adhered to, which is defined by

$$\frac{dx}{dt} = \mathbf{A}(t)x(t) + \mathbf{b}(t)u(t) + \mathbf{b}(t)\mathrm{T}_{\mathrm{exh}}(t)$$

With this, we can mathematically formulate the problem as shown below with the addition of the initial condition and defining the output:

$$min \int_0^\infty mC_p u(t)emf(t)dt \tag{2.10}$$

$$\frac{dx}{dt} = \mathbf{A}(t)x(t) + \mathbf{b}(t)u(t) + \mathbf{b}(t)\mathrm{T}_{\mathrm{exh}}(t) \tag{2.11}$$

$$y(t) = \mathbf{c}x(t) \tag{2.12}$$

$$y(t) \geq T_{lim} \tag{2.13}$$

$$0 \leq u(t) \leq u_{max} \tag{2.14}$$

$$x(0) = x_0 \tag{2.15}$$

3

Time normalization

Transforming the time-axis can be used when analysing non-linear systems, the idea is to make a non-linear system appear to be linear on the new time-axis. This approach is closely related to the concept of *orbital flatness* introduced by M. Fliess[18] where changing the time-axis was used to simplify the analysis of the system. Time normalization can similarly be used to convert an LTV system into an LTI in its new time domain. This idea serves as the foundation for introducing time normalization in this work, as it is significantly easier to construct an MPC for an LTI model than for an LTV.

3.1 Concept of Time Scaling

Consider a standard system with a single natural variable (t), defined as $\frac{d\mathbf{x}}{dt} = \mathbf{Ax} + \mathbf{Bu}$. This system can be transformed to evolve on a new time-axis by introducing a time scaling function.

$$\frac{dt}{d\tau} = s(t, x) \tag{3.1}$$

Equation (3.1) is a function that depends on both time and state, it can be used to transform a system to a different *time*-axis, by substituting the independent variable "t" with the fictitious variable τ. It is important to note that $s(t, x)$ must be chosen such that $0 < s(t, x) < \infty$. Meaning that the new time τ must be monotonically increasing with respect to the natural time t, as the new time should not move backwards [10].

Lastly, the transformation must be bijective to ensure that we can map between natural time and τ-time as well as from τ-time back to natural time, with a one-to-one correspondence in both directions.

By substituting (3.1) into the model we can replace t with τ, hence the normalized system evolves in the fictitious time τ.

$$\frac{dx}{d\tau} = \frac{dx}{dt}\frac{dt}{d\tau} = (\mathbf{Ax} + \mathbf{Bu})s(t, x) \tag{3.2}$$

Notice in the new formulation, the transformation enters on the right side as a multiplier, which often is chosen such that it cancels out undesired behaviours in the

dynamics. The result of this is that the new time-axis will evolve non-uniformly. Below is an example taken from [12] of how this concept can be utilized to simplify a non-linear system into a linear one formed in the new normalized time τ.

Example of time normalization

Consider the two systems A and B which are presented below in state-space form.

System A

$$\frac{d}{dt}\begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} x_2 \\ -x_1 - 0.5x_2 \end{bmatrix}$$

System B

$$\frac{d}{dt}\begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} x_2(1.5 + \sin(16x_1)) \\ (-x_1 - 0.5x_2)(1.5 + \sin(16x_1)) \end{bmatrix}$$

In Figures 3.1 and 3.2 the phase plot of both System A and B can be seen.

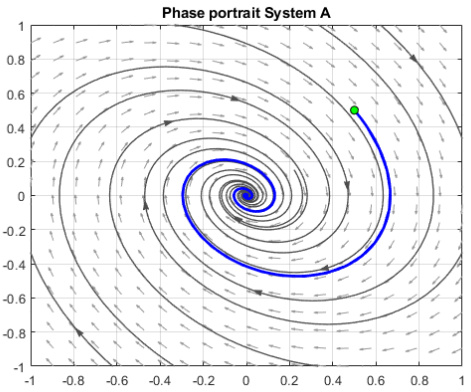

Figure 3.1: Phase portrait of System A, with initial condition [0.5,0.5] highlighted in blue.

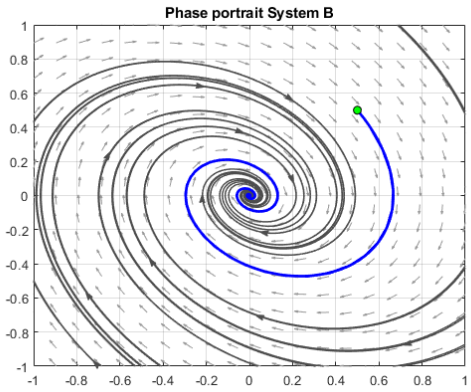

Figure 3.2: Phase portrait of system B, with initial condition [0.5,0.5] highlighted in blue.

Even though the state-space equations differ, the trajectories in the phase planes are equivalent.

So let us introduce the new time scale, denoted as τ. The time transformation function $s(t, x)$ is chosen as:

$$\frac{dt}{d\tau} = \frac{1}{1.5 + \sin(16x_1)}$$

By applying this time transformation to system B, the system can be expressed in terms of the new variable τ.

$$\frac{dx}{d\tau} = \begin{bmatrix} x_2(1.5 + \sin(16x_1)) \\ (-x_1 - 0.5x_2)(1.5 + \sin(16x_1)) \end{bmatrix} \frac{1}{1.5 + \sin(16x_1)}$$

Effectively transforming the system's dynamics, the state-space representation of System B takes the same form as System A:

$$\frac{d}{d\tau} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} x_2 \\ -x_1 - 0.5x_2 \end{bmatrix}$$

Thus given any initial conditions, the trajectory of the states for both System A (in natural time) and System B (in τ-time) will be identical but evolving at different rates, or formally

$$x_A(t) = x_B(\tau(t)) \; \forall \; x_A(0) = x_B(0)$$

3.2 EMF Based Time Normalization

The same concept can be applied to the thermal model presented in Section 2.1, which is linear but time-varying. By time normalization, the system can be represented as an LTI system with a different independent variable. Consider the thermal model developed in Section 2.1, presented below again for convenience.

$$\frac{dx}{dt} = \mathbf{A}(t)\mathbf{x}(t) + \mathbf{b}(t)u(t) + \mathbf{b}(t)\mathrm{T}_{\mathrm{exh}}(t)$$

Notice that the only time-varying variable is emf(t), and it enters the model as a multiplier on all non-zero elements of both the $\mathbf{A}$(t) and $\mathbf{b}$(t) matrix.

$$\frac{dx}{dt} = \mathrm{emf}(t)\left(\mathbf{A_n}\mathbf{x}(t) + \mathbf{b_n}u(t) + \mathbf{b_n}\mathrm{T}_{\mathrm{exh}}(t)\right) \tag{3.3}$$

emf(t) can therefore be extracted. The notation for both $\mathbf{A_n}$ and $\mathbf{b_n}$ indicate that the matrices are now constant as shown below.

$$\mathbf{A_n} = \begin{bmatrix} -\alpha & 0 & \dots & 0 \\ \alpha & -\alpha & 0 & \dots \\ 0 & \alpha & -\alpha & 0 \\ \vdots & 0 & \ddots & \ddots \end{bmatrix} \tag{3.4}$$

$$\mathbf{b_n} = \begin{bmatrix} \alpha \\ 0 \\ \vdots \\ 0 \end{bmatrix} \tag{3.5}$$

$$\tag{3.6}$$

The time-varying emf(t) can now be absorbed into the time axis by introducing a time scaling function containing emf, under the assumption that emf$(t) > 0 \, \forall t$ to satisfy the conditions described earlier.

$$s(t) = \frac{dt}{d\tau} \coloneqq \frac{\text{emf}_{\text{nom}}}{\text{emf}(t)} \tag{3.7}$$

Where emf_{nom} is added to closer reflect to the original system, meaning when emf(t) is the nominal value $s(t) = 1$ the systems evolves at the same rate. By moving the system to the τ-domain, the time-varying dynamics is removed.

$$\frac{dx}{d\tau} = \frac{dx}{dt}\frac{dt}{d\tau} \tag{3.8}$$

$$\frac{dx}{d\tau} = \frac{\text{emf}_{\text{nom}}}{\cancel{\text{emf}(t)}}\cancel{\text{emf}(t)}\left(\mathbf{A_n}\mathbf{x}(t) + \mathbf{b_n}u(t) + \mathbf{b_n}\text{T}_{\text{exh}}(t)\right) \tag{3.9}$$

$$\frac{dx}{d\tau} = \text{emf}_{\text{nom}}\left(\mathbf{A_n}\mathbf{x}(\tau) + \mathbf{b_n}u(\tau) + \mathbf{b_n}\text{T}_{\text{exh}}(\tau)\right) \tag{3.10}$$

Under the assumption that the derivatives are smooth functions the chain rule can be utilized and we can from there derive the system dynamics in τ-domain. The resulting system in τ becomes an LTI system, as seen in Equation (3.10). In a sense, the result will be that τ compresses or stretches time depending on emf. Large emf will cause time to move "faster" while small emf will make it "slower".

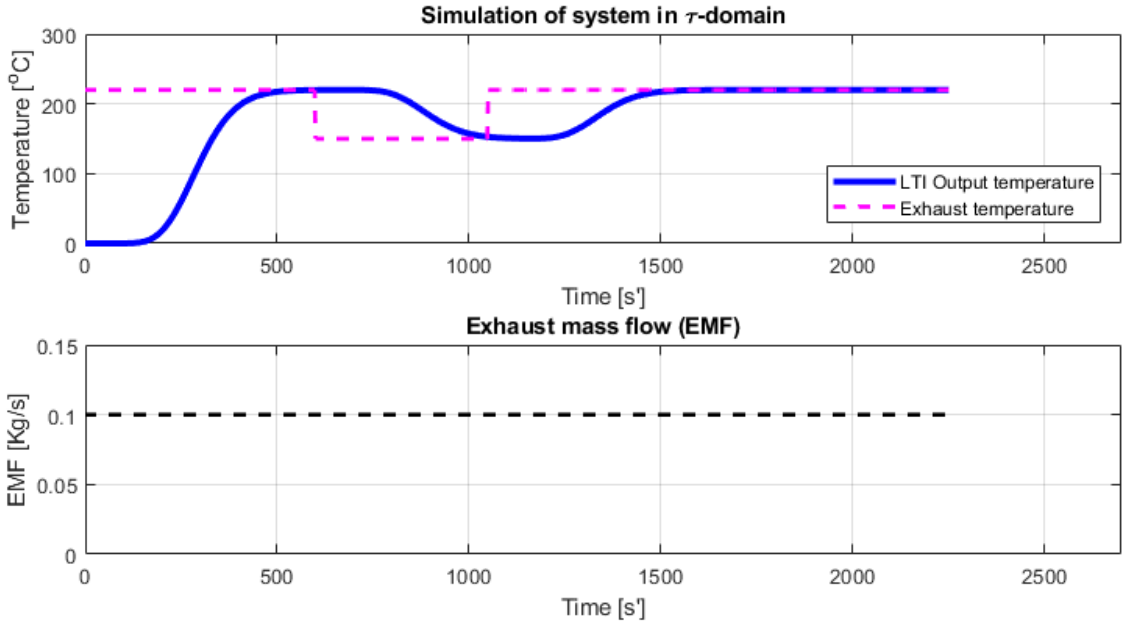

Figure 3.3: Simulation of time normalized system during a cooldown period.

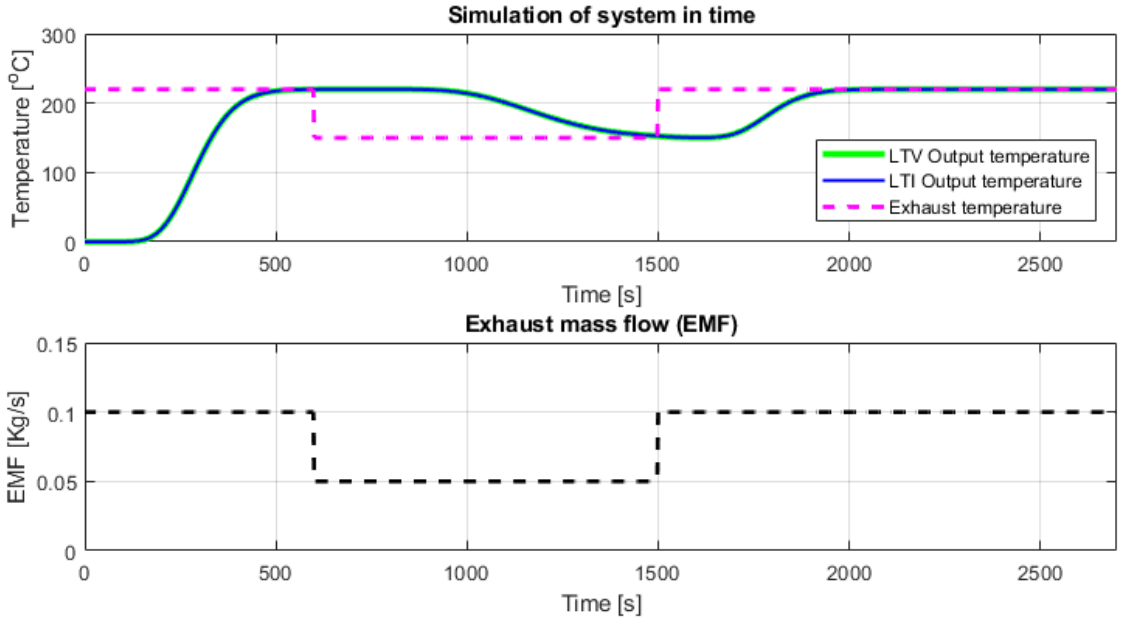

Figure 3.4: Simulation of system during a cooldown period.

