

## Article

# Model Predictive Supervisory Control for Integrated Emission Management of Diesel Engines

Johannes Ritzmann <sup>1,\*</sup>, Christian Peterhans <sup>2</sup>, Oscar Chinellato <sup>2</sup>, Manuel Gehlen <sup>2</sup> and Christopher Onder <sup>1</sup>

<sup>1</sup> Department of Mechanical Engineering and Process Control, ETH Zürich, 8092 Zürich, Switzerland; onder@idsc.mavt.ethz.ch

<sup>2</sup> FPT Motorenforschung AG, 9320 Arbon, Switzerland; christian.peterhans@ivecogroup.com (C.P.); oscar.chinellato@gmail.com (O.C.); manuel.gehlen@ivecogroup.com (M.G.)

\* Correspondence: jritzman@idsc.mavt.ethz.ch

**Abstract:** In this work, a predictive supervisory controller is presented that optimizes the interaction between a diesel engine and its aftertreatment system (ATS). The fuel consumption is minimized while respecting an upper bound on the emitted tailpipe NO<sub>x</sub> mass. This is achieved by optimally balancing the fuel consumption, the engine-out NO<sub>x</sub> emissions, and the ATS heating. The proposed predictive supervisory controller employs a two-layer model predictive control structure and solves the optimal control problem using a direct method. Through experimental validation, the resulting controller was shown to reduce the fuel consumption by 1.1% at equivalent tailpipe NO<sub>x</sub> emissions for the nonroad transient cycle when compared to the operation with a fixed engine calibration. Further, the controller's robustness to different missions, initial ATS temperatures, NO<sub>x</sub> limits, and mispredictions was demonstrated.

**Keywords:** integrated emission management; variable engine calibration; pollutant emissions; aftertreatment system; supervisory control; model predictive control



**Citation:** Ritzmann, J.; Peterhans, C.; Chinellato, O.; Gehlen, M.; Onder, C. Model Predictive Supervisory Control for the Integrated Emission Management of Diesel Engines. *Energies* **2022**, *15*, 2755. <https://doi.org/10.3390/en15082755>

Academic Editor: Dimitrios C. Rakopoulos

Received: 18 March 2022

Accepted: 7 April 2022

Published: 8 April 2022

**Publisher's Note:** MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Copyright:** © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

## 1. Introduction

Today, diesel combustion engines are used in a variety of applications. As a byproduct of combustion, they emit NO<sub>x</sub>, which has adverse effects on human health and on the environment. Therefore, stringent pollutant legislation has been introduced in recent decades to limit the emission of NO<sub>x</sub> [1]. At the same time, the fuel consumption must be minimized, in order to limit the operational cost as well as the CO<sub>2</sub> emissions. Unfortunately, these two goals are conflicting [2] and necessitate the careful control of the engine operation.

To extend this inherent trade-off, modern engine systems are equipped with an aftertreatment system (ATS) to combat NO<sub>x</sub> and other pollutant emissions. For diesel engines, selective catalytic reduction (SCR) systems are typically used, whereby NO<sub>x</sub> is reduced using ammonia [3]. For these systems to operate effectively, their temperature should surpass a lower limit of approximately 200 °C.

As the engine operation influences the operation of the ATS, a joint optimization of the engine and the ATS is required. Such an approach is known in the literature as integrated emission management (IEM) [4], whereby a supervisory controller is used to coordinate the operation of the engine and the ATS during a mission.

In a previous publication by the authors, offline optimization was used to investigate this problem [5]. It was shown that, in order to achieve optimal performance, the fuel consumption, the engine-out NO<sub>x</sub> emissions, and the enthalpy provided to the ATS must be carefully managed during a mission.

In this paper, the previously obtained optimal input trajectories are applied at the testbench in a feedforward fashion. An error in the emitted tailpipe NO<sub>x</sub> mass is observed and attributed to model mismatch. A predictive supervisory controller is therefore required

that optimizes the engine operation online and compensates for model mismatch and mispredictions using feedback control.

### 1.1. Literature Review

Early works neglected the ATS entirely [6,7]. As a result, good performance was achieved using the equivalent emission minimization strategy, whereby a feedback controller sets the equivalence factor used to determine the optimal trade-off between fuel consumption and NO<sub>x</sub> emissions. The method is closely related to the equivalent consumption minimization strategy (ECMS) applied to the energy management problem of hybrid electric vehicles [8].

In later works [9,10], the NO<sub>x</sub> reduction by the ATS was considered, but the effect of the engine operation on the ATS temperature was neglected. In [9], an ECMS-like causal controller is presented that adapts an equivalence factor for the tailpipe NO<sub>x</sub> emissions. Subsequently, the equivalence factor for the engine-out NO<sub>x</sub> emissions was determined, based on the measured NO<sub>x</sub>-reduction efficiency of the ATS. In [10], a model predictive controller is used to set the operation of the engine and the ATS, whereby the effect of the control action on the ATS temperature is neglected. While these methods are straightforward to implement, the fact that they do not optimize the ATS temperature dynamics could result in a suboptimal operation, especially for missions with a low ATS temperature.

In other works [11–14], the benefit of heating the ATS is approximated offline. This approximation is then used by the supervisory controller when determining the optimal engine operation online. In [11], a feedforward controller based on Pontryagin's minimum principle was implemented, whereby a heuristic method was used to determine the equivalence factor for the ATS heating. In [12], this controller was extended to include feedback control based on the measured tailpipe NO<sub>x</sub> emissions. In [13], dynamic programming was used to determine the optimal operation for several missions offline. Based on the results, a mapping was fitted that relates the equivalence factor for the ATS heating to the current value of the ATS temperature. This mapping was then used in an ECMS-like online controller in [14].

To the best of the authors' knowledge, no publication exists in which an online supervisory controller is developed that uses predictive information about the upcoming mission to minimize the fuel consumed during the mission, while respecting legislative limits on the tailpipe NO<sub>x</sub> emissions and explicitly considering the effect of the engine operation on the ATS temperature.

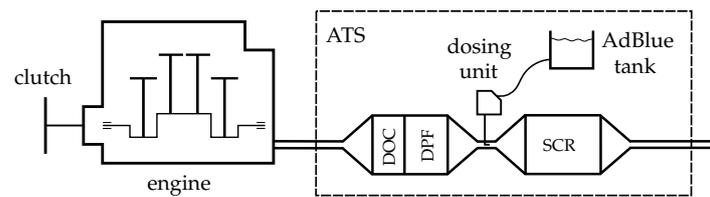
### 1.2. Contribution

The contribution of this work is twofold. First, a real-time-capable optimization method to solve the IEM problem by solving a nonlinear program (NLP) is presented. Thereby, continuously differentiable look-up elements are used to describe a preoptimized engine operation, i.e., a Pareto front with respect to fuel consumption, NO<sub>x</sub> emissions, and enthalpy provided to the ATS. Due to its computational efficiency, the resulting optimization can be performed in real time.

Second, a predictive supervisory controller is designed that employs this optimization method in combination with predictive information about the upcoming mission in a model predictive control framework. Thereby, the vehicle's fuel consumption is minimized, while a limit on the emitted tailpipe NO<sub>x</sub> mass is respected. Based on the predictive information, the controller adapts the engine operation to the mission and optimally sets the trade-off between the fuel consumption, the NO<sub>x</sub> emissions, and the ATS heating. The proposed controller was experimentally validated at an engine testbench.

## 2. System Description

A schematic of the engine system considered in this work is shown in Figure 1. It consists of a diesel engine equipped with an ATS. The main characteristics of the engine system are listed in Table 1.



**Figure 1.** Schematic of engine system layout. The aftertreatment system (ATS) features a diesel oxidation catalyst (DOC), a diesel particulate filter (DPF), and a selective catalytic reduction (SCR) system.

**Table 1.** Main characteristics of the engine system.

Engine type	inline 4-cylinder diesel
Displacement	3 L
Max. power	130 kW
Homologation	EURO 6 step D

For the experiments, the engine was fitted to a dynamic testbench. An AVL KMA 4000 fuel flow sensor was used to measure the fuel consumption. The stock  $\text{NO}_x$  sensors were used to measure the engine-out  $\text{NO}_x$  concentration and the tailpipe  $\text{NO}_x$  concentration. The stock air mass flow sensor was used to measure the air mass flow. The temperature measured downstream of the ATS using the stock sensor was used to characterize the ATS temperature. The  $\text{NO}_x$  mass flow was calculated based the  $\text{NO}_x$  concentration and the exhaust mass flow, which was obtained as the sum of the air mass flow and the fuel mass flow. The  $\text{NO}_x$ -reduction efficiency of the ATS was calculated based the measured  $\text{NO}_x$  concentrations.

The engine considered in this work is overactuated. Consequently, some engine control inputs can be adapted online in order to achieve a certain objective. In this work, the five engine control inputs start of injection  $\varphi_{\text{soi}}$ , fuel rail pressure  $p_{\text{rail}}$ , variable geometry turbine actuator position  $u_{\text{vgt}}$ , exhaust gas recirculation valve position  $u_{\text{egr}}$ , and exhaust flap position  $u_{\text{flap}}$ , are considered as degrees of freedom that can be optimized while providing the requested torque  $T_e$  at the current speed  $\omega_e$ . The resulting vector of engine control inputs is

$$\mathbf{u} = \begin{bmatrix} \varphi_{\text{soi}} \\ p_{\text{rail}} \\ u_{\text{vgt}} \\ u_{\text{egr}} \\ u_{\text{flap}} \end{bmatrix}. \quad (1)$$

As a simplification, only  $\text{NO}_x$  emissions were considered in this work. The engine operation was chosen such that the emitted carbon monoxide, unburned hydrocarbons, and particulate matter can be oxidized in the diesel oxidation catalyst (DOC) or trapped in the diesel particulate filter (DPF).

### 2.1. Engine Model

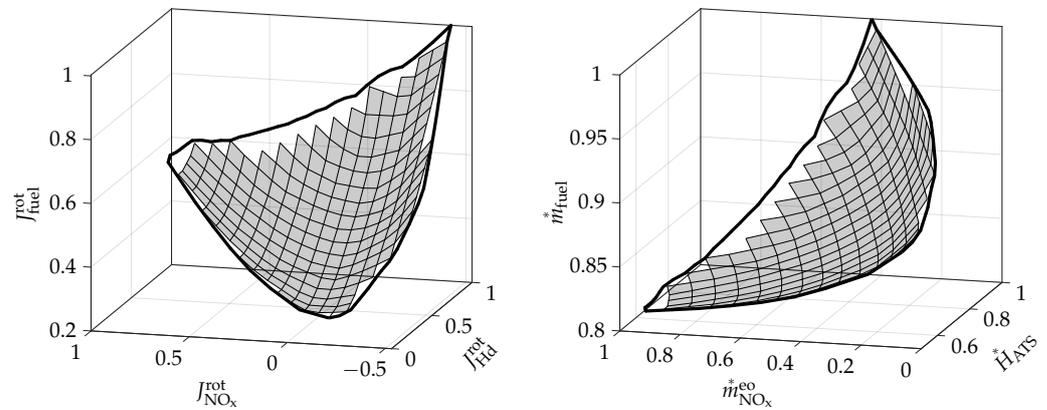
In this work, a reduced-order engine model resulting from a preoptimization of the engine operation was adopted from [5]. The benefit of the preoptimization is that the computational complexity of the subsequent online optimization is reduced. This is achieved by reducing the number of inputs from the five engine control inputs of (1) to two engine strategy inputs  $J_{\text{Hd}}^{\text{rot}}$  and  $J_{\text{NO}_x}^{\text{rot}}$ , by limiting the engine operation to the Pareto front defined by a multi-objective optimization. Furthermore, the preoptimization guarantees the feasibility of the engine operation according to a set of limits, which no longer need to be checked during the subsequent online optimization.