The system is then simulated both in τ-time and natural time, where the system in τ-time is an LTI and in natural time LTV. In Figure 3.3 the simulation in τ-domain can be seen and in Figure 3.4 the simulation in natural time with time-varying dynamics and the same output as the previous figure but artificially plotted against natural time. It is evident that the output from the simulation in τ is compressed during the period of low emf but unchanged when emf is nominal.

3.3 Application of Closed-Loop Design

Time normalization can be used to simplify a plant, synthesize a controller, and then transform both back to the original time domain. This is done using the time scaling function to move between the normalized time domain (τ-domain) and the actual time domain (t-domain). Throughout this section, we assume $\text{emf}_{\text{nom}} = 1$ for notational simplicity. In this section, it will be shown that the stability and performance of an LTI-controller synthesized in tau will be preserved when transformed back to natural time.
Plant model:

$$\frac{dx}{d\tau} = \mathbf{A_n x} + \mathbf{B_n} u + \mathbf{B}_\omega \omega \tag{3.11}$$
$$y = \mathbf{C_n x} + \mathbf{D}_\omega \omega \tag{3.12}$$

Which describe the plant dynamics in τ-domain. Consider a general LTI controller in τ-domain can be defined using controller states.

Controller model:

$$\frac{d}{d\tau} x_c = \mathbf{A_c x_c} + \mathbf{B_c} y \tag{3.13}$$
$$u = \mathbf{C_c x_c} + \mathbf{D_c} y \tag{3.14}$$

We may substitute and get 4 equations which will describe the closed-loop system

$$\frac{dx}{d\tau} = \mathbf{A_n x} + \mathbf{B_n C_c x_c} + \mathbf{B_n D_c C_n x} + \mathbf{B_n D_c D}_\omega \omega + \mathbf{B}_\omega \omega \tag{3.15}$$
$$\frac{dx_c}{d\tau} = \mathbf{A_c x_c} + \mathbf{B_c C_n x} + \mathbf{B_c D}_\omega \omega \tag{3.16}$$
$$y = \mathbf{C_n x} + \mathbf{D}_\omega \omega \tag{3.17}$$
$$u = \mathbf{C_c x_c} + \mathbf{D_c C x} + \mathbf{D_c D}_\omega \omega \tag{3.18}$$

Which can be put to state-space form as

$$\frac{d}{d\tau}\begin{bmatrix} x \\ x_c \end{bmatrix} = \begin{bmatrix} \mathbf{A_n} + \mathbf{B_n D_c C_n} & \mathbf{B_n C_c} \\ \mathbf{B_c C_n} & \mathbf{A_c} \end{bmatrix}\begin{bmatrix} x \\ x_c \end{bmatrix} + \begin{bmatrix} \mathbf{B_n D_c D_w} + \mathbf{B}_\omega \\ \mathbf{B_c D}_\omega \end{bmatrix}\omega(t) \tag{3.19}$$
$$\begin{bmatrix} y \\ u \end{bmatrix} = \begin{bmatrix} \mathbf{C_n} & \mathbf{0} \\ \mathbf{D_c C_n} & \mathbf{C_c} \end{bmatrix}\begin{bmatrix} x \\ x_c \end{bmatrix} + \begin{bmatrix} \mathbf{D}_\omega \\ \mathbf{D_c D}\omega \end{bmatrix}\omega(t) \tag{3.20}$$

The inverse transformation can then be used on the system in τ-domain, yielding the resulting system in natural time.

$$\frac{d\tau}{dt} = \text{emf}(t) \tag{3.21}$$

$$\frac{d\mathbf{x}}{dt} = \frac{d\mathbf{x}}{d\tau}\frac{d\tau}{dt} \tag{3.22}$$

$$\frac{d}{dt}\begin{bmatrix} x \\ x_c \end{bmatrix} = \text{emf}(t) \cdot \left(\begin{bmatrix} \mathbf{A_n} + \mathbf{B_n D_c C_n} & \mathbf{B_n C_c} \\ \mathbf{B_c C_n} & \mathbf{A_c} \end{bmatrix} \begin{bmatrix} x \\ x_c \end{bmatrix} + \begin{bmatrix} \mathbf{B_n D_c D_w} + \mathbf{B_w} \\ \mathbf{B_c D_\omega} \end{bmatrix} \omega(t) \right) \tag{3.23}$$

$$\begin{bmatrix} y \\ u \end{bmatrix} = \begin{bmatrix} \mathbf{C_n} & 0 \\ \mathbf{D_c C_n} & \mathbf{C_c} \end{bmatrix} \begin{bmatrix} x \\ x_c \end{bmatrix} + \begin{bmatrix} \mathbf{D}_\omega \\ \mathbf{D_c D}\omega \end{bmatrix} \omega(t) \tag{3.24}$$

It can be seen from the above equations that a controller constructed in τ-domain will also be scaled properly in the time-domain. The closed-loop structure is unchanged between t and τ. The only difference is the scaling factor, emf(t), meaning that performance and stability are preserved, only stretched or compressed in time. As such it is meaningful to construct a controller in τ and then map it back and use it in natural time.

In this thesis, we aim to use an MPC, since MPC is not an LTI-controller it cannot be expressed in state-space form. Thus, it is not possible to apply the same formulation as presented above. However, as will be presented in a later chapter, the principles discussed above are indeed applicable to the MPC as well.

3.4 Problem Formulation in Normalized Time

As previously demonstrated, time normalization simplifies the thermal control problem by transforming the LTV system into an LTI system in τ-domain. This transformation alters the problem formulation and its parts. First, the objective function is transformed,

$$s(t) = \frac{dt}{d\tau} = \frac{\text{emf}_{\text{nom}}}{\text{emf}(t)} \implies dt = \frac{\text{emf}_{\text{nom}}}{\text{emf}(t)} d\tau \tag{3.25}$$

$$V_N(u) \coloneqq \int_0^\infty mC_g u(t)\text{emf}(t)dt \tag{3.26}$$

Which gives the corresponding objective function in normalized time,

$$V_N = \int_0^\infty mC_g \text{emf}_{\text{nom}} u(\tau)\, d\tau. \tag{3.27}$$

As one can see, the objective function simplifies, as the time-varying emf(t) is replaced by the constant emf_{nom}. In a similar manner, the system dynamics are transformed as described in Section 3.2, where the $\mathbf{A}$ and $\mathbf{B}$ matrices are no longer

time-varying.

$$\frac{dx}{d\tau} = \mathrm{emf}_{\mathrm{nom}}(\mathbf{A_n}x(\tau) + \mathbf{b_n}u(\tau) + \mathbf{b_n}\mathrm{T}_{\mathrm{exh}}(\tau)) \tag{3.28}$$

$$y(\tau) = \mathbf{C}x(\tau) \tag{3.29}$$

$$\tag{3.30}$$

Finally, the constraint remains the same as they are amplitude invariant and the transformation does not alter them.

$$\begin{aligned} y(\tau) &\geq T_{lim} \\ 0 \leq u(\tau) &\leq u_{max} \\ x(0) &= x_0 \end{aligned}$$

4

Model Predictive Control

In this chapter, we evaluate two different approaches to implementing an MPC for the thermal regulation of the muffler, one traditional implementation and one utilizing emf based time normalization as described in Chapter 3. MPC is well suited for our application due to the constraint-oriented nature of SCR thermal regulation and the inherent constraint handling of MPC.

4.1 Introduction to Model Predictive Control

MPC relies on an internal model of the system it aims to control and operates by repeatedly solving an optimization problem over a finite prediction horizon. At each time step, the controller minimizes a cost function subject to system dynamics and constraints, applies only the first control action, and then re-optimizes at the next step. One of the key advantages of MPC - compared to traditional control methods is its explicit handling of constraints on both the system states and control inputs [19].

The foundation of MPC is the internal model of the dynamics, describing how the states interact and how they are affected by disturbances and control actions. Today's available MPC toolboxes support various models, including LTI, LTV, and non-linear models. However, most toolboxes rely on discrete models, and therefore the dynamics must be discretized.

As MPC performs an optimization, a cost function (V) must be defined. This function can vary depending on the application, but common terms to penalize via the cost function are control action and state deviation from reference. The solver then tries to minimize the cost function by means of control action, u, which is a free variable in the optimization. Each optimization returns an optimal sequence of control input and state trajectory (x^*) where the first control action of the sequence is the optimal action to take at the current time step, denoted as u^*. The length of these vectors depends on the prediction horizon, which is a tunable parameter.

$$u(0:N-1) = \{u(0), u(1), \ldots, u(N-1)\} \tag{4.1}$$
$$x^*(0:N-1) = \{x(0), x(1), \ldots, x(N-1)\} \tag{4.2}$$
$$u^* = u(0) \tag{4.3}$$

4.2 MPC formulation in natural time

In this section we try to formulate an MPC for the thermal regulation problem described in Section 2.1 in natural time.

4.2.1 Discretization

We first need to discretize the system from a continuous time system into a discrete time system. The simplest discrete time approximation is the zero-order-hold discretization (ZOH) illustrated in Figure 4.1 below.

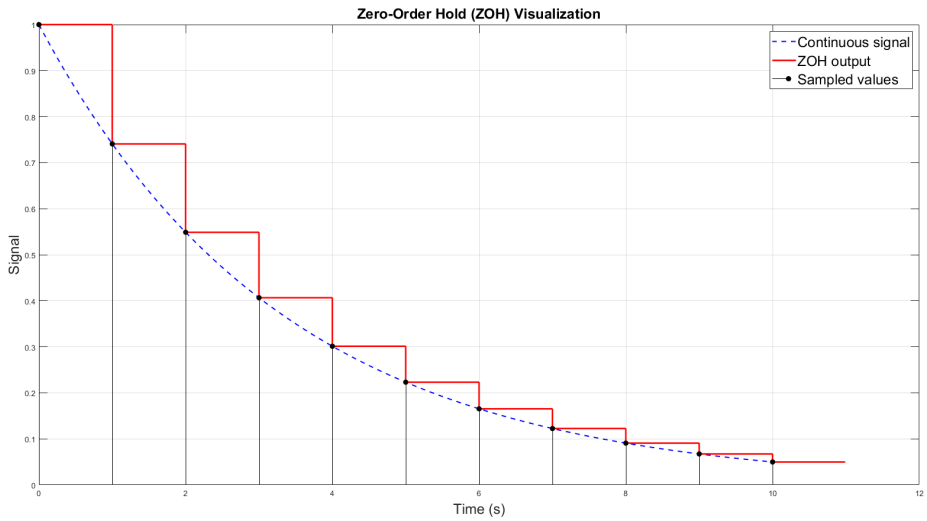

Figure 4.1: Illustration of zero-order hold.

The assumption ZOH relies on is that the signals and dynamics remain constant within each sampling interval. Discrete signals generated by ZOH are equivalent to the continuous time signal at each sampling point from which the digital signal is generated, as illustrated in Figure 4.1 above. A drawback of ZOH is that it may introduce approximation errors if the signal changes rapidly, this can readily be seen in Figure 4.1 by comparing the left-most discretization step to the right-most.

For an LTI-system, ZOH is well defined and there exists an analytical solution that can be used.

Consider a continuous model which we wish to discretize.

$$\dot{x} = \mathbf{A}\mathbf{x}(t) + \mathbf{B}u(t)$$

The derivation for ZOH's analytical solution relies on the use of the derivative of the matrix exponential as

$$\frac{d}{dt}e^{\mathbf{A}t} = \mathbf{A}e^{\mathbf{A}t} \tag{4.4}$$

However, as our system is LTV we in fact have $\dot{x} = \mathbf{A}(t)\mathbf{x}(t) + \mathbf{B}(t)u(t)$. As such we must consider that, unlike in the time-invariant case, the matrix exponential does not behave in the same straightforward manner. Specifically, we must note that:

$$\frac{d}{dt}e^{\mathbf{A}(t)t} \neq \mathbf{A}(t)e^{\mathbf{A}(t)t} \tag{4.5}$$

Continuing the derivation from Equation (4.5) without making an assumption on the emf(t) is impossible. To continue the derivation, we can assume that the emf(t) is piece-wise constant, then denoted as $emf(k)$, we can now use Equation (4.4). From here it is then straightforward as the derivation follows the LTI-case, eventually the ZOH's analytical solution falls out as

$$x(k+1) = \underbrace{e^{emf(k)\mathbf{A_n}\Delta t}}_{\mathbf{A_d}} x(k) + \underbrace{\int_0^{\Delta t} e^{emf(s)\mathbf{A_n}s} ds\, emf(s)\mathbf{B_n}}_{\mathbf{B_d}} u(s) \tag{4.6}$$

Which can be used to discretize our system. Note that now both the emf(t) and control are assumed piecewise constant as

$$emf(s) = emf(k\Delta t),\ k\Delta t \leq s < (k+1)\Delta t$$
$$u(s) = u(k\Delta t),\ k\Delta t \leq s < (k+1)\Delta t$$

This discretization relies on one problematic assumption. It assumes emf(t) is constant during each sampling interval, which can be a good approximation with a small interval and a slowly varying emf(t) . However, emf(t) changes quite fast, requiring a small sampling interval which is computationally expensive or the model becomes less representative of the true dynamics(illustrated generically in Figure 4.1).

4.2.2 Constraints

The constraints are the most important part of the problem formulation as they are the main goal of thermal regulation. There are two inequality constraints, one on the output and one on the input which reflect the nature of the underlying problem. For now, they are kept general but will in later chapters be given numerical values.

$$y(k) \geq T_{lim} \tag{4.7}$$
$$0 \leq u(k) \leq u_{max} \tag{4.8}$$

4.2.3 Cost Function

In typical MPC formulations, the primary control objective is often reflected directly in the cost function. However, in this case, the main objective is instead enforced through a hard constraint, while the cost function tries to minimize the added thermal energy. This formulation is non-standard and reflects the unique nature of thermal regulation in the SCR. The thermal energy is expressed as a linear function of both the emf and the control input u, as shown below.

$$\min_{u(0:N-1)} V_N \coloneqq \sum_{k=1}^{N} mC_g u(k) emf(k) \tag{4.9}$$

An important observation here is that the time-varying emf is part of the thermal energy, meaning that the cost is also time-varying.

4.2.4 Practical Considerations For MPC

In the previous sections, the common steps for formulating an MPC were shown for the LTV system. Implementing an MPC directly for the LTV system revealed several problematic areas, such as the discretization and cost function. In addition to these problems, one additional issue arises in the tuning process of the MPC, and that is the choice of prediction horizon.

The prediction horizon, as described in Section 4.1, specifies the time window for which the solver optimizes. It is important that the time window is sufficiently long for the controller to see the effect of applied control action. Otherwise, the solver will not be able to see that adding extra control improves constraint performance therefore instead apply minimal control action. The required prediction horizon can therefore be formulated:

$$T_{lag}(\text{emf}(t)) \leq NT_s \tag{4.10}$$

Where $T_{lag}(\text{emf}(t))$ is the closed-loop response time, which depends on the emf(t), N is the length of the prediction horizon and T_s is the length of a single time step. The issue, once again is that the emf(t) has a great impact on the response time, meaning that at times the system can respond very fast to an input but also very slow another time. One could assume the worst-case scenario and choose a very long prediction horizon accordingly, but this would increase the computational load drastically. Another alternative is to choose a horizon based on the nominal emf, but that would lead to unpredicted control behaviour when the emf(t) is low and the response is longer than the horizon.

4.3 MPC Formulation Time-Normalized System

Using the time-scaling method described in Chapter 3, the system simplifies to an LTI system. This results in a trivial MPC formulation which will be shown below.

4.3.1 Discretization

The discretization becomes trivial after time-normalization is applied since the normalized system is an LTI system and is not affected by the varying emf(t). Thus the ZOH's analytical solution can be used directly to discretize the system.

$$x(t+\Delta t) = e^{\mathbf{A_n}\Delta t}x(t) + \int_o^{\Delta t} e^{\mathbf{A_n}s}\mathbf{B_n}u(s)ds, \text{where } e^{\mathbf{A}} = \sum_{k=0}^{\infty}\frac{1}{k!}\mathbf{A}^k \tag{4.11}$$

Assuming a piecewise constant control signal

$$\begin{aligned} u(s) &= u(k\Delta t),\ k\Delta t \le s < (k+1)\Delta t, \\ x(k+1) &= \underbrace{e^{\mathbf{A_n}\Delta t}}_{\mathbf{A_d}}x(k) + \underbrace{\int_o^{\Delta t} e^{\mathbf{A_n}s}ds\ \mathbf{B_n}}_{\mathbf{B_d}}u(s) \end{aligned} \tag{4.12}$$

Hence no assumptions on system dynamics needs to be made.

4.3.2 Cost Function in Normalized Time

The cost function also simplifies and is no longer dependent on emf(t) which was the case in Section 3.4.

$$\underset{u(0:N-1)}{min}\ V_N \coloneqq \sum_{k=1}^{N} mC_g u(k)\text{emf}_{\text{nom}} \tag{4.13}$$

Where the minimization is with respect to the sequence of control inputs as before.

$$u(0:N-1) = \{u(0), u(1), \ldots, u(N-1)\} \tag{4.14}$$

$$\tag{4.15}$$

4.3.3 Prediction

The prediction horizon should still be chosen to exceed the closed-loop response as discussed in Section 4.2.4. However, this choice is now simplified as the closed-loop response time is constant in τ.

$$T_{lag} \le NT_s \tag{4.16}$$

4.3.4 Equivalent Trajectory in Time and Tau with MPC

In Section 3.3 it was shown that an LTI controller can effectively be constructed in the normalized τ-domain and scaled to natural time and ensure equivalent performance and stability. In this section, we extend this by proving that the same applies to an MPC. The proposed framework, utilizing emf-based time-normalization is powerful, but the drawback is that the controller becomes abstract, its relation to reality and how to apply the given control action must therefore be made clear.

Below we prove that by transforming the control signal via the integral of emf and calling the optimizer with a **non-uniform** call-rate, equivalent trajectory in natural time as in τ is guaranteed.

Remember, $\text{emf}(t) > 0$, defines the time-normalization

$$\tau(t) = \int_{t_0}^{t_{end}} \text{emf}(t)\, dt \tag{4.17}$$

$\mathbf{A_n}$ and $\mathbf{b_n}$ are the constant matrices, as defined in (3.4)-(3.5), and consider the LTV system

$$\frac{dx}{dt} = \text{emf}(t)\left(\mathbf{A_n}\mathbf{x}(t) + \mathbf{b_n}u(t)\right), \quad x(0) = x_0. \tag{4.18}$$

Define the normalized-time state $\tilde{x}(\tau)$ as the solution to the LTI system.

$$\frac{d\tilde{x}}{d\tau} = \mathbf{A_n}\tilde{\mathbf{x}}(\tau) + \mathbf{b_n}\tilde{u}(\tau), \quad \tilde{x}(0) = x(0) = x_0. \tag{4.19}$$

The control inputs are related by

$$u(t) := \tilde{u}(\tau(t)), \tag{4.20}$$

where $\tilde{u}(\tau)$ is piecewise constant over unit-length intervals:

$$\tilde{u}(\tau) = u_k \quad \text{for } \tau \in [k, k+1), \quad k \in \mathbb{N}. \tag{4.21}$$

Then the solution to the LTV system (4.18) satisfies

$$\mathbf{x}(t) = \tilde{\mathbf{x}}(\tau(t)) \quad \forall\, t \geq t_0. \tag{4.22}$$

In particular, applying a control sequence $\{u_k\}$ for real-time intervals defined implicitly by

$$\int_{t_k}^{t_{k+1}} \text{emf}(t)\, dt = 1 \tag{4.23}$$

yields an equivalent trajectory as the normalized-time LTI.

Remark 1. *The non-uniformity of the control durations in real-time is implicitly defined by the mapping $\tau(t)$. While $\tilde{u}(\tau)$ is held constant for fixed intervals in normalized time, the corresponding $u(t)$ is held for intervals $[t_k, t_{k+1})$ of varying length determined by emf(t). This construction ensures that the two systems evolve identically under the same control sequence.*

As a simple example, imagine the emf-normalized system states evolving at **twice** the speed of the LTV. Intuitively, one can imagine applying the control action on the emf-normalized system for **half** the amount of time as when applying it to the LTV would result in equivalent trajectories. This is essentially what this time-varying call rate to the optimizer does. The final implementation for the LTV can be summarized as follows,

$$u(t) = \begin{cases} \bar{u}(\tau_k), & \text{if } \int_{t_0}^{t_k} \text{emf}(t)\, dt \geq \tau_k \\ u_{prev}, & \text{otherwise} \end{cases} \tag{4.24}$$

4.4 Preview information

A key objective of this work was to evaluate whether the preview of external signals can improve thermal control of the exhaust after-treatment system. To explore this, preview capabilities are implemented within the MPC framework under idealized conditions, assuming perfect knowledge of future disturbances over the prediction horizon. This setup provides a first step for assessing the potential value of preview. The following subsections describe how preview is used in this context and the assumptions made in its absence.

4.4.1 Utilizing Preview Information

In this work, the exhaust temperature is used as the preview signal, as it is the main driver of the system's dynamics. Having information about future temperatures and upcoming cold periods may therefore be beneficial.

As a first step in evaluating the impact of exhaust temperature preview, we assume perfect knowledge within a sliding window (the MPC prediction horizon). While not practically realistic, this assumption allows for an initial investigation into the effectiveness of preview. If the results prove promising, they create an incentive for modelling and estimating the preview.

One important remark is that when using time normalization based on exhaust mass flow, is that preview signals must also be mapped to the normalized time domain. This requires a perfect preview of emf as well, since the transformation to the τ-domain depends on emf.

4.4.2 Assumption on Disturbance When Not Using Preview

Since the MPC should be able to operate without preview information, the assumption of disturbance in the absence of preview information becomes an important consideration. Several plausible assumptions can be made, each with different implications for the controller's behaviour and performance. One straightforward approach is to assume that the current exhaust temperature is constant for the entire prediction horizon.

$$\hat{d}(k+i) = d(k), \quad \forall i = 1, 2, \ldots, N \tag{4.25}$$

This assumption is simple to implement but can lead to reactive control actions depending on how the exhaust temperature changes throughout run-time. For instance, a brief drop in the exhaust temperature might result in the controller unnecessarily engaging actuators to compensate, despite the disturbance being short and momentary.

An alternative assumption is to take a conservative approach and assume that the exhaust temperature will remain below the temperature constraint throughout the prediction horizon.

$$\hat{d}(k+i) = y_{\text{lim}} - C, \quad \forall i = 1, 2, \ldots, N \tag{4.26}$$

Where C is an arbitrary constant set to get the desired offset. While this might result in a controller that will remain above the constraint at most times, it will also yield a controller that is cautious, leading to poor energy efficiency and suboptimal performance.

Another method is a sliding mean of previously observed temperature values and assuming that the mean will be the exhaust temperature through the prediction horizon.

$$\hat{d}(k+i) = \sum_{j=0}^{M-1} w_j d(k-j), \quad \forall i = 1, 2, \ldots, N \tag{4.27}$$

This method could be implemented in different ways, for example by applying higher weights to recently observed values to emphasize that more recent values are more relevant or no weights at all could be used. In the case of using weights, one could use different types of weights to emphasize how fast the relevance should decay, using linear or exponential decay. However, how well this performs depends on the window size and the nature of the temperature variation.

The aforementioned assumptions have their own compromises, naturally there exist more choices on disturbance assumptions than what has been presented above. The choice of disturbance assumption should, in the end, be guided by the application's requirements and available computational resources. Since none of the presented methods is guaranteed to have a better, general performance than the other, we implemented the first mentioned one (4.25). This assumption is easy to implement and neither optimistic nor conservative.

4.5 Optimization Solver: ODYS Embedded MPC

For this work, ODYS Embedded MPC [20] was utilized as the numerical solver and tool to evaluate our framework and the impact of utilizing preview information. The

toolbox fulfils the requirements for this project which were the following:

- Fast, real-time MPC solver
- Embedded deployment capability
- Offer debugging and post-process analysis

Solving the optimization fast is vital for automotive and HDV applications to ensure good performance, especially in our application due to the slow system response time leading to a long prediction horizon. ODYS is both very fast and memory efficient, utilizing sparsity and reoccurring structures in the optimization. The solver is highly flexible and supports LTI, LTV, and non-linear systems. In addition, solver settings and behaviour, such as cost function, tolerance, constraints, and weights, are highly customizable which gives a solid foundation to expand and improve upon this work.

5

Application & Analysis

This chapter presents the practical application of the developed model and controller, including the numerical parameters used in their implementation. It begins by analysing how variations in the exhaust mass flow affect the model's response time, offering insights into the implications of time-varying emf. To support the proposed EMF-based time normalization approach for model predictive control, a numerical verification is conducted through simulation. Finally, the chapter evaluates the controller's performance under various disturbances and emf conditions, highlighting its robustness and examining interesting behaviours.

5.1 Model Parameters

In Section 2.1 the thermal model was introduced. The system was divided into n-stones where each stone propagates the inlet gas further in the system until the last stone. The equation is repeated below for the reader's convenience.

$$\mathrm{C_p} m \dot{T}_i = \mathrm{C_g} \mathrm{emf}(t) \left(T_{i-1} - T_i\right) \rightarrow \dot{T}_i = \alpha \mathrm{emf}(t) \left(T_{i-1} - T_i\right)$$

Where the variable α was defined as $\alpha \coloneqq \frac{\mathrm{C_g}}{\mathrm{C_p} m}$, m is the weight of a single stone. $\mathrm{C_g}$ is the heat capacity for the exhaust gas and $\mathrm{C_p}$ is the corresponding heat capacity of the ceramic, an average of the whole muffler. In this thesis, a value of $\alpha = 0.5$ was used, which corresponds to a time constant of 20 seconds for a single stone with nominal emf ($\mathrm{emf_{nom}} = 0.1$). A time constant of 20 seconds is used as it reflects the real system dynamics accurately. Since the simulated system consists of 15 stones, the resulting model is of order 15, with a rise time of approximately 300 seconds for nominal emf.

The muffler model was then presented in state-space form as below.

$$\dot{\mathbf{x}} = \begin{bmatrix} -0.5\mathrm{emf}(t) & 0 & \dots & 0 \\ 0.5\mathrm{emf}(t) & -0.5\mathrm{emf}(t) & 0 & \dots \\ 0 & 0.5\mathrm{emf}(t) & -0.5\mathrm{emf}(t) & 0 \\ \vdots & 0 & \ddots & \ddots \end{bmatrix} \begin{bmatrix} T_1(t) \\ T_2(t) \\ \vdots \\ T_n(t) \end{bmatrix} + \begin{bmatrix} 0.5\mathrm{emf}(t) \\ 0 \\ \vdots \\ 0 \end{bmatrix} u + \begin{bmatrix} 0.5\mathrm{emf}(t) \\ 0 \\ \vdots \\ 0 \end{bmatrix} \mathrm{T_{exh}} \tag{5.1}$$

$$\mathbf{y}(t) = \begin{bmatrix} 0 & 0 & \dots & 1 \end{bmatrix} \begin{bmatrix} T_1(t) \\ T_2(t) \\ \vdots \\ T_n(t) \end{bmatrix} \tag{5.2}$$

With the corresponding matrices presented in equations 2.4- 2.7, u being the control signal and T_{exh}, the exhaust temperature which acted as a disturbance on the system by cooling or heating up the system. Furthermore, the matrices are time-dependent due to the time-varying emf(t) which affects the system by slowing down or speeding up the dynamics. Lastly, only the last stone was considered the output as the SCR is located at the end of the muffler.

5.2 EMF Impact on Model Lag

The emf is a time-varying multiplier that alters the characteristics of the plant dynamics. It speeds up or slows down the dynamics, which is illustrated in Figure 5.1. The response time of an open-loop step response has been simulated with three different emf values $(0.05, 0.15, 0.29)$, which closely reflects the range of values in the different simulation scenarios presented.