The objectives of the preoptimization are minimizing the fuel consumption, minimizing the engine-out NO<sub>x</sub> emissions, and maximizing the enthalpy provided to the ATS. The latter is defined as

$$\dot{H}_{ATS}^* = \dot{m}_{exh}^* c_p (\vartheta_{exh} - \vartheta), \quad (2)$$

where  $\dot{m}_{exh}^*$  is the exhaust mass flow,  $c_p$  is the specific heat capacity of the exhaust gas,  $\vartheta_{exh}$  is the exhaust gas temperature, and  $\vartheta$  is the ATS temperature.

The preoptimized engine model uses look-up elements to characterize the Pareto front. To achieve a numerically robust interpolation, a scaled and rotated version of the Pareto front is stored using the variables  $J_X^{rot}$ , where  $X \in \{Hd, NO_x, fuel\}$ . An example of such a Pareto front for a given speed, torque, and ATS temperature is shown on the left of Figure 2. The Pareto front is described by a five-dimensional look-up element for  $J_{fuel}^{rot}$  with inputs  $(\omega_e, T_e, \vartheta, J_{Hd}^{rot}, J_{NO_x}^{rot})$ , while the bounds on  $J_{NO_x}^{rot}$  are described by two four-dimensional look-up elements for  $J_{NO_x, min}^{rot}$  and  $J_{NO_x, max}^{rot}$  with inputs  $(\omega_e, T_e, \vartheta, J_{Hd}^{rot})$ .



**Figure 2.** Exemplary Pareto front at an engine speed of 2000 rpm, an engine torque of 200 Nm, and an ATS temperature of 67 °C. The gray surface shows the identified Pareto front, while the solid black lines show the bounds on  $J_{NO_x}^{rot}$ . On the left, the Pareto front described by scaled and rotated values is shown. On the right, the Pareto front transferred back to physically meaningful values is shown. The values in the plot on the right were normalized.

To retrieve the desired engine outputs, the rotation is reversed to obtain the variables  $J_X$  and the scaling is reversed using

$$\dot{H}_{ATS}^* = (1 - J_{Hd}) \dot{H}_{ATS, min}^* + J_{Hd} \dot{H}_{ATS, max}^* \quad (3)$$

$$\dot{m}_{NO_x}^{eo} = (1 - J_{NO_x}) \dot{m}_{NO_x, min}^{eo} + J_{NO_x} \dot{m}_{NO_x, max}^{eo} \quad (4)$$

$$\dot{m}_{fuel} = (1 - J_{fuel}) \dot{m}_{fuel, min}^* + J_{fuel} \dot{m}_{fuel, max}^*, \quad (5)$$

where the fuel mass flow is denoted  $\dot{m}_{fuel}^*$  and the engine-out NO<sub>x</sub> mass flow is denoted  $\dot{m}_{NO_x}^{eo}$ . The scaling values  $\dot{H}_{ATS, min}^*$ ,  $\dot{H}_{ATS, max}^*$ ,  $\dot{m}_{NO_x, min}^{eo}$ ,  $\dot{m}_{NO_x, max}^{eo}$ ,  $\dot{m}_{fuel, min}^*$ , and  $\dot{m}_{fuel, max}^*$  are characterized by three-dimensional look-up elements with inputs  $(\omega_e, T_e, \vartheta)$ . In Figure 2, on the right, the Pareto front transferred back to physically meaningful values is shown.

To operate the engine, the engine control inputs must be determined. For this, a smoothed inverse mapping was developed that returns the engine control inputs as a function of the selected engine strategy inputs. An equivalent mapping is also developed for the exhaust mass flow, which is required to evaluate the ATS model. The mappings are stored in the form of look-up elements with inputs  $(\omega_e, T_e, \vartheta, J_{Hd}^{rot}, J_{NO_x}^{rot})$ .

## 2.2. Aftertreatment System Model

The ATS model used in this work was adopted from [5]. The thermal model consists of a single lumped brick temperature  $\vartheta$ , whose dynamics are

$$\frac{d}{dt}\vartheta = \frac{\dot{H}_{ATS} - \dot{Q}_{hl}}{C_{ATS}} \quad (6)$$

$$\dot{Q}_{hl} = \alpha S (\vartheta - \vartheta_{amb}), \quad (7)$$

where  $C_{ATS}$  is the heat capacity of the ATS. The conductive heat loss to the environment is denoted  $\dot{Q}_{hl}$ , where  $\alpha$  is the heat transfer coefficient,  $S$  is the outer surface area of the ATS, and  $\vartheta_{amb}$  is the ambient temperature.

The chemical model consists of a static look-up element that describes the  $\text{NO}_x$ -reduction efficiency  $\eta_{\text{DeNO}_x}$  as a function of the ATS temperature and the exhaust mass flow. The resulting tailpipe  $\text{NO}_x$  mass flow is

$$\dot{m}_{\text{NO}_x}^{\text{tp}} = \left(1 - \eta_{\text{DeNO}_x}(\vartheta, \dot{m}_{\text{exh}}^*)\right) \dot{m}_{\text{NO}_x}^{\text{eo}}. \quad (8)$$

The dynamics of the emitted tailpipe  $\text{NO}_x$  mass are

$$\frac{d}{dt}m_{\text{NO}_x}^{\text{tp}} = \dot{m}_{\text{NO}_x}^{\text{tp}}. \quad (9)$$

## 3. Real-Time-Capable Optimization Method

The objective of the control task is to minimize the fuel consumed during a given mission, while respecting an upper bound on the emitted tailpipe  $\text{NO}_x$  mass at the end of the mission. The resulting optimal control problem (OCP) is

$$\underset{J_{\text{Hd}}^{\text{rot}}, J_{\text{NO}_x}^{\text{rot}}}{\text{minimize}} \int_0^{t_f} \dot{m}_{\text{fuel}}^*(\omega_e, T_e, \vartheta, J_{\text{Hd}}^{\text{rot}}, J_{\text{NO}_x}^{\text{rot}}) dt \quad (10a)$$

$$\text{subject to} \quad \frac{d}{dt}\vartheta = f_\vartheta(\omega_e, T_e, \vartheta, J_{\text{Hd}}^{\text{rot}}, J_{\text{NO}_x}^{\text{rot}}) \quad (10b)$$

$$\frac{d}{dt}m_{\text{NO}_x}^{\text{tp}} = \dot{m}_{\text{NO}_x}^{\text{tp}}(\omega_e, T_e, \vartheta, J_{\text{Hd}}^{\text{rot}}, J_{\text{NO}_x}^{\text{rot}}) \quad (10c)$$

$$\vartheta(0) = \vartheta_0 \quad (10d)$$

$$m_{\text{NO}_x}^{\text{tp}}(0) = 0 \quad (10e)$$

$$m_{\text{NO}_x}^{\text{tp}}(t_f) \leq \bar{m}_{\text{NO}_x}^{\text{tp}} \quad (10f)$$

$$J_{\text{Hd}}^{\text{rot}} \in [0, 1] \quad (10g)$$

$$J_{\text{NO}_x}^{\text{rot}} \geq J_{\text{NO}_x, \text{min}}^{\text{rot}}(\omega_e, T_e, \vartheta, J_{\text{Hd}}^{\text{rot}}) \quad (10h)$$

$$J_{\text{NO}_x}^{\text{rot}} \leq J_{\text{NO}_x, \text{max}}^{\text{rot}}(\omega_e, T_e, \vartheta, J_{\text{Hd}}^{\text{rot}}), \quad (10i)$$

where  $t_f$  is the final time, the function  $f_\vartheta$  represents (6), and  $\bar{m}_{\text{NO}_x}^{\text{tp}}$  is the limit on the emitted tailpipe  $\text{NO}_x$  mass. The time dependency of all variables was omitted to improve readability.

In this work, the mission is assumed to be set by the driver and must be followed exactly. It is therefore not part of the optimization and consequently  $\omega_e$  and  $T_e$  are exogenous inputs. The system states are  $\vartheta$  and  $m_{\text{NO}_x}^{\text{tp}}$ . The inputs to be optimized are  $J_{\text{Hd}}^{\text{rot}}$  and  $J_{\text{NO}_x}^{\text{rot}}$ .

In [5], dynamic programming (DP) was used to solve the OCP, which is a reasonable choice of method when a global optimality guarantee is important and computational time is not limited. However, its use in an online model predictive controller is severely limited by the constraint on the maximal computation time that must be met to allow for feedback control. In fact, examples from the literature show that the time required to solve the OCP with DP for a mission with a duration of roughly 30 min for similar problems is in the order

of a few hours to one week on a standard PC [15–18]. This clearly disqualifies DP as an online optimization method.

A more efficient way to solve the OCP is to employ a direct method, whereby the OCP is discretized and transformed to a finite-dimensional, constrained nonlinear optimization problem, i.e., an NLP [19]. Considering further that many established NLP solvers exist, this approach is well suited for the task at hand.

Common NLP solvers use first-order derivatives and approximate second-order derivatives using first-order derivatives. Hence, the model equations must be at least once continuously differentiable. Evaluating the look-up elements from Sections 2.1 and 2.2 using linear interpolation results in piecewise affine functions that are not continuously differentiable. A replacement for the look-up tables must therefore be found that can characterize the underlying data with the same accuracy, but is continuously differentiable.

In the following, continuously differentiable look-up elements are introduced and their application as a replacement for the look-up elements from Sections 2.1 and 2.2 is outlined. Next, OCP (10a–i) is discretized and an NLP is formulated. The result obtained by the real-time-capable NLP optimization is compared to a benchmark obtained using DP. Finally, the result from the NLP optimization is validated experimentally.

### 3.1. Continuously Differentiable Look-Up Elements

In this work, a piecewise quadratic interpolant with a continuous derivative is used. Using the terminology from [20], this corresponds to a Hermite quadratic interpolant. The resulting continuously differentiable look-up elements are referred to as cdElements. Their implementation is outlined in Appendix A.

The selected interpolation method is compared to the two predominant interpolation methods, namely linear interpolation and cubic spline interpolation, in Table 2. Compared to linear interpolation, Hermite quadratic interpolation has the advantage of continuous differentiability, whereas, compared to cubic spline interpolation, it has the advantage of less computations required for the evaluation.

**Table 2.** Comparison of interpolation methods.

Method	Continuously Differentiable	Evaluation Complexity
linear	no	low
Hermite quadratic	once	intermediate
cubic spline	twice	high

The model presented in Section 2.1 uses a five-dimensional look-up element to describe the preoptimized engine operation. According to OCP (10a–i), the model equations only have to be continuously differentiable with respect to  $\vartheta$ ,  $J_{\text{Hd}}^{\text{rot}}$ , and  $J_{\text{NO}_x}^{\text{rot}}$ , as only these can be influenced by the optimization. As  $\omega_e$  and  $T_e$  are given and cannot be affected by the optimization, continuous differentiability in these directions is not required.

Therefore, local continuously differentiable mappings with inputs ( $\vartheta$ ,  $J_{\text{Hd}}^{\text{rot}}$ ,  $J_{\text{NO}_x}^{\text{rot}}$ ) are stored for a number of ( $\omega_e$ ,  $T_e$ ) points on a regular grid. When evaluating the global model, the local models are evaluated at the four ( $\omega_e$ ,  $T_e$ ) grid points surrounding the current point and bilinear interpolation is used to return the corresponding values for the current point. At each ( $\omega_e$ ,  $T_e$ ) grid point, two three-dimensional cdElements were used to characterize  $J_{\text{fuel}}^{\text{rot}}$  and  $\dot{m}_{\text{exh}}^*$ , two two-dimensional cdElements were used to characterize the upper and lower bound on  $J_{\text{NO}_x}^{\text{rot}}$ , and six one-dimensional cdElements were used to characterize the normalization values. Further, a two-dimensional cdElement was used to characterize the NO<sub>x</sub>-reduction efficiency map of the ATS used in (8).