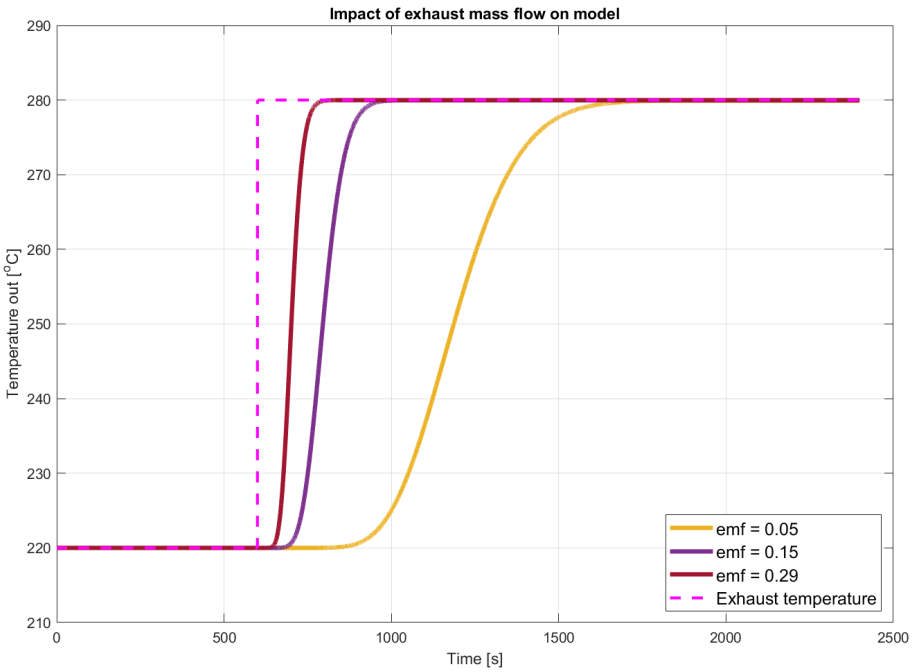

Figure 5.1: Comparison of system response with different exhaust mass flow.

In Figure 5.1 it is evident that higher emf yields a faster response time. The brown trajectory, corresponding to a constant emf $= 0.29\text{kg} \cdot s^{-1}$, reached steady state after roughly 840 seconds. On the other hand, the yellow trajectory has the slowest response time with an emf of $0.05\text{kg} \cdot s^{-1}$, and it reached steady state at 1750 seconds. As shown, the value of the emf has a great impact on the speed of the system. However, we have seen in previous sections that the usage of EMF-based time normalization can help to circumvent this issue.

5.3 Parameters of Model Predictive Controller

As established in Section 4.3.4, the MPC was constructed entirety in the τ-domain, i.e it does not operate in natural time but rather in the normalized τ-time. This means that the internal model used by the MPC was simulated in τ-domain with the unit [s'], which is normalized seconds. In τ-domain, emf does not vary with time and the dynamics of the internal model are therefore time-invariant.

5.3.1 Constraints

The MPC includes two primary constraints, as previously discussed in the problem formulation in Section 2.2.The first constraint applies to the control input, which must remain positive ($u > 0$) and within a realistic upper limit based on commercially available actuators. This upper limit was set to the nominal value of 60°C. Thus, setting the constraint on u as;

$$0°\text{C} \leq u \leq 60°\text{C}$$

The second constraint is based on the system objective: to maintain the SCR's inlet temperature (the system's output) above a critical threshold, below which the NO_x conversion drops significantly. This threshold was set to,

$$y \geq 200°\text{C}$$

5.3.2 Tunable Parameters

Tuning an MPC is essential for achieving good performance, but is rather complex since there exist several parameters and weights that significantly influence its behaviour. Key tunable parameters include the length of the prediction horizon, formulation of the cost function, cost term weights, and the number and distribution of control moves during the horizon. While an initial estimate of suitable values can be obtained from the problem formulation and system dynamics, empirical testing is ultimately necessary to achieve good performance.

In this section, selected values will be presented and motivated. The cost function which was used is presented below. Since there is only one control input and one output constraint, the notation is reduced to the scalar case. In accordance with the objective function presented in Section 4.2.3 and 4.3.2, which showed that the added thermal energy was a linear function of $u(t)$ with a weight of $W_u = 1$. The cost function at its core is then simply a linear cost on control.

$$V = \sum_{k=0}^{p-1} u_k W_u \tag{5.3}$$

To improve the numerical conditioning of the optimization problem, a small quadratic weight, ($W_u^2 = 0.01$) was introduced. It ensures that the cost function is strictly convex with respect to control input, which corresponds to a positive definite Hessian. This helps the solver converge more reliably to a unique solution. To soften the output constraint, both a quadratic slack weight W_ϵ^2 and a linear slack weight W_ϵ were introduced. Relying solely on a quadratic slack term can often result in small but persistent constraint violations. The addition of a linear slack weight helps drive these violations to zero. The combination of both quadratic and linear penalization of the slack variables ensures a strictly convex optimization with respect to slack and results in a similar formulation as the Hubert loss[21]. Including slack variables is essential to soften the constraints and ensure that the optimization problem remains feasible and minimize the constraint violation.

$$V = \frac{1}{2}\epsilon^2 W_\epsilon^2 + \epsilon W_\epsilon + \sum_{k=0}^{p-1} \frac{1}{2} u_k^2 W_u^2 + u_k W_u \tag{5.4}$$

The slack weights were set roughly two orders of magnitude higher than the weight for control effort, to prioritize constraint satisfaction over minimizing control input. Both W_u and W_u^2 were normalized with respect to the length of the prediction horizon in order to be able to compare different simulations where different prediction horizons were used.

Additionally, the prediction horizon must be set appropriately. In Section 4.3.3, Equation (4.16) provided the formulation for the minimum length of the prediction horizon the controller needs to have in normalized time. Since the controller is designed in τ-domain, the closed-loop response is constant, with a rise time of approximately 300 s' when 15 stones were used. Choosing a prediction horizon of less than 300 s' leads to unpredictable behaviours. Essentially, the controller cannot see the effects of the applied control when simulating its internal model, such lengths should therefore be avoided.

When determining the upper bound of the prediction horizon, the prediction horizon should be chosen such that it minimizes the computational cost whilst still capturing the entire necessary dynamics to properly control the system. For systems with lag, a prediction horizon 2-3 times longer than the lag time is typically adequate for predictable behaviour. This allows the controller to adequately observe and account for the delayed effects of its control actions within its internal prediction.

An interesting observation of choosing the prediction horizon in τ-domain is that a consequence of this is that the prediction horizon in natural time will vary depending on the emf(t). The final choice of prediction horizon will be further motivated in the coming sections, but as of now, it is chosen to be 600 s'.

Finally, the number of free control variables, u_j was set to 100 over the prediction horizon, structured using blocking moves such that each control input is held

constant for 5 time steps. This accelerates the optimization and decreases the computational load with minimal impact on performance.

To summarize all selected values for MPC setup and tuning variables, they are presented in the table below.

Table 5.1: Selected MPC parameters

Parameter	Value
Prediction horizon	600
Free control variables	100
Blocking moves	5
Quadratic control weight, W_u^2	0.01
Linear control weight, W_u	1
Quadratic slack weight, W_ϵ^2	1e4
Linear slack weight, W_ϵ	2e4
$u_{\min}$ (Hard constraint)	0°C
$u_{\max}$ (Hard constraint)	60°C
$y_{\min}$ (Softened constraint)	200°C

5.4 Numerical Verification of EMF Based Time Normalization

In Section 4.3.4 it was proved that an MPC could be constructed in τ-domain, similarly to what was shown in Section 3.3 for an LTI-controller. In this section, we numerically verify that time normalization yields equivalent trajectories for an MPC. The control strategy is outlined below and presented in the form of a flowchart where full-state feedback is assumed.

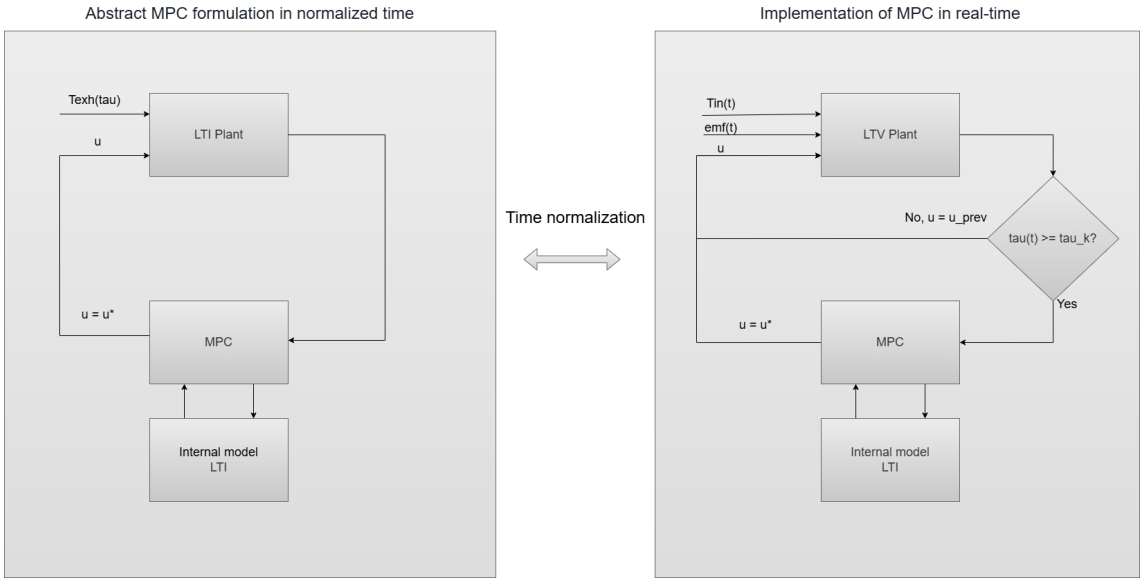

Figure 5.2: Flowchart showing implementation of LTI (left) and LTV(right) system, yielding amplitude invariant and equal trajectory.

The flowchart on the left illustrates the MPC formulation in normalized time, this system has no time-varying aspects affecting the simulation and runs in τ-domain. However, the emf is needed to extract the correct T_{exh}-values in τ-time as it is given in seconds. The system dynamics are integrated in the plant by using explicit Euler, and then a control action is computed in the MPC-block, which simulates its internal model and outputs u^*, which is fed to the plant and applied in the next integration. In this case, the internal model and plant are identical. Thus, in the time normalized domain, this is a perfectly standard MPC formulation.

The right flowchart represents the real-time implementation which is affected by emf(t). As with the normalized time simulation, we start at the plant where the dynamics are integrated. This is followed by a condition block that checks whether $\tau(t) \geq \tau_k$. As shown in Section 4.3.4, the MPC is called only when this condition is met, since the MPC operates entirely in the τ-domain. The condition block is not needed in the normalized time implementation, as the plant and internal model both run in τ-domain and therefore have the same pacing. If the condition-block is satisfied, the MPC is called, else the previously computed value is reused. Yet again, the MPC simulates its internal model which is in τ-domain and outputs u^*.

However, correct T_{exh} needs to be fed to the controller, hence emf preview is needed if the preview of T_{exh} is used.

As can be seen in Figure 5.2 and explained above, the implementations are inherently different depending on the domain. The normalized time system is not affected by the time-varying emf, whilst the real-time simulation is affected by it.

By simulating the system in both domains and post-processing the τ-domain result and displaying it on the natural time axis, it can be seen that the trajectories are equivalent, which is illustrated in Figure 5.3 below.

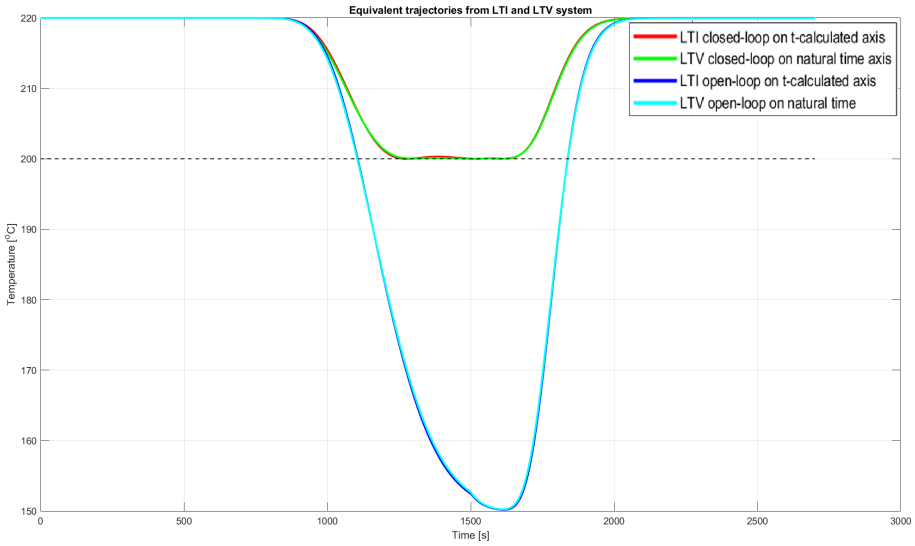

Figure 5.3: Cooldown-period simulation results from both normalized and real-time system, to show that the trajectories given from both system are equivalent.

In Figure 5.3 both the closed-loop and open-loop responses are shown for the normalized and real-time simulation. It can be seen that they have identical trajectories (disregarding some numerical error), hence confirming what has been discussed before. The same analysis can also be done for a more complicated set of inputs, a transient which is shown in the next three figures.