### 3.2. NLP Formulation

In order to transform OCP (10a–i) to an NLP, it was discretized in time using a constant time step  $\Delta t = 1$  s and Euler forward integration. Multiple shooting was used for the ATS

temperature, as the state dynamics (10b) and (10c) depend on the ATS temperature in a nonlinear way [21]. For the emitted tailpipe NO<sub>x</sub> mass state, on the other hand, single shooting was used, as the state dynamics do not depend on the emitted tailpipe NO<sub>x</sub> mass state. The state dynamics (10c), the initial condition (10e), and the terminal constraint (10f) are combined to formulate (11d). The resulting NLP is

$$\underset{\vartheta, J_{\text{Hd}}^{\text{rot}}, J_{\text{NO}_x}^{\text{rot}}}{\text{minimize}} \quad \sum_{k=0}^N \dot{m}_{\text{fuel}}^*(\cdot) \Delta t \quad (11a)$$

$$\text{subject to} \quad \vartheta[0] = \vartheta_0 \quad (11b)$$

$$\vartheta[k+1] - \vartheta[k] - f_{\vartheta}(\cdot) \Delta t = 0 \quad k = 0, \dots, N \quad (11c)$$

$$\sum_{k=0}^N \dot{m}_{\text{NO}_x}^{\text{tp}}(\cdot) \Delta t \leq \bar{m}_{\text{NO}_x}^{\text{tp}} \quad k = 0, \dots, N \quad (11d)$$

$$J_{\text{Hd}}[k] \in [0, 1] \quad k = 0, \dots, N \quad (11e)$$

$$J_{\text{NO}_x}^{\text{rot}}[k] \geq J_{\text{NO}_x, \text{min}}^{\text{rot}}(\cdot) \quad k = 0, \dots, N \quad (11f)$$

$$J_{\text{NO}_x}^{\text{rot}}[k] \leq J_{\text{NO}_x, \text{max}}^{\text{rot}}(\cdot) \quad k = 0, \dots, N, \quad (11g)$$

where the arguments of  $\dot{m}_{\text{fuel}}^*$ ,  $f_{\vartheta}$ ,  $\dot{m}_{\text{NO}_x}^{\text{tp}}$ ,  $J_{\text{NO}_x, \text{min}}^{\text{rot}}$ , and  $J_{\text{NO}_x, \text{max}}^{\text{rot}}$  are equivalent to those in OCP (10a–i) but have been omitted to improve readability. The number of considered time steps  $N$  is

$$N = \frac{t_f}{\Delta t} - 1. \quad (12)$$

In this work, the NLP solver IPOPT [22] was used to solve NLP (11a–g).

### 3.3. Comparison to DP Benchmark

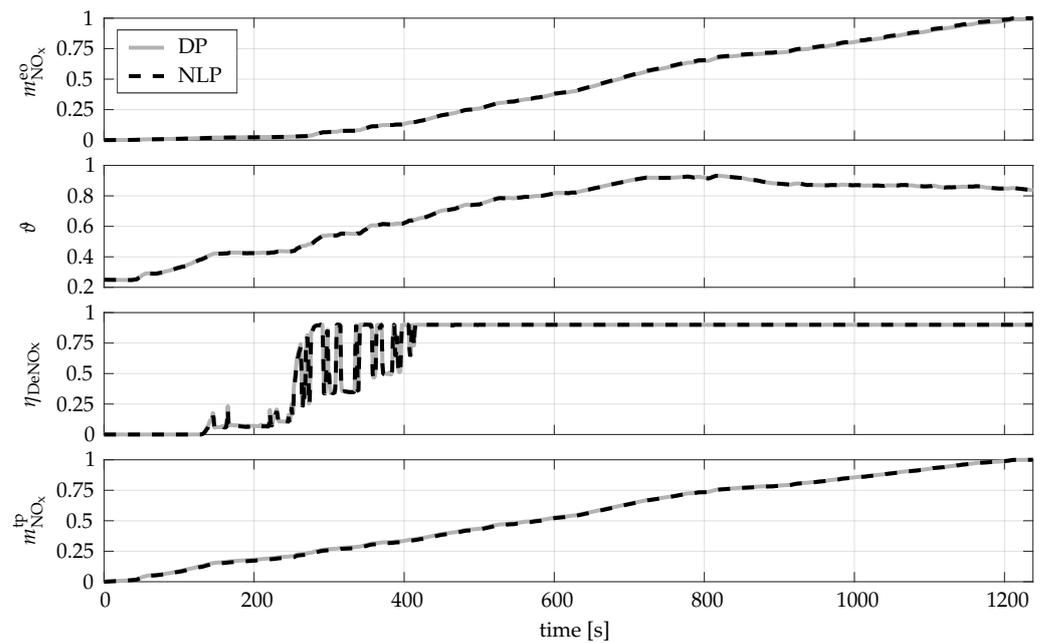
A drawback of using a direct method is that there are no global optimality guarantees for nonconvex problems such as NLP (11a–g). To check the quality of the solution obtained from the NLP optimization, it was compared to a benchmark obtained using DP. As the goal of the DP optimization is to find a benchmark, a fine discretization was used. Specifically, 601 grid points were used for  $\vartheta$  and 41 for  $J_{\text{Hd}}$  and  $J_{\text{NO}_x}^{\text{rot}}$ . An adapted version of the DP algorithm presented in [23] was used to solve the problem.

The comparison of the trajectories resulting from the DP and the NLP optimizations is shown in Figure 3. No discernible difference in the trajectories can be observed. For a closer analysis, the results listed in Table 3 are considered. The computation times reported were obtained on a PC with a 2.8 GHz quad-core processor and 32 GB of RAM. The results show that the NLP solution fulfills the constraint on the emitted tailpipe NO<sub>x</sub> mass exactly. The consumed fuel mass was even 0.049% lower than that of the DP solution. This can be attributed to the fact that even though a fine discretization was used in the DP optimization, the error introduced by the discretization is still nonzero, while the NLP can set the inputs continuously and manages to find a superior solution for the considered case. While this example does not prove that the NLP solver converges to the global optimum in general, it indicates that the resulting solution is of high quality. Similar results were obtained for further missions.

The developed optimization method is a viable option for online optimization, as the time required to solve the NLP is significantly shorter than the duration of the cycle.

### 3.4. Experimental Validation of Solution

Next, the obtained solution is validated at the testbench by applying the obtained optimal input trajectories in a feedforward fashion. The reference trajectories and the measured trajectories shown in Figure 4 exhibit the same general trends, but it can be observed that the simple control-oriented model is not able to fully capture the behavior of the more complex real system. From the comparison, three main sources of model mismatch can be identified.



**Figure 3.** Comparison of the solution to the nonlinear program (NLP) to the dynamic programming (DP) benchmark for the nonroad transient cycle (NRTC) with  $\vartheta(0) = 100^\circ\text{C}$ . Shown are the emitted engine-out  $\text{NO}_x$  mass  $m_{\text{NO}_x}^{\text{eo}}$ , the ATS temperature  $\vartheta$ , the  $\text{NO}_x$ -reduction efficiency  $\eta_{\text{DeNO}_x}$ , and the emitted tailpipe  $\text{NO}_x$  mass  $m_{\text{NO}_x}^{\text{tp}}$ . All values were normalized.

**Table 3.** Comparison of the DP and NLP optimizations for the NRTC with  $\vartheta(0) = 100^\circ\text{C}$ . The consumed fuel mass is denoted  $m_{\text{fuel}}$  and the required computation time is denoted  $t_{\text{calc}}$ . The results of the DP optimization were used for the normalization.

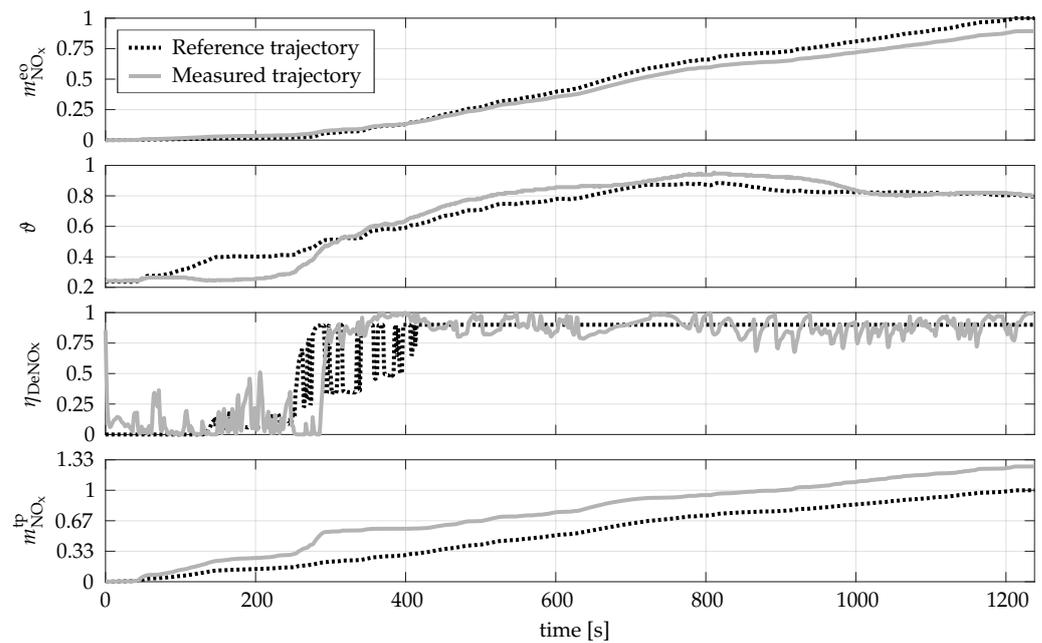
Method	$t_{\text{calc}}$	$m_{\text{fuel}}[\%]$	$m_{\text{NO}_x}^{\text{eo}}[\%]$	$m_{\text{NO}_x}^{\text{tp}}[\%]$
DP	15.8 h	100	100	100
NLP	23.4 s	99.951	100.943	100.000

The first error source is the static engine model. The observed error in the emitted engine-out  $\text{NO}_x$  mass was  $-9.3\%$ , while the error in the consumed fuel was  $+2.9\%$ .

The second error source is the simplified thermal model of the ATS that uses a single lumped temperature. Its effect is clearly visible during the initial 300 s of the considered mission. As the engine produces hot exhaust gases, the modeled ATS temperature rises. The temperature measured after the ATS, however, remains low, as there is a non-negligible axial temperature distribution in the real system that is not captured by the model. After this initial warm-up phase, the modeled temperature captures the actual dynamics reasonably well.

The third error source is the static map used to characterize the  $\text{NO}_x$ -reduction efficiency in the chemical model of the ATS. Neglecting the dynamics of the ammonia storage of the SCR, combined with the error in the ATS temperature, leads to a mismatch in  $\eta_{\text{DeNO}_x}$ . The mismatch in the timing of the sharp rise in  $\eta_{\text{DeNO}_x}$  at around 300 s is caused mainly by the error in the ATS temperature, while the mismatch at higher temperatures from 400 s onward is attributed mainly to the neglected ammonia storage dynamics.

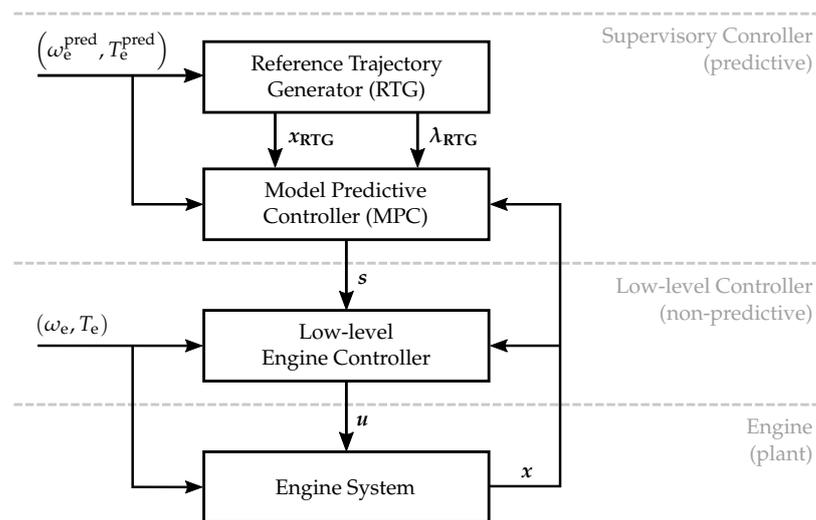
The emitted tailpipe  $\text{NO}_x$  mass is affected by all three of the error sources, resulting in an overall error of  $+27.9\%$ . In conclusion, the simple control-oriented model captures the basic system behavior, but feedback control is required in order to compensate for model mismatch during operation.



**Figure 4.** Comparison of the reference trajectories resulting from solving optimal control problem (10a–i) (dotted black) and the corresponding measured trajectories obtained by applying the input trajectories in a feedforward fashion (solid gray) for the NRTC with  $\vartheta(0) = 100^\circ\text{C}$ . All values were normalized.

#### 4. Predictive Supervisory Controller

The controller structure used in this work is shown in Figure 5. It consists of a predictive supervisory controller that passes a strategy input vector  $\mathbf{s}$  to a non-predictive low-level controller that sets the engine control input vector  $\mathbf{u}$ . The vector of measured system states  $\mathbf{x}$  is used as feedback. Here, bold symbols are used to denote multidimensional vectors.



**Figure 5.** Controller structure consisting of a predictive supervisory controller and a non-predictive low-level engine controller.

The predictive supervisory controller, in turn, contains two layers. In both layers, an optimization problem related to OCP (10a–i) is solved to determine the optimal operation over the prediction horizon. The reference trajectory generator (RTG) considers the entire mission and is run once at the beginning of the mission. It provides reference trajectories that are tracked by the model predictive controller (MPC), which optimizes the engine

operation periodically for a shorter prediction horizon and considers feedback of the measured system states. The MPC's prediction horizon is chosen to be much shorter than that of the RTG, resulting in a shorter computation time and allowing for frequent updates.

At the testbench, the predictive supervisory controller was implemented in Matlab and ran on a PC with a 2.4 GHz dual-core processor and 16 GB of RAM, while the low-level engine controller ran on an ETAS ES910 rapid prototyping unit. INCA's Matlab integration package was used for the communication between the predictive supervisory controller and the low-level engine controller. The stock engine ECU was used to set the engine control inputs.

In the following, the elements of the controller are discussed in more detail.

#### 4.1. Reference Trajectory Generator

The task of the RTG is to optimize the operation of the engine system over the entire mission and provide reference trajectories that the MPC will follow. To achieve this, the RTG solves OCP (10a–i) for the entire mission using the optimization method presented in Section 3. The resulting optimal trajectories for the states as well as the corresponding costates are passed to the MPC as

$$\mathbf{x}_{\text{RTG}} = \begin{bmatrix} \vartheta_{\text{RTG}} \\ m_{\text{NO}_x, \text{RTG}}^{\text{tp}} \end{bmatrix} \quad \text{and} \quad \boldsymbol{\lambda}_{\text{RTG}} = \begin{bmatrix} \lambda_{\vartheta, \text{RTG}} \\ \lambda_{m, \text{RTG}} \end{bmatrix}. \quad (13)$$

In this work, the predictive information is considered to be of high quality, i.e., there is limited benefit in updating the reference trajectories during the mission. However, the RTG could be updated if a change in the mission or a significant offset of the current states from their respective reference is detected.

#### 4.2. Model Predictive Controller

The task of the MPC is to correct for errors introduced by model mismatch and misprediction in an optimal fashion. To achieve this, an OCP similar to OCP (10a–i) is solved at every update time  $t_u$  and the resulting strategy inputs are passed to the low-level controller. The OCP is initialized with the current ATS temperature  $\vartheta_0$  and emitted tailpipe  $\text{NO}_x$  mass  $m_{\text{NO}_x, 0}^{\text{tp}}$ . The terminal state constraints are chosen such that the emitted tailpipe  $\text{NO}_x$  mass does not exceed that of the reference and the ATS temperature is at least at the reference value at the end of the MPC's prediction horizon  $t_p$ .

The OCP solved by the MPC is

$$\underset{J_{\text{Hd}}^{\text{rot}}, J_{\text{NO}_x}^{\text{rot}}, \epsilon_{\vartheta}, \epsilon_m}{\text{minimize}} \quad \int_{t_0}^{t_0+t_p} \dot{m}_{\text{fuel}}^*(\omega_e, T_e, \vartheta, J_{\text{Hd}}, J_{\text{NO}_x}^{\text{rot}}) dt - \lambda_{\vartheta, \text{RTG}} \epsilon_{\vartheta} + w_m \epsilon_m \quad (14a)$$

$$\text{subject to} \quad \frac{d}{dt} \vartheta = f_{\vartheta}(\omega_e, T_e, \vartheta, J_{\text{Hd}}, J_{\text{NO}_x}^{\text{rot}}) \quad (14b)$$

$$\frac{d}{dt} m_{\text{NO}_x}^{\text{tp}} = \dot{m}_{\text{NO}_x}^{\text{tp}}(\omega_e, T_e, \vartheta, J_{\text{Hd}}, J_{\text{NO}_x}^{\text{rot}}) \quad (14c)$$

$$\vartheta(t_0) = \vartheta_0 \quad (14d)$$

$$\vartheta(t_0 + t_p) \geq \vartheta_{\text{RTG}} - \epsilon_{\vartheta} \quad (14e)$$

$$m_{\text{NO}_x}^{\text{tp}}(t_0) = m_{\text{NO}_x, 0}^{\text{tp}} \quad (14f)$$

$$m_{\text{NO}_x}^{\text{tp}}(t_0 + t_p) \leq \bar{m}_{\text{NO}_x, \text{RTG}}^{\text{tp}} + \epsilon_m \quad (14g)$$

$$J_{\text{Hd}} \in [0, 1] \quad (14h)$$

$$J_{\text{NO}_x}^{\text{rot}} \geq J_{\text{NO}_x, \text{min}}^{\text{rot}}(\omega_e, T_e, \vartheta, J_{\text{Hd}}) \quad (14i)$$

$$J_{\text{NO}_x}^{\text{rot}} \leq J_{\text{NO}_x, \text{max}}^{\text{rot}}(\omega_e, T_e, \vartheta, J_{\text{Hd}}) \quad (14j)$$

$$\epsilon_{\vartheta}, \epsilon_m \geq 0. \quad (14k)$$

The time dependency of all variables was omitted to improve readability. OCP (14a–k) is solved by adapting the method presented in Section 3. The resulting MPC is a zone-MPC [24], whereby the fuel consumption is minimized, while an upper bound on the emitted tailpipe NO<sub>x</sub> mass and a lower bound on the terminal ATS temperature are respected.

The terminal constraint on the emitted tailpipe NO<sub>x</sub> mass is given by  $\bar{m}_{NO_x,RTG}^{tp}$ , which represents the emitted tailpipe NO<sub>x</sub> mass of the RTG reference trajectory at the end of the MPC’s prediction horizon. A soft constraint was formulated using the slack variable  $\epsilon_m$ , in order to always allow the solver to find a solution to the OCP. The activation of the slack variable is penalized linearly using  $w_m$ . To avoid violation of the soft constraint when the problem is feasible, a large, constant weight  $w_m$  was used. This guarantees that constraint (14g) is fulfilled unless it is physically impossible to do so and that the emitted tailpipe NO<sub>x</sub> mass is minimized when the non-slacked problem becomes infeasible.

When the entire mission is considered, there is no constraint on the terminal ATS temperature, as it does not influence the tailpipe NO<sub>x</sub> emissions. However, the MPC’s prediction horizon does generally not reach the end of the mission, meaning that the ATS temperature at the end of the MPC’s prediction horizon will influence tailpipe NO<sub>x</sub> emissions later on in the mission. It can therefore be beneficial to enforce a terminal constraint on the ATS temperature to shift the engine operation to utilize more fuel to heat up the ATS in the considered prediction horizon, as this can improve the NO<sub>x</sub>-reduction efficiency for the remainder of the mission. To account for this, constraint (14e) is introduced, where  $\vartheta_{RTG}$  represents the ATS temperature of the RTG reference trajectory at the end of the MPC’s prediction horizon. As for the tailpipe NO<sub>x</sub> emissions, a soft constraint was formulated using the slack variable  $\epsilon_\vartheta$ . As constraint (14e) is only used to steer the MPC in the direction of the reference trajectory, a less aggressive penalty on  $\epsilon_\vartheta$  can be used. In [25], it was shown that the costate from the RTG can be used to quantify the sensitivity of the objective of OCP (10a–i) to a drift from the reference trajectory in a first-order approximation. The ATS temperature costate of the RTG reference at the end of the MPC’s prediction horizon  $\lambda_{\vartheta,RTG}$  is therefore used as the weight for the penalty on  $\epsilon_\vartheta$ .

#### 4.2.1. MPC Tuning

The MPC’s tuning parameters are the update time  $t_u$  and the prediction horizon  $t_p$ . They were selected by considering the controller performance for a grid of combinations thereof and ensuring that the average computation time required to solve OCP (14a–k) is well below the update time. In the following,  $t_u = 20$  s and  $t_p = 200$  s were used.

Note that the MPC’s update time  $t_u = 20$  s is larger than the sampling time  $\Delta t = 1$  s. Therefore, a trajectory of strategy inputs is passed to the low-level engine control rather than a scalar value.

#### 4.2.2. Offset Correction

An offset correction based on [26] was introduced to improve the MPC’s tracking performance. Thereby, the MPC’s terminal constraints are adapted using the current offset. For the emitted tailpipe NO<sub>x</sub> mass, it results in

$$e_m = m_{NO_x,RTG}^{tp}(t_0) - m_{NO_x}^{tp}(t_0) \tag{15}$$

$$\bar{m}_{NO_x,RTG}^{tp} = m_{NO_x,RTG}^{tp}(t_0 + t_p) + e_m, \tag{16}$$

while, for the ATS temperature, this results in

$$e_\vartheta = \vartheta_{RTG}(t_0) - \vartheta(t_0) \tag{17}$$

$$\vartheta_{RTG} = \vartheta_{RTG}(t_0 + t_p) + e_\vartheta. \tag{18}$$

The theory behind this correction guarantees offset-free tracking in steady state for linear MPC [21]. Even though the references never reach a steady state and nonlinear MPC is used here, the method was found to improve the tracking behavior of the MPC.

#### 4.3. Low-Level Engine Controller

In this work, two realizations of the low-level engine controller are considered. The task of both realizations is to set the engine control input vector as stated in (1). However, the two realizations differ in the strategy input vector that they receive and in how the engine control inputs are determined. The two realizations are introduced in Sections 5.1 and 6.1, respectively.