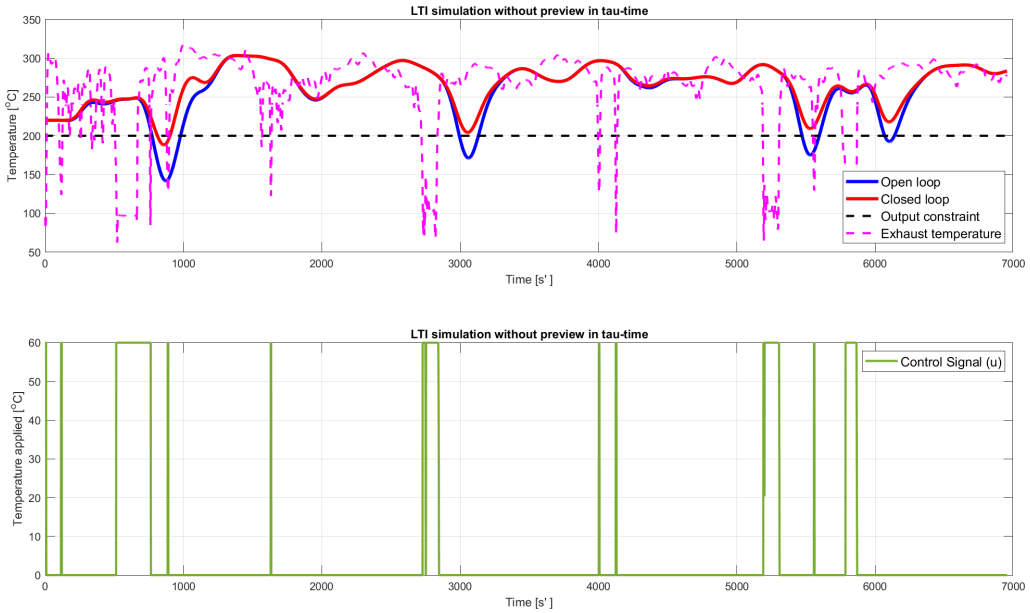

Figure 5.4: Transient Cycle normalized time results, showing no stretching or compressing of trajectories.

In Figure 5.4 system is simulated in normalized time with transient data. If the results from this simulation are put on the real time-axis we get the following figure.

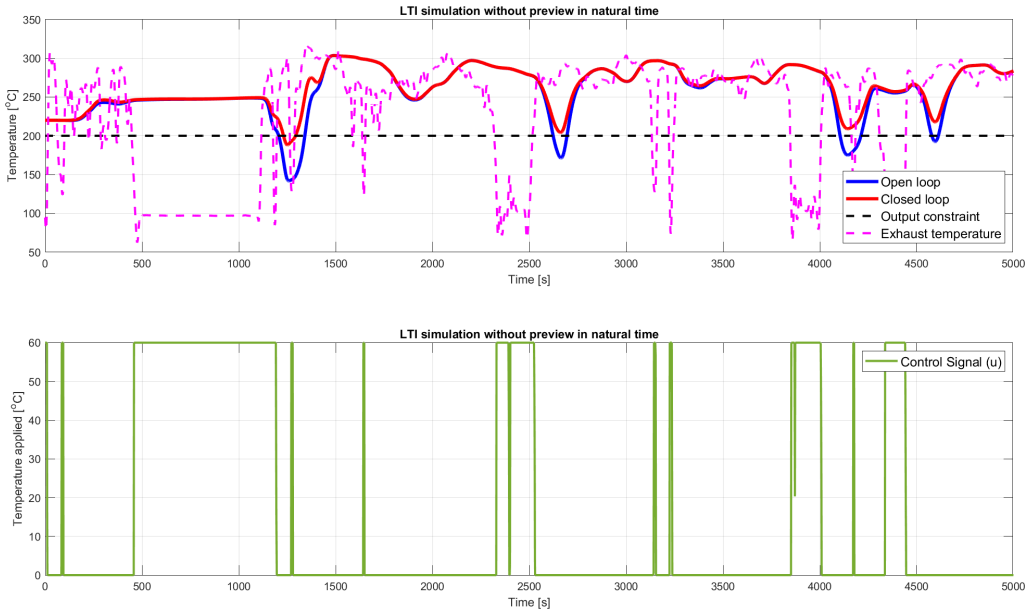

Figure 5.5: Transient Cycle LTI results on time-axis, showing stretching or compressing of trajectories.

Equivalently, the results from simulating the system in real-time can be seen in Figure 5.6.

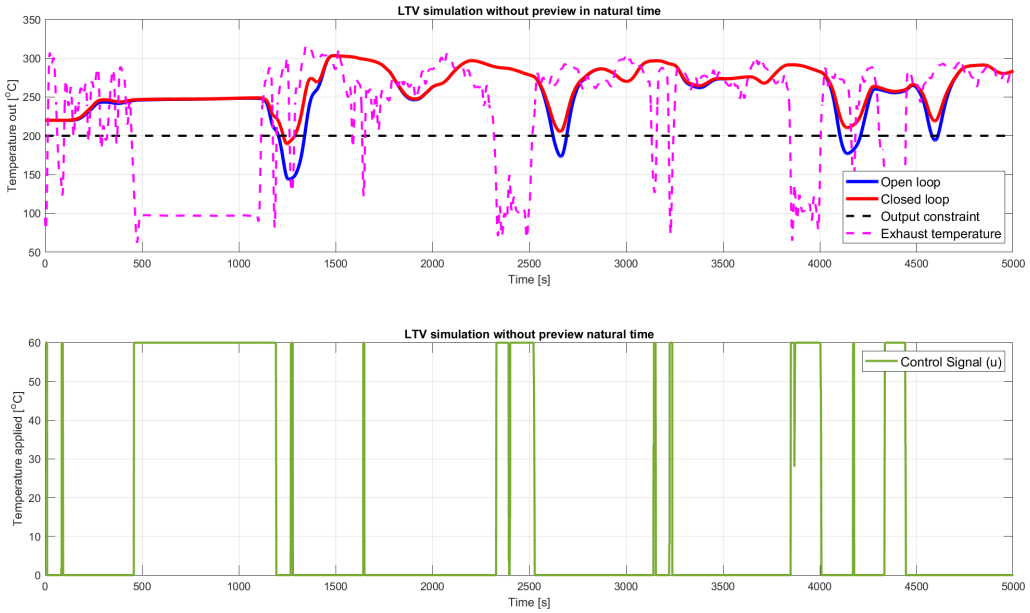

Figure 5.6: Transient Cycle results when LTV system was simulated.

Comparing figures 5.5 and 5.6 it is clear that the same control signal is applied to both plants, closed and open-loop follow the same trajectory with equivalent minimum and maximum values. As such equivalent trajectories between the two domains have now been shown, both theoretically and numerically.

Consequently, this means that it does not matter in which domain the analysis of the constructed controller is made, in the real-time or the fictitious time. The main advantage of analysing the performance in the normalized time domain is that the time-varying emf does not influence the system and a complete analysis can be done in normalized time without losing any insights. The developed framework allows us to use an MPC controller with an LTI internal model to control an LTV system, providing an efficient formulation for muffler thermoregulation. This approach eliminates the need to use more complex methods such as robust-MPC or LTV-MPC.

5.5 Controller Performance

In this section, the performance of the controller is evaluated. A range of simulation results are presented to illustrate the system's behaviour and control reaction. Including how the controller responds to challenges such as a too short prediction horizon, a medium prediction horizon compared to a long prediction horizon. Both cases with and without preview are also presented and analysed to emphasize what is gained when using preview. The implications of these findings are explored to offer insight into how model assumptions and preview length influence the closed-loop behaviour and control quality.

5.5.1 Impact of Prediction Horizon

In this section, we motivate the choice of prediction horizon more rigorously and show why a prediction horizon equal to the closed-loop rise-time (300 s') and below is inadequate for an effective controller. The simulations in this section are performed **without** utilizing preview information, and thus under the assumption that the current T_{exh} is constant during the prediction horizon.

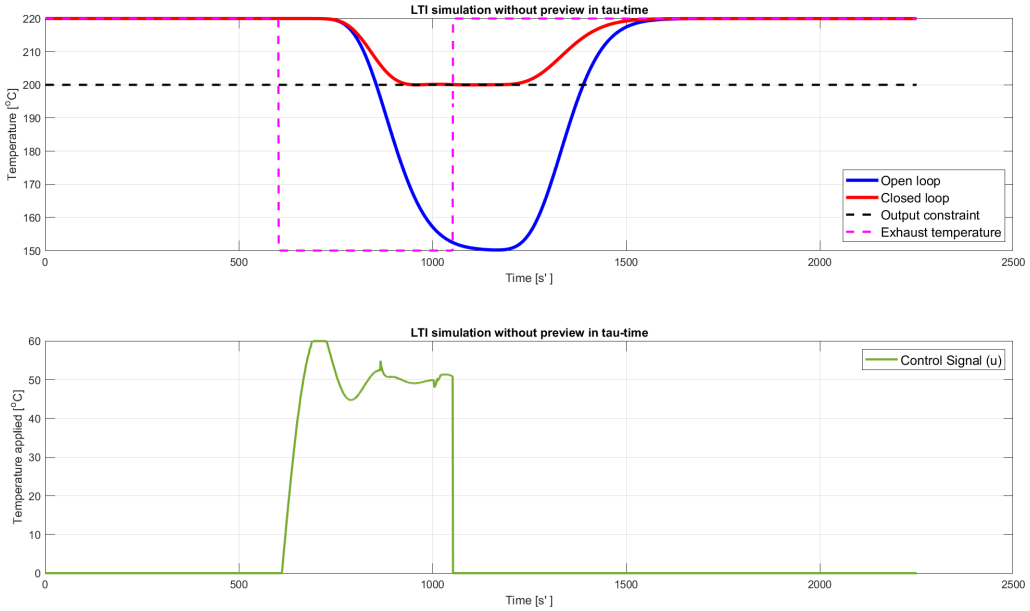

Figure 5.7: Cooldown-period simulation with a sufficiently long prediction horizon of 600 s'. Top panel shows system response (both open and closed-loop), minimum output constraint and exhaust temperature,while the bottom panel shows applied control signal.

The controller effectively controls the system, as shown in Figure 5.7. With a prediction horizon of 600 s', the length of the prediction horizon is two times that of the time constant of the normalized system. This is more than sufficiently long to capture the dynamics of the system. As such, the controller is able to anticipate the output behaviour accurately by simulating the response of its internal model.

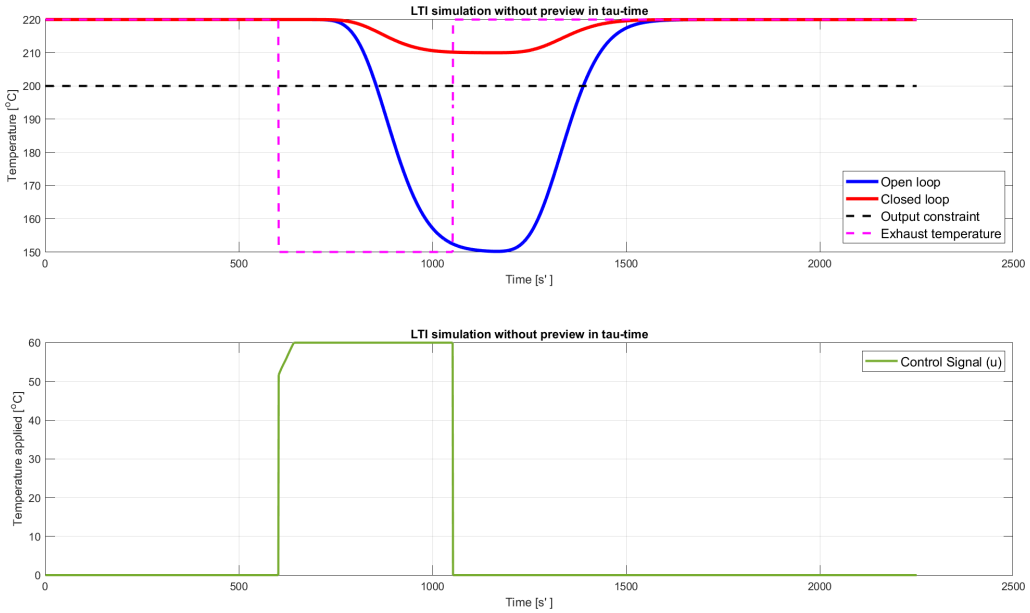

Figure 5.8: Cooldown-period simulation with a prediction horizon of 300 s'. Top panel shows system response (both open and closed-loop), minimum output constraint and exhaust temperature, while the bottom panel shows applied control signal.

Choosing a prediction horizon equal to the closed-loop response time (300 s'), the controller still satisfies the constraint, but it becomes more aggressive as illustrated in Figure 5.8 by the saturation of the control signal. We can observe that the system output lies above the constraint, the most energy-efficient is to lie on the constraint. Thus we are using excessive energy. This behaviour can be further analysed using the post-processing tool from ODYS, which allows a more in-depth analysis of what the controller does.

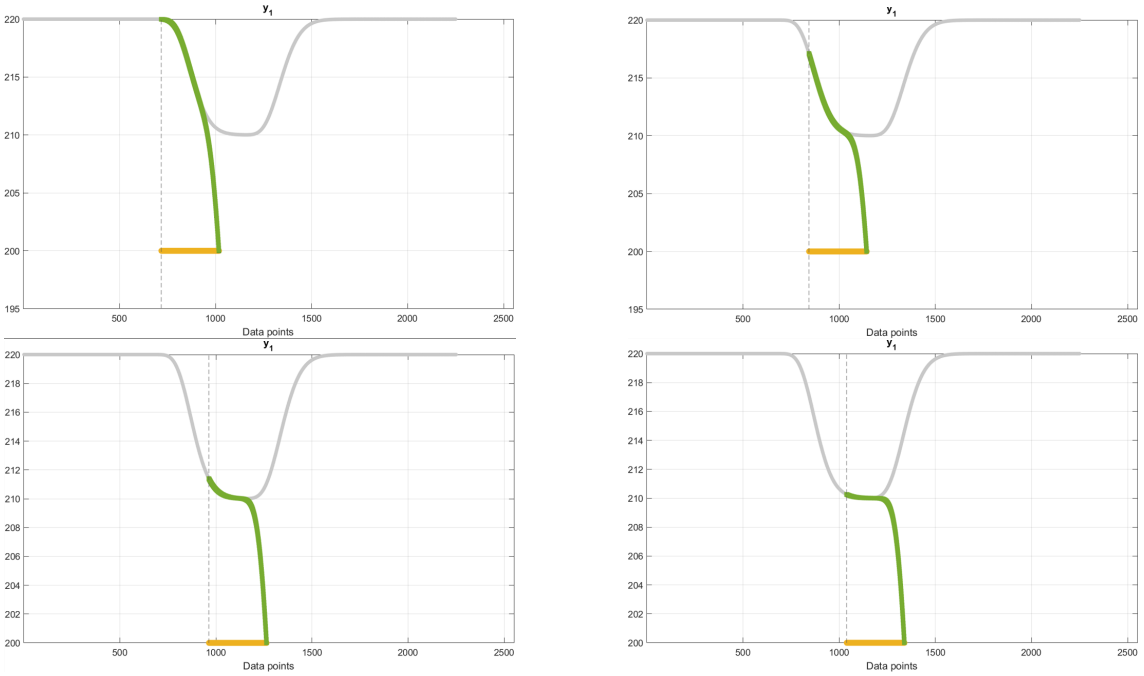

Figure 5.9: 4 plots showing different stages of the internal optimal trajectory(green line) and how last step in horizon lies on the output constraint (yellow line) despite saturated control.

In Figure 5.9 this discrepancy can be explained as the controller within its horizon observes that on the last prediction step it will lie on the constraint. Therefore, we observe this distance to the constraint which we did not observe before. When the prediction horizon is set equal to the size of the lag, the controller is unable to observe the impact of the control actions beyond the delay. This, in combination with the constraint on the control action results in unpredictable behaviour and a suboptimal solution.

Decreasing the prediction horizon even further leads to increasingly unpredictable behaviour, as the impact on the output from the control action further decreases. The next figure illustrates the results when the prediction horizon is 150 s'.

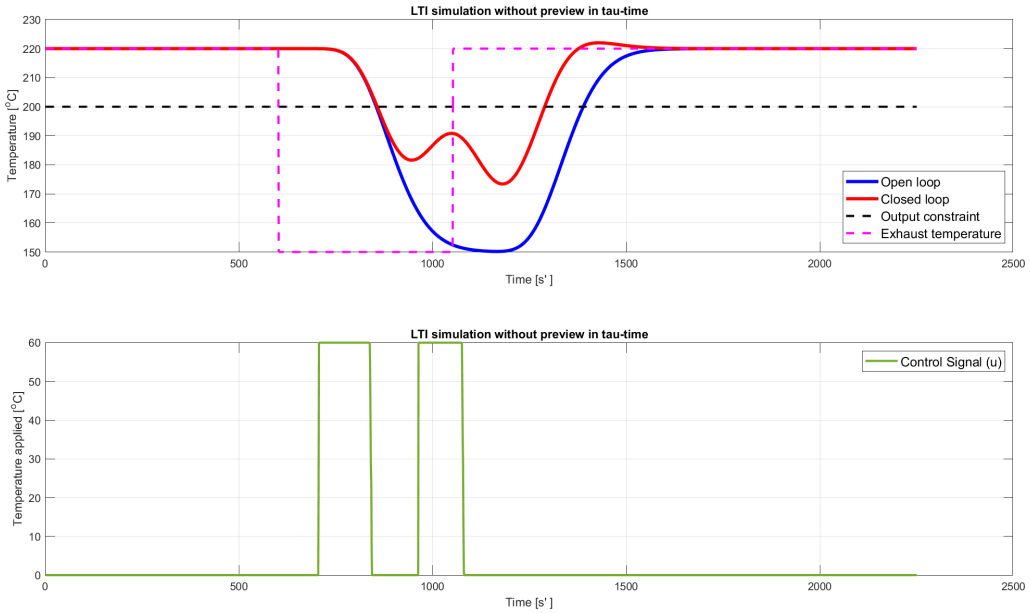

Figure 5.10: Cooldown-period simulation with a too short prediction horizon of 150 s'. Top panel shows system response (both open and closed-loop), minimum output constraint and exhaust temperature, while the bottom panel shows applied control signal.

The controller is not able to control the system adequately, as the constraints are violated. The control signal resembles a "Bang-Bang"-controller with two intervals where the control signal saturates. The observed closed-loop and control signal are unexpected and to gain insights into why this behaviour occurs, ODYS post-processing tool can once again be used. In figures 5.11-5.14 the optimal output trajectory computed by the controller within its prediction horizon at the beginning and end of each control actuation is illustrated. Each figure illustrates the start and end of actuation in Figure 5.10, this corresponds to k$\in$ [707, 838] for the first actuation and k$\in$ [964, 1076] for the second actuation. The green line represents the MPC's optimal output within the prediction horizon, the yellow line indicates the output constraint and the gray line corresponds to the entire trajectory, which aligns with the closed-loop in Figure 5.10.

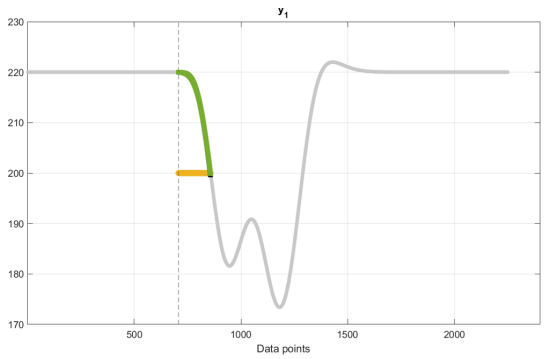

Figure 5.11: Optimal output trajectory at step k=707.

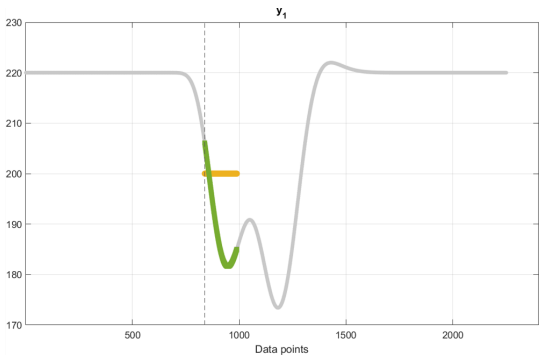

Figure 5.12: Optimal output trajectory at step k=838.

Figure 5.11 shows the step where the first control action is applied. The controller reacts as soon as it predicts that it will violate the output constraint within the prediction horizon. The controller continues to heat the muffler to minimize the cost until the temperature starts to increase again within its horizon.

In Figure 5.12 we can see the moment at which the controller stops applying a control signal. At this point, it observes that the largest slack cost is within the horizon (at the minimum of the predicted trajectory). As the cost function is driven by the worst-case deviation, the optimal strategy to minimize the cost is to set $u = 0$. As the controller cannot reduce the worst-case deviation within its horizon even if it's saturating the control signal, hence the controller stops applying any heat to the muffler.

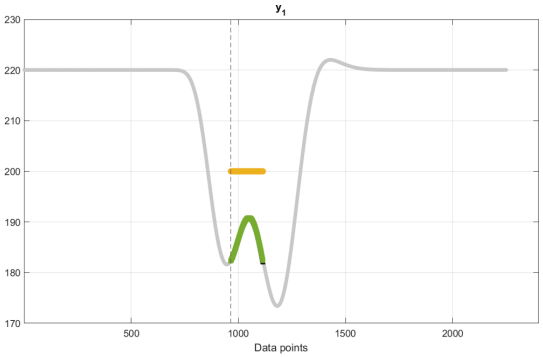

Figure 5.13: Optimal output trajectory at step k=964.

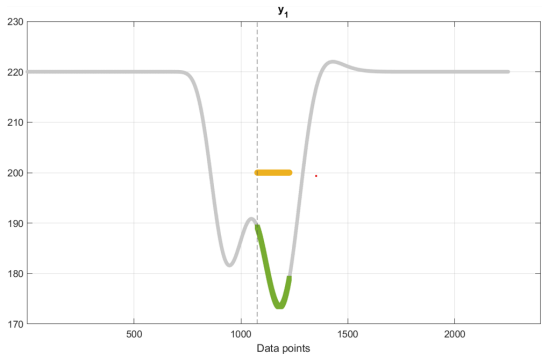

Figure 5.14: Optimal output trajectory at step k=1076.

In Figure 5.13 we can see which incentives are driving the controller to apply heat again. Within the prediction horizon, the magnitude of the constraint violation begins to increase. The previous minimum temperature is no longer the largest deviation, as the predicted trajectory now is concave and the predicted temperature continues to decrease. It now tries to minimize the cost by applying heat to the system again, this results in the second activation of the control signal.

Then, in Figure 5.14 the optimal trajectory is again convex, and the controller cannot decrease the maximum deviation. Consequently, to minimize the cost the controller should not apply any heat. However, at this stage, the disturbance has started to heat the system again and no additional heat is required from the controller.

As demonstrated by the results above, when the prediction horizon is shorter than the system's lag, the controller cannot anticipate the long-term effects and gives unpredictable behaviour. It starts to apply a control action too late but also for too short of a period to accurately see its effects, which results in poor performance. This motivates the reason to consider a prediction horizon of 600 s' for the remainder of this thesis, and why a horizon less than closed-loop response should be avoided.

5.5.2 Impact of Preview Information

One of the main objectives in this thesis is how preview information can improve control of the thermal regulation problem. As presented before, one assumption that is needed to utilize preview information of T_{exh} is that preview of emf(t) is available. When using preview, we assume that perfect preview of both emf(t) and T_{exh} is available. Simulations that do not utilize preview information rely on the assumption that the current T_{exh} is constant for the entire prediction horizon, as explained in Section 4.4.1.

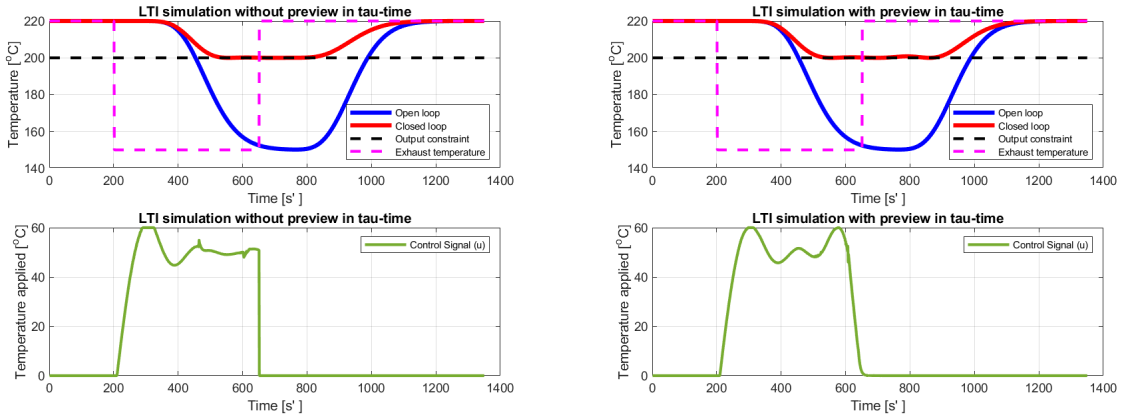

(a) Not utilizing preview. **(b)** Utilizing preview.

Figure 5.15: Comparison between utilizing preview and not for Cooldown-period simulation with a prediction horizon of 600 s' and $u_{max} = 60°C$. Top panel shows system response (both open and closed-loop), minimum output constraint and exhaust temperature, while bottom panel shows applied control signal.

In Figures 5.15a and 5.15b the results of not using preview information and using preview information are illustrated. Evidently, there is not a large difference between the two figures. Both simulations meet the constraint of the minimum temperature on the output, and both control signals appear similar.

An explanation for these results resides in the problem itself. As the T_{exh} drops to 150°C, which is within what the controller can supply as $u_{max} = 60°C$. That is, the disturbance is within the region of which the controller can compensate the difference between the output constraint and disturbance, by applying some control. Hence, no extra action is required to adhere to the constraint. This is confirmed in the figure below.

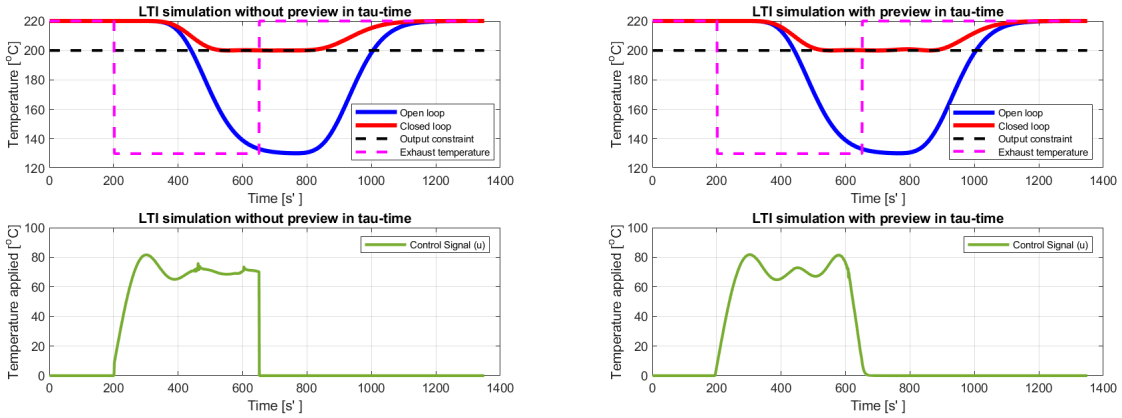

(a) Not utilizing preview. **(b)** Utilizing preview.

Figure 5.16: Comparison between utilizing preview and not for Cooldown-period simulation with a prediction horizon of 600 s', no constraint on $u_{\max}$ and T_{exh} = 130°C for cooling phase . Top panel shows system response (both open and closed-loop), minimum output constraint and exhaust temperature,while bottom panel shows applied control signal.