### 5. Controller Performance without Misprediction

In this section, the performance of the predictive supervisory controller with perfect predictive information, i.e., without mispredictions, is analyzed experimentally. In order to obtain comparable and repeatable results, the ATS was preconditioned by emptying the ammonia storage and cooling it to 100 °C before each experiment. As the initial ATS temperature is well below the temperature where the injection of AdBlue is enabled, the result is similar to a cold start.

#### 5.1. Low-Level Engine Controller

The low-level engine controller used in this section evaluates the smoothed inverse mapping introduced in Section 2.1 to determine the engine control inputs. The maps are evaluated at the current engine speed and torque, the current ATS temperature, and the engine strategy inputs selected by the MPC. The strategy input vector becomes

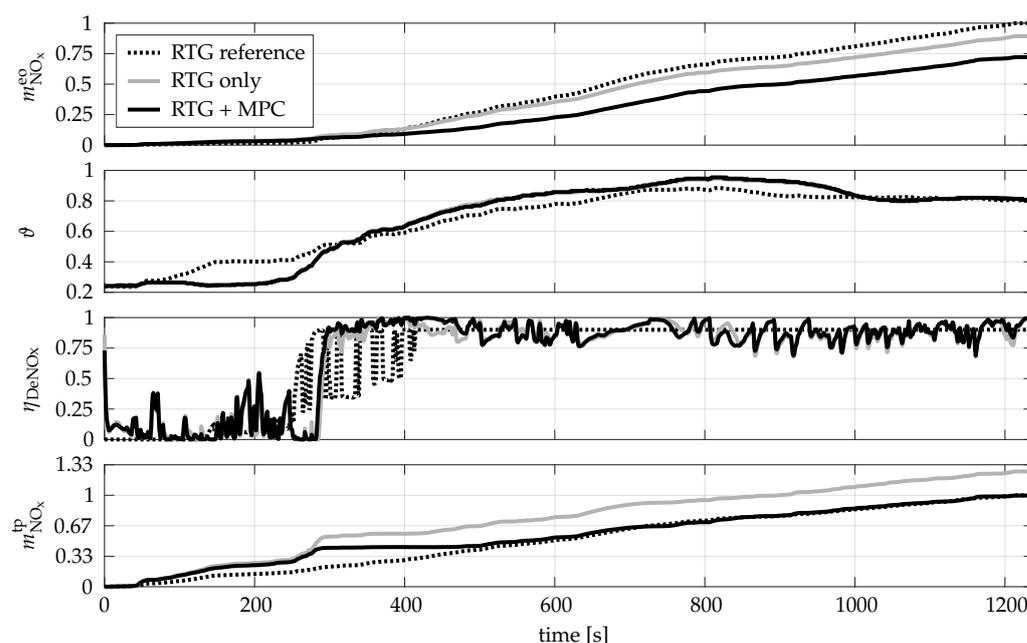
$$\mathbf{s} = \begin{bmatrix} J^{\text{rot}} \\ J^{\text{Hd}} \\ J^{\text{rot}}_{\text{NO}_x} \end{bmatrix}. \quad (19)$$

Note that the evaluation of the smoothed inverse mapping is not part of the optimization itself, but is carried out afterwards. Hence, continuous differentiability is not required and the smoothed input maps from Section 2.1 can be applied directly. This realization of the low-level engine controller is computationally inexpensive as finding the optimal engine control inputs is reduced to evaluating a mapping obtained offline.

#### 5.2. Reference Tracking

First, the reference tracking capability of the controller is investigated; see Figure 6. For the first 300 s, the operations with and without the MPC were equivalent. Due to the model error discussed in Section 3.4, the emitted tailpipe NO<sub>x</sub> mass lies above the reference. During this time, the ATS temperature was lower than expected, resulting in a lower NO<sub>x</sub>-reduction efficiency. Considering further that the operation without the MPC already produced minimal engine-out NO<sub>x</sub> emissions, it becomes evident that a saturation is active, and even when feedback is introduced via the MPC, the emitted tailpipe NO<sub>x</sub> mass cannot be reduced further in this part of the mission.

At 300 s, the measured NO<sub>x</sub>-reduction efficiency rose rapidly and the operation without the MPC moved to an engine operation with higher engine-out NO<sub>x</sub> emissions. The operation with the MPC, on the other hand, continued to operate the engine with minimal engine-out NO<sub>x</sub> emissions until the tailpipe NO<sub>x</sub> reference was reached at around 500 s. From there on, the operation with the MPC also switched to an engine operation with higher engine-out NO<sub>x</sub> emissions and continued to track the tailpipe NO<sub>x</sub> reference until the end of the mission.



**Figure 6.** Resulting trajectories for the NRTC. The dotted gray trajectories were obtained by the RTG optimization. The solid gray trajectories show the measured operation with only the RTG, i.e., without feedback from the MPC. The solid black trajectories show the measured operation with the full predictive supervisory controller, i.e., with feedback from the MPC. All values were normalized.

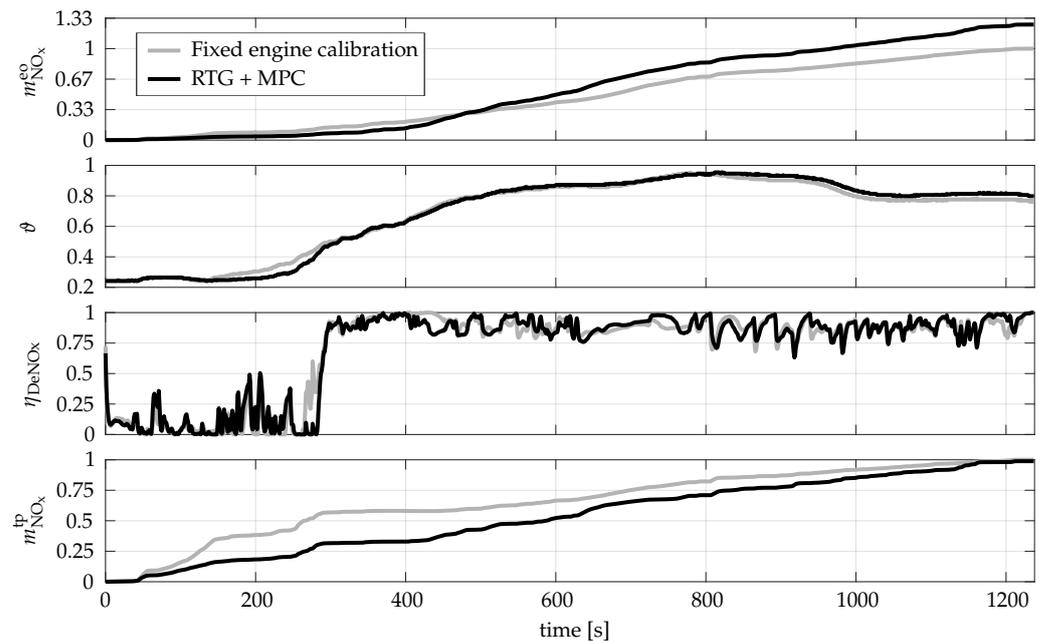
The result shows that the effect of the model mismatch can be compensated for by the MPC when the operation is not limited by a saturation such as the lower limit on the engine-out  $\text{NO}_x$  emissions. When the reference trajectories can no longer be tracked exactly, the MPC follows them as well as the system allows. This can be observed during the initial 300 s of the mission, where the controller tolerated a violation of the terminal constraint on  $\theta$  in order to get as close to the  $m_{\text{NO}_x}^{\text{tp}}$  reference as possible.

### 5.3. Comparison to Fixed Engine Calibration

Next, the performance of the predictive supervisory controller is compared to that achieved using the fixed engine calibration of the stock ECU; see Figure 7. In the absence of a supervisory controller, the engine-out  $\text{NO}_x$  emissions depend solely on the mission. High tailpipe  $\text{NO}_x$  emissions result when the ATS is inactive (initial 300 s) and low tailpipe  $\text{NO}_x$  emissions result once the ATS has reached its operating temperature (from 300 s onward).

The tailpipe  $\text{NO}_x$  mass emitted by the fixed engine calibration was used as the limit for the predictive supervisory controller. During the initial 300 s of the cycle, the predictive supervisory controller adapted the engine operation to the low  $\text{NO}_x$ -reduction efficiency by reducing the engine-out  $\text{NO}_x$  emissions. At the end of this phase, the tailpipe  $\text{NO}_x$  mass emitted by the predictive supervisory controller was roughly half of that achieved with the fixed engine calibration. Once the ATS was fully operational, the predictive supervisory controller reacted by allowing higher engine-out  $\text{NO}_x$  emissions, thereby reducing the fuel consumption.

The final emitted tailpipe  $\text{NO}_x$  mass of the predictive supervisory controller was equivalent to that achieved with the fixed engine calibration, while the fuel consumption was reduced by 1.1%.



**Figure 7.** Comparison of the measurement obtained with the fixed engine calibration (gray) to that obtained with the predictive supervisory controller (black) for the NRTC. Using the predictive supervisory controller reduced the fuel consumption by 1.1%. All values were normalized.

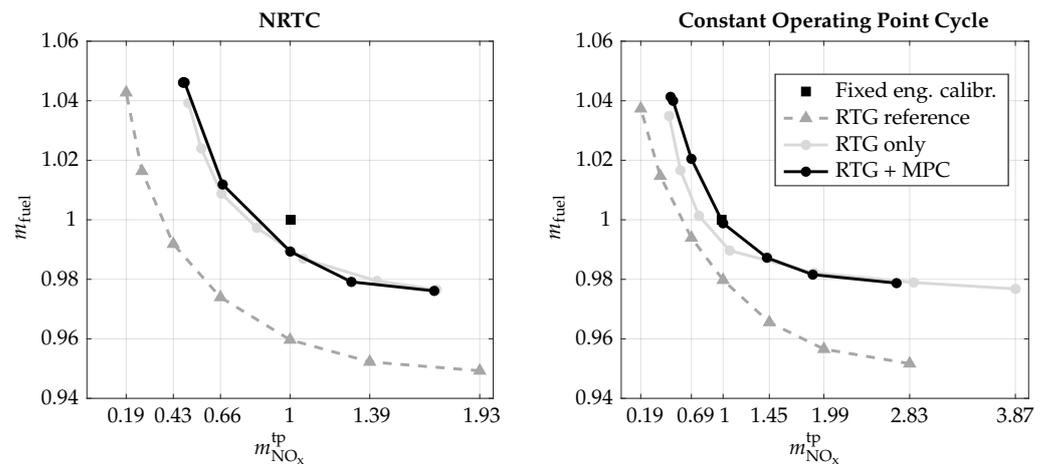
#### 5.4. Controllability Limits

On the left of Figure 8, the resulting operation for the nonroad transient cycle (NRTC) and different NO<sub>x</sub> limits is shown. The difference between the RTG reference (triangles) and the measurements (circles) results from model mismatch. The results obtained without feedback (gray circles) always show a deviation from the tailpipe NO<sub>x</sub> limit, which is reduced when feedback is introduced (black circles). The tailpipe NO<sub>x</sub> limit was only respected and fully exploited for  $\bar{m}_{\text{NO}_x}^{\text{tp}} = 0.66$  and  $\bar{m}_{\text{NO}_x}^{\text{tp}} = 1$ , while more stringent limits were violated and more relaxed limits were not fully exploited.