In figures 5.16a and 5.16b the simulation has been modified by lowering T_{exh} during the cooling phase to 130°C. In addition, the constraint on $u_{\max}$ has been removed, allowing the MPC to supply as high control signal as needed. Despite these changes, the results are similar to the previous case. The controller compensates for the disturbance to satisfy the output constraint. Since the control input is unrestricted, the controller simply delivers the required action to offset the disturbance. When the disturbance lies within the range the controller can handle, defined by the following condition.

$$T_{exh} + u_{\max} \geq y_{\min} \tag{5.5}$$

The use of preview provides limited benefit for constraint satisfaction. The next figure illustrates a case where preview information improves the performance as condition 5.5 is not satisfied.

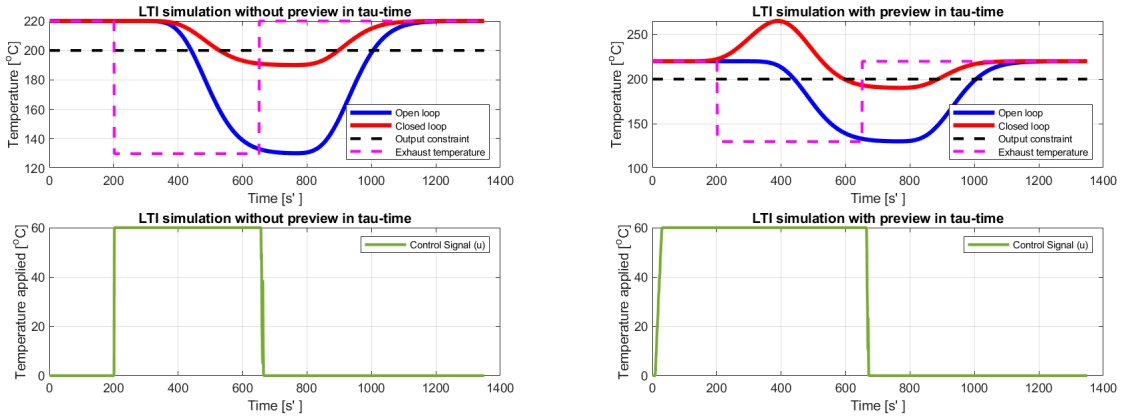

(a) Not utilizing preview. **(b)** Utilizing preview.

Figure 5.17: Comparison between utilizing preview and not for Cooldown-period simulation with a prediction horizon of 600 s' and $u_{\max} = 60°C$. Top panel shows system response (both open and closed-loop), minimum output constraint and exhaust temperature,while bottom panel shows applied control signal.

In Figure 5.17 the T_{exh} is still decreased to 130°C, however the constraint on $u_{\max}$ has been reintroduced. It can be observed in Figure 5.17a that the controller saturates the control signal for the duration of the cooldown period. The controller also violates the output constraint as it cannot apply any more heat. As preview information is not available, no proactive actions are taken, which align with what has been observed before.

However, Figure 5.17b reveals an interesting behaviour. The controller proactively begins to heat the muffler before the exhaust temperature drops, i.e it preheats. This is illustrated by the fact that the control signal saturates almost immediately, as the controller anticipates the upcoming drop in exhaust temperature. As a result, the output temperature increases to counteract the decrease of T_{exh}, yielding a longer time above the output constraint. For the two previous simulations, preheating could not be observed as condition 5.5 was satisfied. On the contrary, for preheating to occur the following condition must be satisfied:

$$T_{exh} + u_{\max} < y_{\min} \tag{5.6}$$

The drop in T_{exh} needs to be larger than what $u_{\max}$ can supply. In addition to condition (5.6), the duration of the cold period needs to be sufficiently long for the controller to deem it beneficial to preheat the muffler. In this case, condition (5.6) holds and the cold period is sufficiently long. Hence the controller will try to preheat the muffler before the cooldown period starts to minimize the cost. However, preheating is not a guarantee that the output constraint will hold. In this case, the cooling period is too long for the controller to be able to keep the constraint.

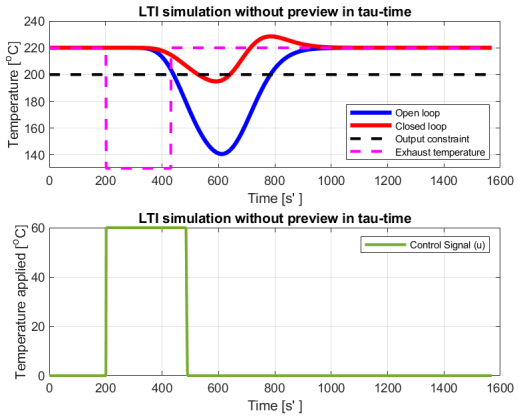

(a) Not utilizing preview.

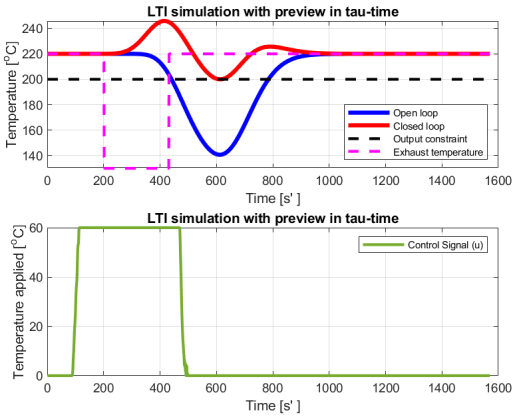

(b) Utilizing preview.

Figure 5.18: Comparison between utilizing preview and not for Cooldown-period simulation with a prediction horizon of 600 s' and $u_{\max} = 60°C$. Top panel shows system response (both open and closed-loop), minimum output constraint and exhaust temperature,while bottom panel shows applied control signal.

It might occur that the cooldown phase is short enough for the controller utilizing preview to not violate the constraint, whilst the controller which does not, violate the constraint. This is the case in figures 5.18a and 5.18b. Where the cold period of T_{exh} has been shortened to $230s'$(equivalent to 460 s) compared to $450s'$ (equivalent to 900 s) and $T_{exh} = 130°C$.

The structure of the problem can give an indication of when preview information is useful. What has been identified to give good insights regarding if the controller can meet the constraint, is the temperature of the output before T_{exh} decreases below the output constraint. The duration of the cold period and how much T_{exh} decreases below the constraint during this period, linking back to conditions 5.5 and 5.6.

A decrease of T_{exh} can be beyond what the controller can supply. However, if the duration is short, the controller might still be able to preheat the muffler and thus meet the constraint, as was illustrated in figure 5.18b. On the other hand, if the decrease is within the limit of which the controller can supply the problem becomes trivial, and utilizing preview is not beneficial in terms of improved constraint performance. The problem becomes infeasible for the controller when condition 5.6 is true and a sufficiently long cooling period occurs, then the controller is guaranteed to violate the output constraint.

Realistic scenario: Driving cycle data

Up to this point, the results presented have been based on a simple example, where the exhaust temperature (T_{exh}) only has taken two distinct values. A more realistic scenario is the driving cycle data, which is normalized data coming from test cells. As we now will present performance metrics the following figures will illustrate the

real-time simulation instead of the normalized time. Below is the driving cycle 1 data presented without and with preview, followed by a table summarizing key metrics.

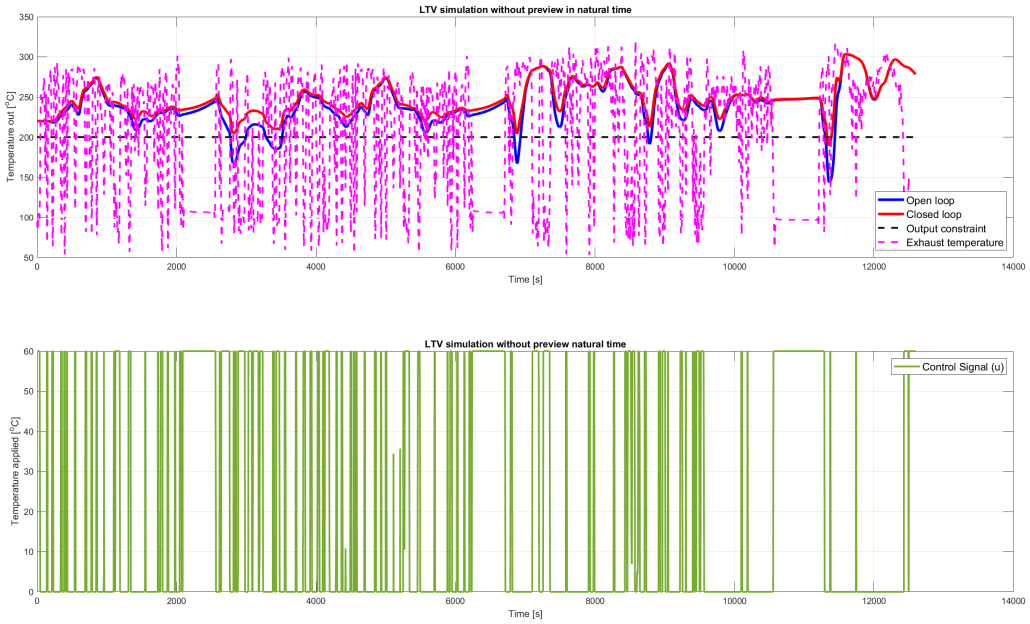

Figure 5.19: Simulation of driving cycle 1 data without preview. Top panel shows system response (both open and closed-loop), minimum output constraint and exhaust temperature,while the bottom panel shows applied control signal.

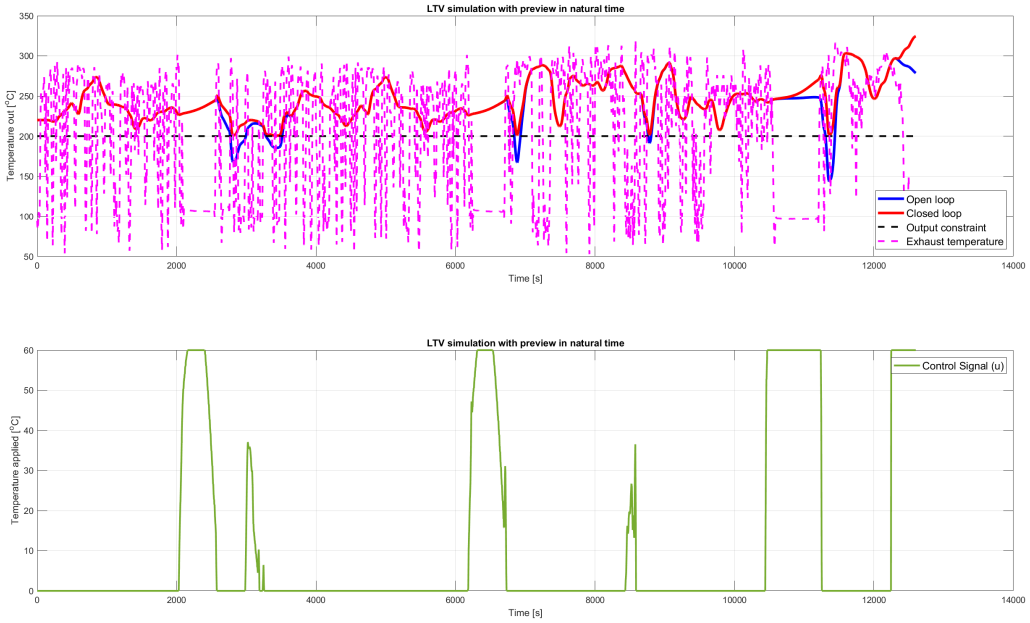

Figure 5.20: Simulation of driving cycle 1 data with preview. Top panel shows system response (both open and closed-loop), minimum output constraint and exhaust temperature,while the bottom panel shows applied control signal.

Table 5.2: Performance comparison of simulations with and without preview

Scenario: Transient 1	**Percentage Below Constraint (%)**
Without preview	0.47
With preview	0
Energy Usage Reduction with Preview: 37.04%	

Due to the assumption that is used when not utilizing preview, fast actuation can be observed in Figure 5.19 compared to Figure 5.20. Table 5.2 illustrates the key differences between the two simulations. From the table it can be seen that the controller without preview violates the constraint for a small duration. Furthermore, it appears that the controller with preview uses less energy.

To further validate what was shown for driving cycle 1 simulation the next figures illustrate the driving cycle 2 data.

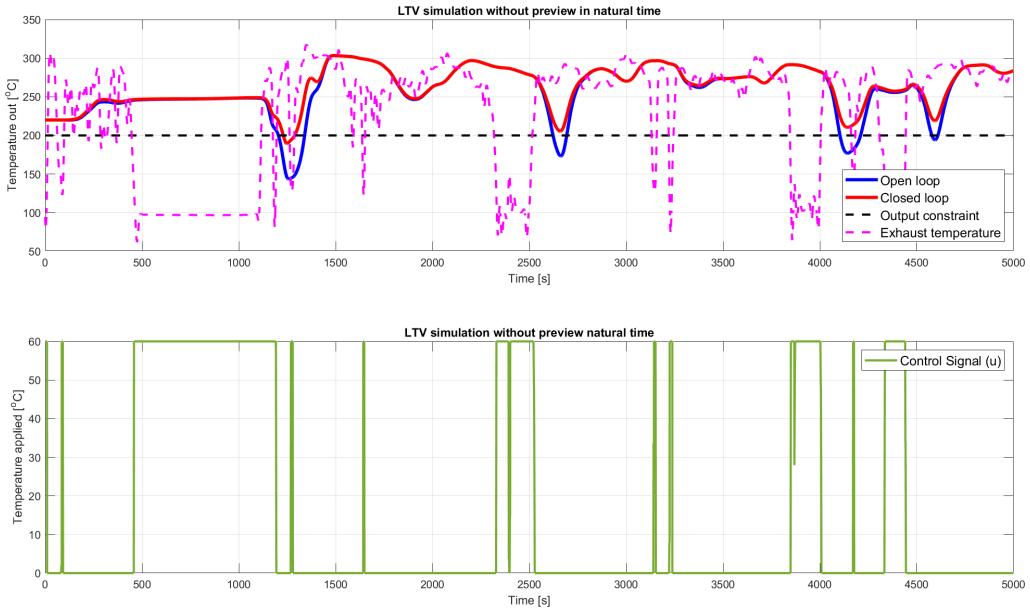

Figure 5.21: Simulation of driving cycle 2 data without preview. Top panel shows system response (both open and closed-loop), minimum output constraint and exhaust temperature,while the bottom panel shows applied control signal.