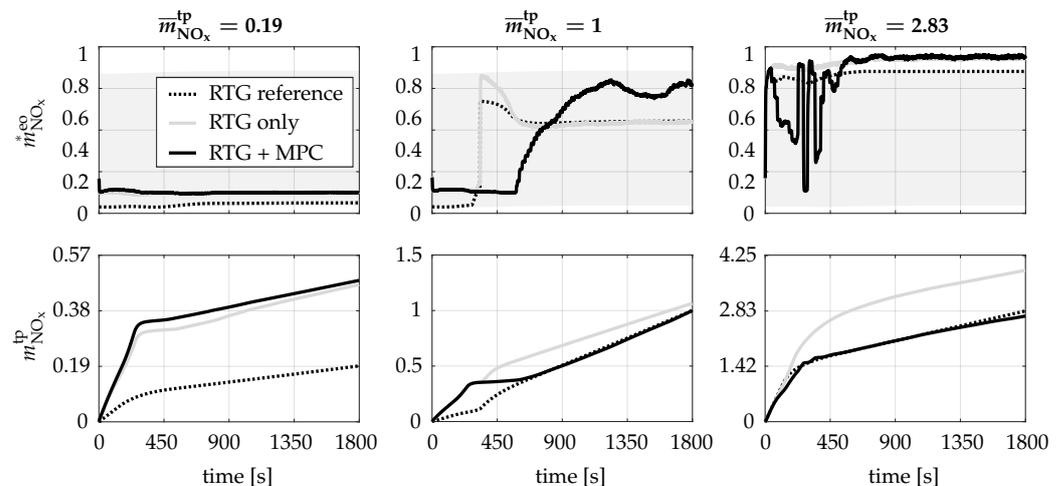
To investigate this behavior, the constant operating point cycle is considered. This mission consists of a constant engine speed of 1900 rpm and a constant engine torque of 200 Nm for the duration of 1800 s. Note that as the ATS is preconditioned to  $\vartheta(0) = 100$  °C, a dynamic operation results as the ATS warms up, even though the engine speed and torque are constant. The benefit of considering such a cycle is that any changes in the engine operation can be attributed to the predictive supervisory controller, rather than to a change in the engine speed or torque. The results for the constant operating point cycle are shown on the right of Figure 8.

As for the NRTC, intermediate tailpipe NO<sub>x</sub> limits were respected and fully exploited by the predictive supervisory controller, whereas stringent limits were violated and relaxed limits were not fully exploited. Figure 9 shows the resulting trajectories for a representative of each of the three cases.

For an intermediate NO<sub>x</sub> limit, i.e.,  $\bar{m}_{\text{NO}_x}^{\text{tp}} = 1$ , the RTG set the engine-out NO<sub>x</sub> emissions to the minimum during the initial 300 s of the cycle. Due to model mismatch, these low engine-out NO<sub>x</sub> emissions could not be achieved at the testbench, resulting in the emitted tailpipe NO<sub>x</sub> mass rising faster than the reference. The predictive supervisory controller reacted by keeping the engine-out NO<sub>x</sub> emissions at the minimum until the tailpipe NO<sub>x</sub> reference was reached at around 600 s. From there on, the tailpipe NO<sub>x</sub> reference was tracked well, even allowing the engine-out NO<sub>x</sub> emissions to exceed those of the RTG.



**Figure 8.** Pareto fronts obtained by setting different limits on the emitted tailpipe  $\text{NO}_x$  mass. The black square marker shows the operation with the fixed engine calibration. The gray triangular markers show the results from the RTG optimization for different  $\text{NO}_x$  limits. The gray circular markers show the corresponding measured operation with the RTG result used in a purely feedforward fashion, i.e., without feedback from the MPC. The black circular markers show the corresponding measured operation with the full predictive supervisory controller, i.e., with feedback from the MPC. The plot on the left shows the results for the NRTC, while the plot on the right shows the results for a mission consisting of a constant speed of 1900 rpm and a constant torque of 200 Nm for the duration of 1800 s. All values were normalized.



**Figure 9.** Resulting trajectories for the constant operating point cycle for different  $\text{NO}_x$  limits. The gray regions in the upper row of plots show the modeled feasible range for the engine-out  $\text{NO}_x$  mass flow. The dotted black trajectories were obtained by the RTG optimization. The solid gray trajectories show the measured operation with only the RTG, i.e., without feedback from the MPC. The solid black trajectories show the measured operation with the full predictive supervisory controller, i.e., with feedback from the MPC. All values were normalized.

For minimal tailpipe  $\text{NO}_x$  emissions, i.e.,  $\bar{m}_{\text{NO}_x}^{\text{tp}} = 0.19$ , the RTG operated the engine at minimal engine-out  $\text{NO}_x$  emissions throughout the cycle. As previously observed, the actual engine-out  $\text{NO}_x$  emissions were significantly higher. As a consequence, the emitted tailpipe  $\text{NO}_x$  mass exceeded the final limit within the first 300 s of the cycle and the predictive supervisory controller continued to keep  $\text{NO}_x$  emissions as low as possible. The observed violation of the  $\text{NO}_x$  limit must therefore be attributed to a lower limit on the engine-out  $\text{NO}_x$  emissions that is higher than the minimal engine-out  $\text{NO}_x$  emissions

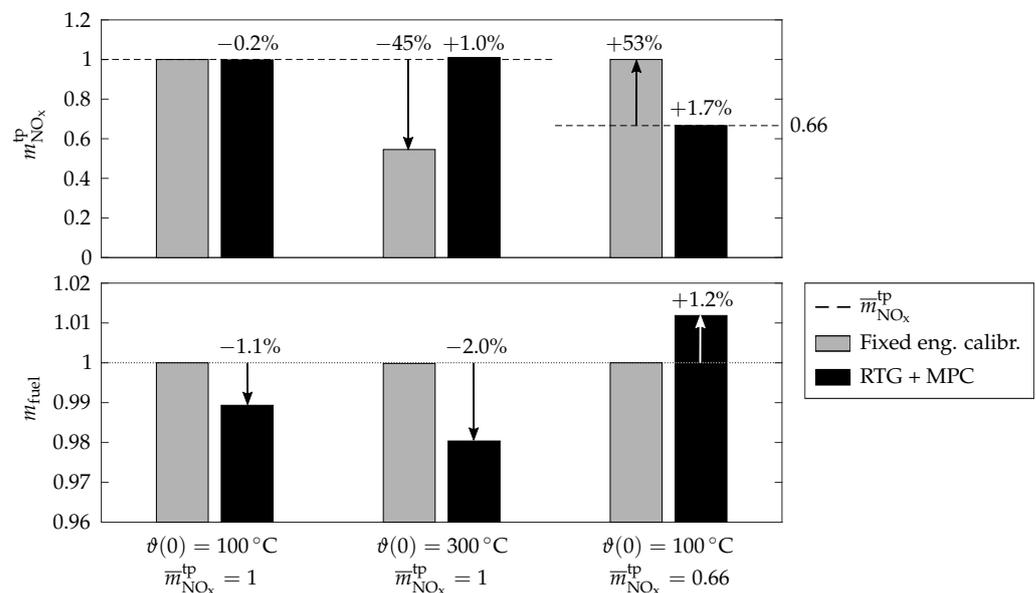
of the model. The MPC is confronted with a lower saturation that limits its reference tracking ability.

For maximal tailpipe NO<sub>x</sub> emissions, i.e.,  $\bar{m}_{\text{NO}_x}^{\text{tp}} = 2.83$ , the RTG operated the engine at maximal engine-out NO<sub>x</sub> emissions throughout the cycle. The resulting operation of the predictive supervisory controller can be divided into two phases. For the initial 500 s, the MPC controller adapted the engine-out NO<sub>x</sub> emissions in such a way that the tailpipe NO<sub>x</sub> reference was tracked. From 500 s onward, the engine-out NO<sub>x</sub> emissions were at the maximum. Due to model mismatch, the actual engine-out NO<sub>x</sub> emissions exceeded the value predicted by the model. However, as the NO<sub>x</sub>-reduction efficiency was also higher than modeled, the resulting tailpipe NO<sub>x</sub> emissions still fell below the reference. While the MPC tried to compensate for this mismatch, in this case, it is confronted with an upper saturation that limits its reference tracking ability.

In summary, the predictive supervisory controller showed good reference tracking behavior when the system bounds allowed for sufficient controllability. When the reference obtained from the RTG could not be achieved with the real system due to saturation, the predictive supervisory controller continued to track the reference as well as the real system allowed.

### 5.5. Benefit of Online Control

Finally, the potential of the predictive supervisory controller to adapt to different conditions and tailpipe NO<sub>x</sub> limits is analyzed; see Figure 10. The comparison on the left corresponds to the one analyzed in Section 5.3. The measured operation obtained with the predictive supervisory controller shows that the tailpipe NO<sub>x</sub> limit was respected and that the fuel consumption was reduced by 1.1%.



**Figure 10.** Comparison of the measured operation for the NRTC for different initial ATS temperatures and tailpipe NO<sub>x</sub> limits. The results obtained with the fixed engine calibration are shown in gray, while the results obtained with the predictive supervisory controller are shown in black. The results on the left and the right correspond to points on the Pareto front on the left of Figure 8. All values were normalized with respect to the operation with the fixed engine calibration for  $\vartheta(0) = 100^\circ\text{C}$  and  $\bar{m}_{\text{NO}_x}^{\text{tp}} = 1$ .

For the comparison in the middle, the same tailpipe NO<sub>x</sub> limit as before was set, but the ATS was preconditioned to a temperature of 300 °C and the corresponding steady-state value for the ammonia storage. As a result, the average NO<sub>x</sub>-reduction efficiency was much higher. The operation obtained with the fixed engine calibration resulted in a final emitted

tailpipe NO<sub>x</sub> mass 45% below the limit. The fixed engine calibration is too conservative for this mission and fuel could have been saved by adapting it. This was achieved by the predictive supervisory controller, which exploited the fact that the tailpipe NO<sub>x</sub> limit was less restrictive, resulting in a fuel reduction of 2.0% while exceeding the tailpipe NO<sub>x</sub> limit by 1.0%.

For the comparison on the right, the initial ATS temperature was 100 °C but the tailpipe NO<sub>x</sub> limit was reduced to 0.66. The operation obtained with the fixed engine calibration reached a final emitted tailpipe NO<sub>x</sub> mass 53% above the limit. The supervisory controller, on the other hand, adapted to the new tailpipe NO<sub>x</sub> limit and only exceeded it by 1.7%. This was achieved at the cost of a 1.2% increase in fuel consumption.

In summary, the predictive supervisory controller was shown to adapt well to changes in the initial ATS temperature and the tailpipe NO<sub>x</sub> limit, adapting to minimize the fuel consumption, while never exceeding the tailpipe NO<sub>x</sub> limit by more than 1.7%.

## 6. Controller Performance with Misprediction

The drawback of the low-level engine controller used in the previous section is that it is not robust against mispredictions. The task of the supervisory controller is to select the optimal trade-off between the individual engine objectives. This is achieved by selecting the engine strategy inputs  $J_{\text{Hd}}^{\text{rot}}$  and  $J_{\text{NO}_x}^{\text{rot}}$ , which describe a relative position on the Pareto front. However, when mispredictions occur, the form of the Pareto front and the trade-off described by  $J_{\text{Hd}}^{\text{rot}}$  and  $J_{\text{NO}_x}^{\text{rot}}$  changes, resulting in an engine operation that will be suboptimal for the current mission. Furthermore, as the bounds on  $J_{\text{NO}_x}^{\text{rot}}$  change with the engine speed and torque, it is possible that the selected point lies outside the bounds. As a consequence, the engine strategy inputs  $J_{\text{Hd}}^{\text{rot}}$  and  $J_{\text{NO}_x}^{\text{rot}}$  are not suited to an operation with misprediction. An alternative low-level controller, suited to deal with misprediction, is introduced and validated in this section.