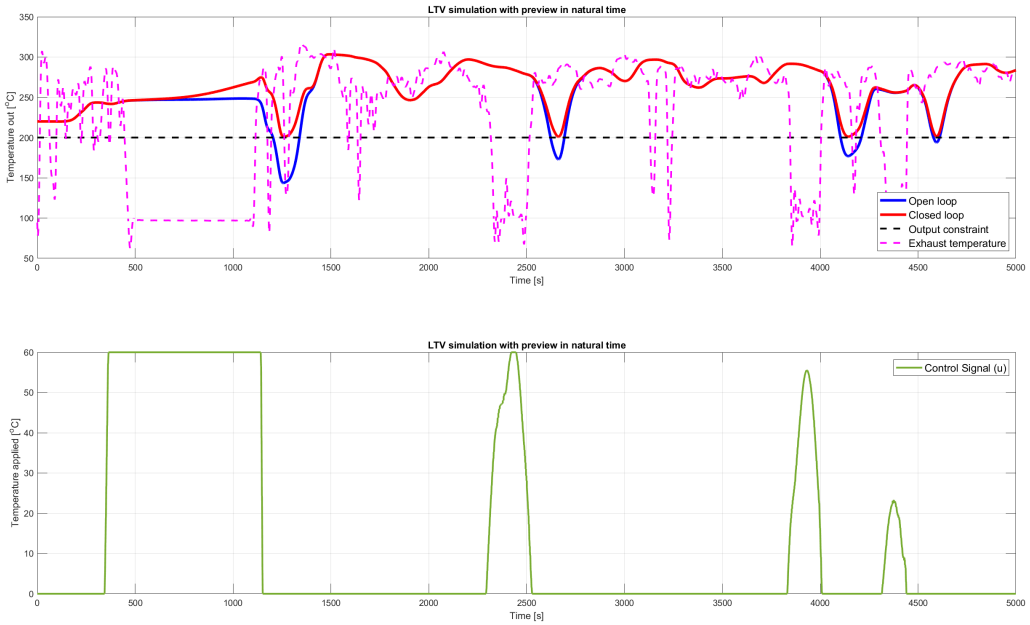

Figure 5.22: Simulation of driving cycle 2 data. Top panel shows system response (both open and closed-loop), minimum output constraint and exhaust temperature,while the bottom panel shows applied control signal.

Table 5.3: Performance comparison of simulations with and without preview

Scenario: Transient 2	**Percentage Below Constraint (%)**
Without preview	1.2
With preview	0.0
Energy Usage Reduction with Preview: 14.6%	

In Figure 5.21 the transient is illustrated again with a controller that does not utilize preview information. For both transient simulations without preview, it can be observed that the controller does not try to preheat and thus violates the constraint. Also, the controller reacts to quick changes in the exhaust temperature which the controller using preview information did not react to. This is due to the assumption of the disturbance, i.e. it is constant for the entire prediction horizon based on the current disturbance value. By using a different assumption on the disturbance, the spikes in the control signal could disappear. This is also the reason why it appears that the controller without preview uses more energy, all the small actuations add up.

In Figure 5.22 the results with preview information are illustrated, we can readily identify where the controller preheats the muffler to avoid constraint violation as the open-loop violates constraint on four distinct time instances while the closed-loop does not. Again it appears in table 5.3 that the controller with preview uses less energy and that the controller without preview violates the constraint.

From this discussion, it is important to note that, the controller without preview is using the previewing controller's weights, shown in section 5.3.2. Thus the compari-

son is not entirely fair, as it should be tuned for the case without preview as well. By doing so the comparison becomes a bit more difficult. As it depends on the tuning, by using the tuned controller as a baseline. It could be observed that by increasing the weight on control, the controller would use less energy but this increased the duration of constraint violation. If the weights were lowered for control, thus creating an incentive to use more control, the energy consumption would increase but still violate the output constraint to some extent. The controller without preview could never reach 0% constraint violation by tuning, but it could be more energy efficient.

Even though the controller without preview information violates the constraint, it still shows adequate performance. If the drop in the exhaust temperature is large, it saturates the control to compensate. The performance would have been worse if any of the controllers under-heated, which in fact none of the controllers do. Both controllers with and without preview are functioning as intended, though the controller with preview is better. Due to its ability to anticipate future events and take them into consideration.

5.6 Persistent feasibility

From this discussion comes an important consideration, *persistent feasibility*. This means that the controller should be able to find a feasible solution (satisfying constraints and limitations) at each time step, by avoiding control actions that could lead the system out of the feasible region beyond the prediction horizon. For the thermal regulation problem studied in this thesis, maintaining persistent feasibility is not trivial due to the time-varying dynamics, actuator limitations, and disturbance acting on the system. Thus it is important to consider what happens with the exhaust temperature beyond the horizon with and without preview.

If the exhaust temperature could be reasonably bounded, one could design a conservative controller by assuming a worst-case temperature drop beyond the horizon. This would result in always ending the horizon by preheating the muffler. However, in practice this approach has limitations, as the exhaust temperature cannot be reliably bounded as the drop could be too large and as such satisfy condition 5.6 and having a large cooling period beyond the horizon. If that is the case, it has been shown that the controller will be unable to maintain the minimum required temperature, regardless of how much it preheats during the horizon.

This highlights a critical insight: for systems with hard constraints and bounded actuation, persistent feasibility cannot always be guaranteed. It depends not only on the length of the prediction horizon, but also on what happens beyond the prediction horizon. However, since our prediction horizon is quite long, events beyond this horizon are unlikely to become a problem for the controller.

5.7 Main Findings

In the previous sections, important behaviours that the controller exhibits have been investigated and are summarized below for convenience.

- Prediction length needs to be **at least** longer than closed-loop response to ensure an adequate and predictable controller.
- Preview **improves** constraint performance in specific cases by **preheating**. Mainly when the exhaust temperature decreases more than the controller is able to supply. This comes at a cost of increased energy consumption.
- Quick drops in disturbance are ignored when using preview, which **decreases** energy usage.

6
Conclusion

This thesis has explored the challenge of maintaining sufficient muffler temperatures in heavy-duty diesel vehicles to ensure effective NOx conversion in the SCR catalyst. A novel approach was presented by leveraging EMF-based time normalization to reformulate the thermal regulation problem into a time-invariant framework, significantly simplifying controller design and improving computational efficiency.

Through modelling and analysis, we demonstrated that time normalization allows an otherwise time-varying system to be expressed in a time-invariant form, making it more tractable for model predictive control. Simulations confirmed that utilizing preview information allows the controller to take proactive measures and prepare the system for upcoming thermal disturbances, extending constraint satisfaction even in challenging scenarios. This came at the cost of increased energy input in general, though in certain scenarios it is actually able to save energy by neglecting very short periods of cold exhaust, whereas the controller without preview would instantly apply additional heat when the exhaust temperature would decrease.

In contrast, when preview information was absent, the controller performed comparably in cases where disturbances remained within its ability to compensate, confirming the robustness of the baseline MPC design. However, the advantages of preview were most evident in edge cases where proactive heating was necessary to avoid constraint violation. Even though these situations might be rare, they have to be taken into consideration when designing the EATS, which typically is overdimensioned. Therefore, implementing preview information can not only reduce harmful emissions, but also reduce the cost of manufacturing the EATS.

7
Future work

While this study provides significant insights into the benefits of time-scaling and utilizing preview information in context of thermal regulation in EATS, several questions remain unanswered.

To begin with, one assumption in this work is the availability of complete preview of both exhaust temperature and emf. These are not trivial to retrieve, but should also be possible with the amount of data available today. With data such as road inclination, traffic and continuous engine data, we believe one could formulate an accurate model of at least future exhaust temperature and possibly exhaust mass flow as well. The quality of preview information is vital as poor quality can lead to a controller performing worse than when not utilizing preview. One possible approach here is to add a uncertainty score to the preview information and take it into consideration. Alternatively, more heuristic approaches such as improving the assumption of constant exhaust temperature during the entire horizon could be tested. Such assumptions could be weighted mean of previous values, gradient of previous values, conservative/optimistic assumption, or a combination of several.

Lastly, Whilst we showed improved constraint satisfaction it is still debatable if the overall impact on the environment is positive as it comes at a cost of increased fuel consumption leading to more exhaust and emissions. This needs to be studied in depth to fully understand the implications.

Bibliography

[1] (2020) Regulation (ec) no 595/2009 of the european parliament. Accesed: Nov 2024. [Online]. Available: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32009R0595

[2] K. Sun, G. Zhang, K. Zhao, W. Sun, G. Li, S. Bai, C. Lin, and H. Cheng, "Study on urea crystallization risk assessment and influencing factors in after-treatment system of diesel engines," *Applied Sciences*, vol. 14, no. 2, 2024. [Online]. Available: https://www.mdpi.com/2076-3417/14/2/684

[3] O. Holmer and L. Eriksson, "Modeling and analytical solutions for optimal heating of aftertreatment systems," *IFAC-PapersOnLine*, vol. 52, no. 5, pp. 523–530, 2019, 9th IFAC Symposium on Advances in Automotive Control AAC 2019. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2405896319307062

[4] J. Dahl, H. Wassén, O. Santin, M. Herceg, L. Lansky, J. Pekar, and D. Pachner, "Model predictive control of a diesel engine with turbo compound and exhaust after-treatment constraints," *IFAC-PapersOnLine*, vol. 51, no. 31, pp. 349–354, 2018, 5th IFAC Conference on Engine and Powertrain Control, Simulation and Modeling E-COSM 2018. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2405896318325357

[5] M. R. Karim, B. Egardt, E. R. Gelso, and N. Murgovski, "Supervisory control for no_x emission compliance of heavy-duty vehicles," *IEEE Transactions on Vehicular Technology*, vol. 70, no. 12, pp. 12 402–12 414, 2021.

[6] M. R. Karim, N. Murgovski, E. R. Gelso, and B. Egardt, "Supervisory framework and model-based control of engine and exhaust aftertreatment system," in *2018 European Control Conference (ECC)*, 2018, pp. 959–964.

[7] J. Kim, T. Lee, C.-J. Kim, and K. Yi, "Model predictive control of a semi-active suspension with a shift delay compensation using preview road information," *Control Engineering Practice*, vol. 137, p. 105584, 2023. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0967066123001533

[8] Thamarai Kannan, Harish Kumar and Ferris, John B., "Discretization-based semi-active suspension control using road preview data," in *Automotive Technical Papers*. SAE International, aug 2024. [Online]. Available: https://doi.org/10.4271/2024-01-5087

[9] Z. Zhang, B. Knauder, M. Ackerl, and J. Pell, "Validation of road-preview-based predictive gear selection on heavy-duty vehicle transmission control unit," 04 2020.

[10] M. Sampei and K. Furuta, "On time scaling for nonlinear systems: Application to linearization," *IEEE Transactions on Automatic Control*, vol. 31, no. 5, pp. 459–462, 1986.

[11] U. Vollmer and J. Raisch, "Control of batch cooling crystallization processes based on orbital flatness," *International Journal of Control*, vol. 76, no. 16, pp. 1635–1643, 2003. [Online]. Available: https://doi.org/10.1080/00207170310001626419

[12] M. Sampei and K. Furuta, "On linearization and control with transformation of the time scale," *IFAC Proceedings Volumes*, vol. 20, no. 5, Part 8, pp. 97–102, 1987, 10th Triennial IFAC Congress on Automatic Control - 1987 Volume VIII, Munich, Germany, 27-31 July. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1474667017550715

[13] B. Kiss and E. Szadeczky-Kardoss, "Tracking control of the orbitally flat kinematic car with a new time-scaling input," in *2007 46th IEEE Conference on Decision and Control*, 2007, pp. 1969–1974.

[14] (2016) Air-quality-in-europe-2016. Accesed: Nov 2024. [Online]. Available: https://www.eea.europa.eu/publications/air-quality-in-europe-2016

[15] (2024, Oct) Soil acidity. Accesed: Nov 2024. [Online]. Available: https://agriculture.vic.gov.au/farm-management/soil/soil-acidity

[16] (2023) Kväveoxider, utsläpp till luft. Accesed: Nov 2024. [Online]. Available: https://www.naturvardsverket.se/data-och-statistik/luft/utslapp/utslapp-av-kvaveoxider-till-luft/

[17] (2024) Ambient (outdoor) air pollution. Accesed: Nov 2024. [Online]. Available: https://www.who.int/en/news-room/fact-sheets/detail/ambient-(outdoor)-air-quality-and-health

[18] P. M. MICHEL FLIESS, JEAN LÉVINE and P. ROUCHON, "Flatness and defect of non-linear systems: introductory theory and examples," *International Journal of Control*, vol. 61, no. 6, pp. 1327–1361, 1995. [Online]. Available: https://doi.org/10.1080/00207179508921959

[19] M. D. Q. Rawlings James B and M. M. Diehl, *Model Predictive Control: Theory, Computation, and Design*. Nob Hill Publishing, 2009.

[20] ODYS S.r.l., "ODYS Embedded MPC," (https://odys.it/embedded-mpc), 2019.

[21] P. J. Huber, "Robust estimation of a location parameter," *Annals of Mathematical Statistics*, vol. 35, no. 1, pp. 73–101, 1964.

DEPARTMENT OF SOME SUBJECT OR TECHNOLOGY
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden
www.chalmers.se

CHALMERS
UNIVERSITY OF TECHNOLOGY