### 6.1. Low-Level Engine Controller

The realization of the low-level engine controller used in this section determines the engine control inputs by solving a static optimization problem online. The optimization performed by the predictive supervisory controller is equivalent to that of the previous section, but the engine strategy inputs provided to the low-level engine controller change to

$$s = \begin{bmatrix} w_{\text{Hd}} \\ w_{\text{NO}_x}^{\text{tp}} \end{bmatrix}, \quad (20)$$

where  $w_{\text{Hd}}$  and  $w_{\text{NO}_x}^{\text{tp}}$  represent equivalence factors for the enthalpy provided to the ATS and the tailpipe NO<sub>x</sub> emissions, respectively. Their interpretation is analogous to that of the equivalence factor used to penalize the use of the battery in the energy management of HEVs.

The values for  $w_{\text{Hd}}$  and  $w_{\text{NO}_x}^{\text{tp}}$  are calculated from the costates of the MPC optimization as

$$w_{\text{Hd}} = -\frac{\lambda_{\theta}}{C_{\text{ATS}}} \quad (21)$$

$$w_{\text{NO}_x}^{\text{tp}} = \lambda_{\text{NO}_x}^{\text{tp}} \cdot \quad (22)$$

The corresponding derivation is shown in Appendix B. The equivalence factor for tailpipe NO<sub>x</sub> emissions is then converted to an equivalence factor for engine-out NO<sub>x</sub> emissions using the current NO<sub>x</sub>-reduction efficiency  $\eta_{\text{DeNO}_x}$ , resulting in

$$w_{\text{NO}_x}^{\text{eo}} = (1 - \eta_{\text{DeNO}_x}) w_{\text{NO}_x}^{\text{tp}} \cdot \quad (23)$$

Finally, the low-level engine controller determines the engine control inputs by minimizing the objective

$$L = \dot{m}_{\text{fuel}}^*(\mathbf{u}) + w_{\text{NO}_x}^{\text{eo}} \dot{m}_{\text{NO}_x}^{\text{eo}}(\mathbf{u}) - w_{\text{Hd}} \dot{H}_{\text{ATS}}(\mathbf{u}), \quad (24)$$

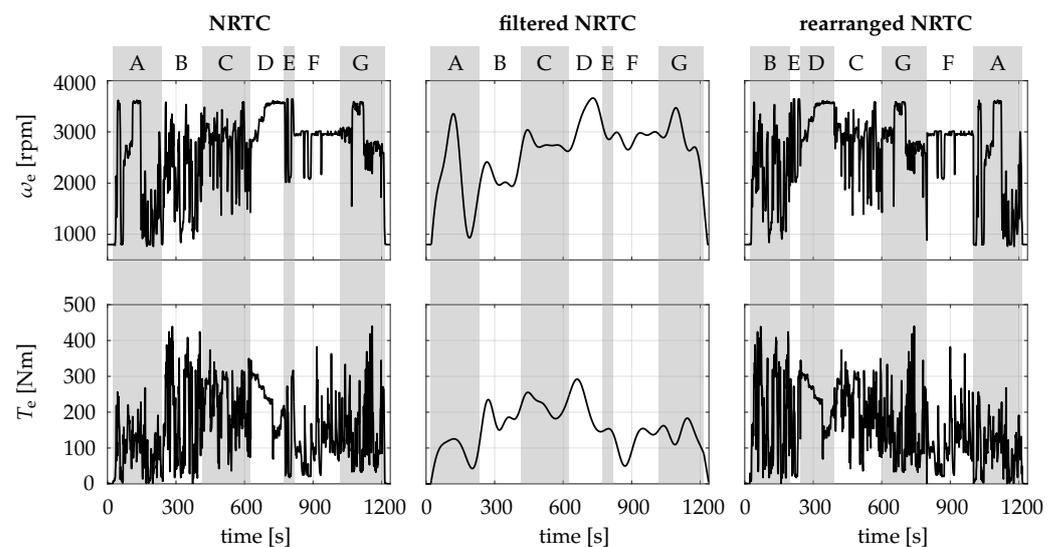
while considering all limits on the engine operation. The objective  $L$  represents a weighted sum of multiple objectives, where the equivalence factor  $w_{\text{NO}_x}^{\text{eo}}$  is used to penalize engine-out  $\text{NO}_x$  emissions and the equivalence factor  $w_{\text{Hd}}$  is used to encourage ATS heating. The optimal engine control inputs are determined using a grid search.

With this realization, the optimal engine control inputs are determined for the actual mission rather than for the predicted one, ensuring robustness toward mispredictions, at the cost of an increased computational demand.

## 6.2. Robustness to Misprediction

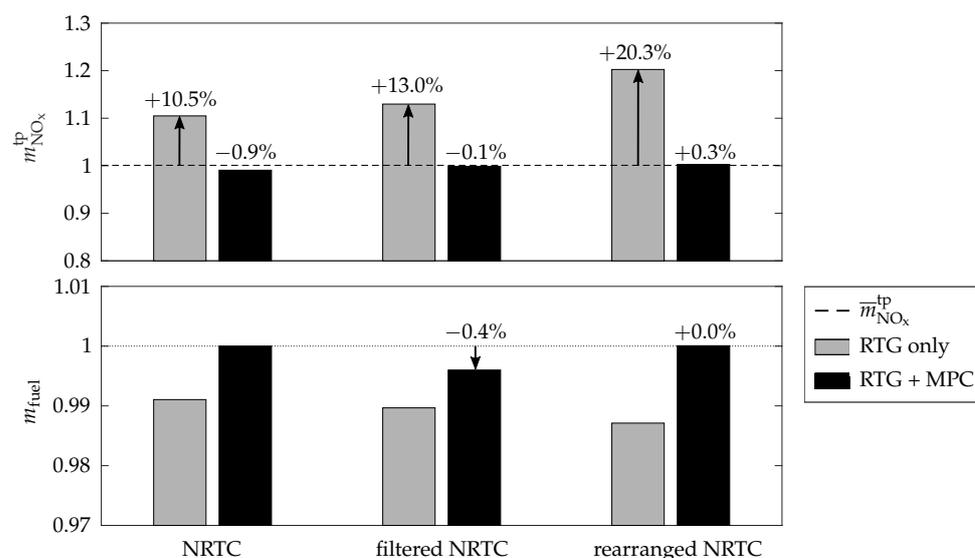
In this section, the performance of the predictive supervisory controller is analyzed for two misprediction cases. To obtain comparable results, the ATS is preconditioned as described in Section 5.

Figure 11 shows the prediction cycles considered in this section. The first misprediction case is shown in the middle. It was obtained by filtering the NRTC using a low-pass filter with a cut-off frequency of 0.01 Hz and represents the case where only averaged information about the upcoming missions is available. The second misprediction case is shown on the right. It was obtained by rearranging the individual segments of the NRTC and represents the case where the timing of the predictive information is wrong. The cycles with misprediction were used by the predictive supervisory controller, while the NRTC was followed during all experiments.



**Figure 11.** Considered misprediction cases. The NRTC is shown on the left. It consists of seven segments that describe the operation of different nonroad applications [27]. These segments are visualized by the gray areas and assigned a capital letter to illustrate their order. The filtered NRTC is shown in the middle. The rearranged NRTC is shown on the right.

The corresponding results are shown in Figure 12. The upper plot shows the error in the emitted tailpipe  $\text{NO}_x$  mass. The effect of misprediction can be analyzed by first considering the operation without feedback from the MPC. For the case without misprediction (left), the error of 10.5% stems solely from model mismatch. For the cycles with misprediction, the error increased to 13.0% for the filtered NRTC (middle) and to 20.3% for the rearranged NRTC (right).



**Figure 12.** Comparison of the measured operation for different misprediction cases. The results obtained without feedback from the MPC are shown in gray, while the results obtained with the full predictive supervisory controller are shown in black. The emitted tailpipe  $NO_x$  mass is normalized with respect to the limit, while the fuel mass is normalized with respect to the operation with the full predictive supervisory controller for the case without misprediction.

When feedback was introduced via the MPC, the predictive supervisory controller never exceeded the tailpipe  $NO_x$  limit by more than 0.3% for all considered cases. This result indicates that the predictive supervisory controller is robust towards mispredictions. Furthermore, the increased feedback control action did not result in an increase in fuel consumption.

In fact, the fuel consumption for the filtered NRTC is lower than that for the case without misprediction. A possible explanation for this is that the misprediction counteracts the model mismatch, resulting in a reference trajectory that results in better performance on the real system. This suggests that using a filtered prediction could improve the performance of the predictive supervisory controller.

## 7. Conclusions

In this work, a predictive supervisory controller for the integrated emissions management of a diesel engine and its aftertreatment system was developed. First, an efficient optimization method was developed to solve the optimal control problem of operating the engine system over a mission. The considered objective was the minimization of the fuel consumption, while respecting an upper bound on the emitted tailpipe  $NO_x$  mass. Next, the resulting optimization method was used in a predictive supervisory controller and the controller performance was evaluated experimentally.

The results show that the predictive supervisory controller successfully adapts the engine operation to the specific mission. At equivalent tailpipe  $NO_x$  emissions, the resulting controller was shown to reduce the fuel consumption by 1.1% for the nonroad transient cycle, compared to the operation with a fixed engine calibration. Further, the controller's robustness to different missions, initial ATS temperatures, and  $NO_x$  limits was demonstrated.

Finally, an adaptation of the low-level engine controller was presented that allows the controller to deal with mispredictions. The results for two misprediction cases showed that the predictive supervisory controller never exceeded the tailpipe  $NO_x$  limit by more than 0.3% and did not require more fuel than for the case without misprediction.

Based on this work, there are three avenues for future research. First, the developed methodology could be investigated further. For example, the type of predictive information

required by the predictive supervisory controller could be analyzed and further experiments with filtered predictive information could be performed. Further, improvements to the model and the optimization method could be considered to lower the computation time or achieve a more accurate result.

Second, the presented methodology could be used for the development of a predictive supervisory controller for series production vehicles. The focus should lie on the integration of the presented optimization method into embedded systems and on obtaining predictive information.

Third, the methodology could be used to obtain an experimental benchmark for the development of simpler, potentially non-predictive supervisory controllers. During the development of such controllers, their performance can be evaluated by comparing it to that of the method developed in this work.

**Author Contributions:** Conceptualization, J.R., C.P., O.C., M.G. and C.O.; methodology, J.R., C.P., O.C. and M.G.; software, J.R., C.P., O.C. and M.G.; investigation, J.R. and C.P.; writing—original draft preparation, J.R.; writing—review and editing, C.P., O.C., M.G. and C.O.; visualization, J.R.; supervision, C.O. All authors have read and agreed to the published version of the manuscript.

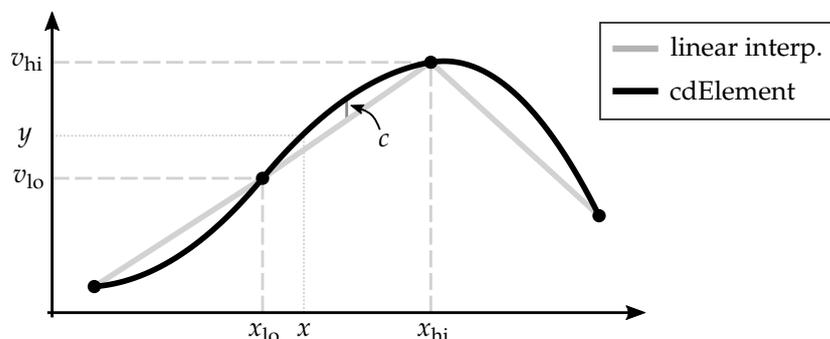
**Funding:** This research was funded by the Swiss Federal Office of Energy (SFOE, grant numbers SI/501742 and SI/501643) and the Swiss Competence Center for Energy Research (SCCER Mobility, contract number 11550025491).

**Acknowledgments:** We thank FPT Motorenforschung AG, the Swiss Federal Office of Energy (SFOE), and the Swiss Competence Center for Energy Research (SCCER) for supporting this project.

**Conflicts of Interest:** The authors declare no conflict of interest.

**Appendix A. Implementation of cdElements**

The basic principle of the continuously differentiable look-up elements (cdElements) is illustrated using the example shown in Figure A1.



**Figure A1.** Illustration of a one-dimensional continuously differentiable look-up element.

The cdElement includes a center  $c$  between any two vertices  $v_{lo}$  and  $v_{hi}$  and a quadratic function is evaluated as

$$\zeta = \frac{x - x_{lo}}{x_{hi} - x_{lo}} \tag{A1}$$

$$y = (1 - \zeta) v_{lo} + \zeta v_{hi} + 4 \zeta (1 - \zeta) c = v_{lo} + (v_{hi} - v_{lo} + 4c) \zeta - 4c \zeta^2 . \tag{A2}$$

where  $\zeta$  is the relative position of  $x$  with respect to  $v_{lo}$  and  $v_{hi}$ , and the center  $c$  is the distance between the quadratic function and the line segment connecting the vertices  $x_{lo}$  and  $x_{hi}$  at  $\zeta = 0.5$ .

During the fitting process of cdElements, continuous differentiability is enforced using respective constraints at each vertex. The example in Figure A1 features four grid points and four vertices. Additionally, the cdElement features three centers: one for each of the three segments. To enforce continuous differentiability, two constraints are considered

during the fitting process. The first constraint enforces that the gradient of the left segment and the gradient of the middle segment must be equal at  $x_{l0}$ . The second constraint enforces that the gradient of the middle segment and the gradient of the right segment must be equal at  $x_{hi}$ . The gradient at the outer bounds of the outermost segments is free.

The resulting evaluation of the cdElement is shown by the solid black line in Figure A1 and is smooth. The first and second derivative of cdElements are computed as

$$\frac{\partial y}{\partial x} = \frac{\partial y}{\partial \zeta} \frac{\partial \zeta}{\partial x} = \frac{1}{x_{hi} - x_{l0}} (v_{hi} - v_{l0} + 4c - 8c\zeta) \quad (A3)$$

$$\frac{\partial^2 y}{\partial x^2} = \frac{\partial}{\partial \zeta} \left( \frac{\partial y}{\partial x} \right) \frac{\partial \zeta}{\partial x} = \frac{-8c}{(x_{hi} - x_{l0})^2}. \quad (A4)$$

The evaluation of the derivatives is computationally inexpensive. The quadratic function between two grid points results in a piecewise affine first derivative and a piecewise constant second derivative. As a result, the cdElements are once continuously differentiable.

For the sake of brevity, the working principle of cdElements is only shown for a one-dimensional look-up element. Nevertheless, the method is extendable to higher-dimensional look-up elements. For a two-dimensional cdElement, the line segment between two vertices is replaced by a surface spanned by four vertices. Four edge centers and one face center are added. Continuous differentiability must be enforced along all four edges. In three-dimensional cdElements, cubes spanned by eight vertices must be considered. Twelve edge centers, six face centers, and one cube center must be added. Continuous differentiability must be enforced along all six faces.

While the method is applicable to multi-dimensional look-up elements, the scalability with dimensionality limits its application. In this work, one-, two-, and three-dimensional cdElements were used.

## Appendix B. Derivation of Equivalence Factors

The Hamiltonian of OCP (14a–k) and the objective of the low-level engine controller presented in Section 6.1 rewritten to consider tailpipe  $\text{NO}_x$  emissions are

$$H_{\text{MPC}} = \dot{m}_{\text{fuel}}^* + \lambda_{\text{NO}_x}^{\text{tp}} \frac{d}{dt} m_{\text{NO}_x}^{\text{tp}} + \lambda_{\vartheta} \frac{d}{dt} \vartheta \quad (A5)$$

$$L = \dot{m}_{\text{fuel}}^* + w_{\text{NO}_x}^{\text{tp}} \dot{m}_{\text{NO}_x}^{\text{tp}} - w_{\text{Hd}} \dot{H}_{\text{ATS}}^* \quad (A6)$$

The equivalence factors  $w_{\text{NO}_x}^{\text{tp}}$  and  $w_{\text{Hd}}$  are found by ensuring that  $L$  represents the same cost function as  $H_{\text{MPC}}$ . The first term is already equivalent. Comparing the second term and considering that  $\frac{d}{dt} m_{\text{NO}_x}^{\text{tp}} = \dot{m}_{\text{NO}_x}^{\text{tp}}$ , this results in

$$w_{\text{NO}_x}^{\text{tp}} = \lambda_{\text{NO}_x}^{\text{tp}}. \quad (A7)$$

Comparing the third term yields

$$-w_{\text{Hd}} \dot{H}_{\text{ATS}}^* = \lambda_{\vartheta} \frac{d}{dt} \vartheta = \frac{\lambda_{\vartheta}}{C_{\text{ATS}}} \left( \dot{H}_{\text{ATS}}^* - \alpha S (\vartheta - \vartheta_{\text{amb}}) \right) \quad (A8)$$

$$w_{\text{Hd}} = -\frac{\lambda_{\vartheta}}{C_{\text{ATS}}} \left( 1 - \frac{\alpha S}{\dot{H}_{\text{ATS}}^*} (\vartheta - \vartheta_{\text{amb}}) \right), \quad (A9)$$

which is approximated by neglecting the convective heat loss to the environment, resulting in

$$w_{\text{Hd}} \approx -\frac{\lambda_{\vartheta}}{C_{\text{ATS}}}. \quad (A10)$$

## References

1. Emission Standards. Available online: <https://dieselnet.com/standards/> (accessed on 24 February 2022).
2. Guzzella, L.; Onder, C. *Introduction to Modeling and Control of Internal Combustion Engine Systems*; Springer Science & Business Media: Berlin/Heidelberg, Germany, 2009.
3. Nova, I.; Tronconi, E. *Urea-SCR Technology for deNOx after Treatment of Diesel Exhausts*; Springer: New York, NY, USA, 2014.
4. Cloudt, R.; Willems, F. Integrated emission management strategy for cost-optimal engine-aftertreatment operation. *SAE Int. J. Engines* **2011**, *4*, 1784–1797. [[CrossRef](#)]
5. Ritzmann, J.; Chinellato, O.; Hutter, R.; Onder, C. Optimal integrated emission management through variable engine calibration. *Energies* **2021**, *14*, 7606. [[CrossRef](#)]
6. Zentner, S.; Aspiron, J.; Onder, C.; Guzzella, L. An equivalent emission minimization strategy for causal optimal control of diesel engines. *Energies* **2014**, *7*, 1230–1250. [[CrossRef](#)]
7. Elbert, P.; Amstutz, A.; Onder, C. Adaptive control for the real driving emissions of diesel engines. *MTZ Worldw.* **2017**, *78*, 68–74. [[CrossRef](#)]
8. Guzzella, L.; Sciarretta, A. *Vehicle Propulsion Systems*; Springer: Berlin/Heidelberg, Germany, 2007; Volume 1.
9. van Dooren, S.; Amstutz, A.; Onder, C.H. A causal supervisory control strategy for optimal control of a heavy-duty diesel engine with SCR aftertreatment. *Control Eng. Pract.* **2022**, *119*, 104982. [[CrossRef](#)]
10. Chen, P.; Wang, J. Nonlinear model predictive control of integrated diesel engine and selective catalytic reduction system for simultaneous fuel economy improvement and emissions reduction. *J. Dyn. Syst. Meas. Control* **2015**, *137*, 081008. [[CrossRef](#)]
11. Mentink, P.; Willems, F.; Kupper, F.; den Eijnden, E.V. *Experimental Demonstration of a Model-Based Control Design and Calibration Method for Cost Optimal Euro-VI Engine-Aftertreatment Operation*; SAE Technical Paper Series; Society of Automotive Engineers: Warrendale, PA, USA, 2013.
12. Ramachandran, S.N.; Hommen, G.; Mentink, P.; Seykens, X.; Willems, F.; Kupper, F. Robust, cost-optimal and compliant engine and aftertreatment operation using air-path control and tailpipe emission feedback. *SAE Int. J. Engines* **2016**, *9*, 1662–1673. [[CrossRef](#)]
13. Maamria, D.; Chaplais, F.; Petit, N.; Sciarretta, A. Numerical optimal control as a method to evaluate the benefit of thermal management in hybrid electric vehicles. *IFAC Proc. Vol.* **2014**, *47*, 4807–4812. [[CrossRef](#)]
14. Maamria, D.; Sciarretta, A.; Chaplais, F.; Petit, N. Online energy management system (EMS) including engine and catalyst temperatures for a parallel HEV. *IFAC-PapersOnLine* **2017**, *50*, 8913–8920. [[CrossRef](#)]
15. van Schijndel, J.; Donkers, M.C.F.; Willems, F.P.T.; Heemels, W. Dynamic programming for integrated emission management in diesel engines. *IFAC Proc. Vol.* **2014**, *47*, 11860–11865. [[CrossRef](#)]
16. Simon, A.; Nelson-Gruel, D.; Charlet, A.; Jaïne, T.; Nouillant, C.; Chamailard, Y. Optimal supervisory control of a diesel HEV taking into account both DOC and SCR efficiencies. *IFAC-PapersOnLine* **2018**, *51*, 323–328. [[CrossRef](#)]
17. Ritzmann, J.; Christon, A.; Salazar, M.; Onder, C. *Fuel-Optimal Power Split and Gear Selection Strategies for a Hybrid Electric Vehicle*; SAE Technical Paper; Society of Automotive Engineers: Warrendale, PA, USA, 2019.
18. Ritzmann, J.; Lins, G.; Onder, C. Optimization method for the energy and emissions management of a hybrid electric vehicle with an exhaust aftertreatment system. *IFAC-PapersOnLine* **2020**, *53*, 13797–13804. [[CrossRef](#)]
19. Aspiron, J. *Optimal Control of Diesel Engines*. Ph.D. Thesis, ETHZ, Zürich, Switzerland, 2013.
20. Kahaner, D.; Moler, C.; Nash, S. *Numerical Methods and Software*; Prentice-Hall, Inc.: Hoboken, NJ, USA, 1989.
21. Albin Rajasingham, T. *Nonlinear Model Predictive Control of Combustion Engines*; Springer: Cham, Switzerland, 2021.
22. Wächter, A.W. *Interior Point Methods for Large-Scale Nonlinear Programming with Applications in Process Systems Engineering*; Carnegie Mellon University: Pittsburgh, PA, USA, 2002.
23. Sundström, O.; Guzzella, L. A generic dynamic programming Matlab function. In Proceedings of the Control Applications, (CCA) & Intelligent Control, (ISIC), St. Petersburg, Russia, 8–10 July 2009; pp. 1625–1630.
24. Maciejowski, J.M. *Predictive Control: With Constraints*; Pearson Education: London, UK, 2002.
25. Salazar, M.; Balerna, C.; Elbert, P.; Grando, F.P.; Onder, C.H. Real-time control algorithms for a hybrid electric race car using a two-level model predictive control scheme. *IEEE Trans. Veh. Technol.* **2017**, *66*, 10911–10922. [[CrossRef](#)]
26. Rawlings, J.B. Tutorial overview of model predictive control. *IEEE Control Syst. Mag.* **2000**, *20*, 38–52.
27. Ettl, J.; Bernhardt, H.; Thuneke, K.; Emberger, P.; Remmele, E. Transient emission and fuel consumption measurements on plant oil tractors. *Landtechnik* **2016**, *71*, 44–54.