Proceeding of the 2004 American Control Conference
Boston, Massachusetts June 30 - July 2, 2004

Estimating the Oxygen Storage Level of a Three-Way Automotive Catalyst

Kenneth R. Muske
Department of Chemical Engineering, Villanova University, Villanova, PA 19085
James C. Peyton Jones
Department of Electrical and Computer Engineering, Villanova University, Villanova, PA 19085

Abstract

A moving horizon estimation strategy is developed for the determination of the oxygen storage level of a three-way automotive catalyst. The ability to store and release oxygen is a critical aspect to the operation of the three-way catalyst component in automotive emission control systems. Because the oxygen storage level cannot be measured directly, it must be estimated from pre-catalyst and post-catalyst air fuel ratio sensor measurements for model-based control and monitoring applications. However, these sensors are subject to distortion and the engine operating conditions that produce low tail-pipe emissions also result in the oxygen storage level becoming unobservable for a significant fraction of typical engine operation. The moving horizon estimator presented in this work, based on the use of a combined catalyst oxygen storage/sensor distortion model, is able to provide reliable estimates under these circumstances. The estimation algorithm is demonstrated using experimental engine operating data.

1. Introduction

An essential component in the reduction of automotive tail-pipe emissions is the use of three-way automotive catalysts in exhaust after-treatment systems. A three-way automotive catalyst reduces engine emissions by oxidizing both the unburned hydrocarbons and carbon monoxide and by reducing the nitrogen oxides contained in the pre-catalyst exhaust feedgas. Key to the operation of three-way catalyst systems is the ability to store and release oxygen resulting from chemisorption/desorption with the cerium oxides contained in the catalyst.

Although improvements in catalyst formulation and substrate design have made a significant contribution to reducing automotive emissions, there is considerable potential for further reduction from advanced control of the catalyst operation. As further environmental legislation imposes increasingly stringent tail-pipe emission regulations, realizing the potential of advanced catalyst control will become more important for regulatory compliance. Model-based techniques offer an attractive advanced control methodology for automotive catalyst systems, but a critical aspect of a model-based approach is the ability to estimate the current oxygen storage state of the catalyst model.

Because the catalyst oxygen storage level cannot be measured directly, it must be estimated based on the available pre- and post-catalyst exhaust gas oxygen (EGO) sensor measurements. A difficulty with these measurements, however, is that they do not provide sufficient information to estimate the oxygen storage state during a significant fraction of the catalyst operation. The post-catalyst sensor can also be subject to distortion that arises from the effects of reversible catalyst deactivation during rich air fuel ratio engine operation. An effective oxygen storage state estimation scheme must therefore provide accurate dynamic estimates while the state is unobservable and reliable model updates when the state becomes observable. In this work, we present a moving horizon estimation strategy for the catalyst oxygen storage level. This estimator is applied using a simplified dynamic catalyst model that captures the underlying physical behavior of the oxygen chemisorption and the reversible deactivation effect on post-catalyst sensor distortion. The computational limitations imposed by implementation on an automotive engine control computer are also considered in the estimator design.

2. Catalyst Operation

Under rich (excess fuel) engine operation, the catalyst oxidizes the hydrocarbons and carbon monoxide present in the incoming feedgas by releasing previously stored oxygen. This oxygen release maintains stoichiometric combustion with commensurately low levels of hydrocarbon and carbon monoxide emissions. Because of the finite storage capacity of the catalyst, however, this process cannot continue indefinitely. When the oxygen release rate of the depleted catalyst can no longer satisfy the feedgas demand, the

0-7803-8335-4/04/$17.00 ©2004 AACC

post-catalyst air fuel ratio will decrease below stoichiometric and hydrocarbon breakthrough will eventually occur. A typical catalyst control system will therefore attempt to switch to lean (excess air) engine operation before this breakthrough condition is encountered. Under lean engine operation, the excess oxygen in the feedgas is now adsorbed onto the catalyst resulting in near-stoichiometric post-catalyst conditions and low tail-pipe emissions. As the oxygen storage capacity of the catalyst approaches its saturation condition, however, the post-catalyst oxygen concentration increases above stoichiometric and breakthrough of nitrogen oxides will eventually occur. A typical catalyst control system will then attempt to switch back to rich engine operation before lean breakthrough. By cycling the engine operation in this way, the oxygen storage capacity of the catalyst can be used as a buffer against breakthrough by compensating for transient oxygen excess or deficiency.

This dynamic behavior is clearly shown in Figure 1 where the post-catalyst air fuel ratio remains essentially at stoichiometric for several seconds after the pre-catalyst air fuel ratio makes a lean to rich step transition. When the oxygen release rate can no longer satisfy the feedgas demand, the post-catalyst air fuel ratio begins to become rich until it eventually matches the pre-catalyst air fuel ratio. The separation between the ratios noted between twenty and seventy seconds is due to sensor distortion in the post-catalyst air fuel ratio sensor. Oxygen storage on the catalyst is clearly shown after the pre-catalyst air fuel ratio rich-to-lean step change made at approximately seventy seconds. We note that the magnitude and duration of the pre-catalyst feedgas transitions are much larger than would be made in practice.

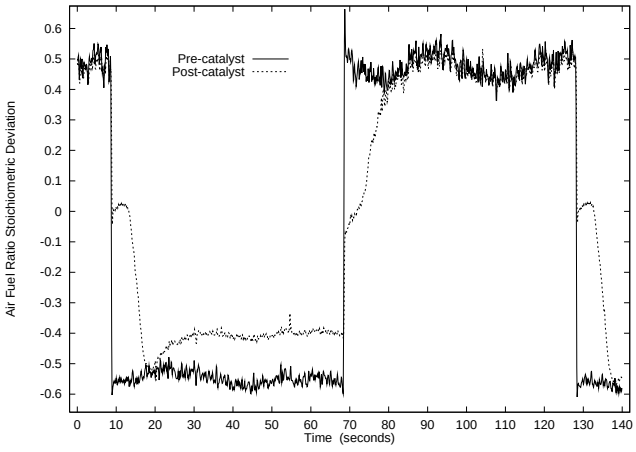

Figure 1: Catalyst oxygen storage level transitions.

3. Catalyst Model

The oxygen storage model used in the present study is described in [1]. The catalyst oxygen storage is represented by the following nonlinear integrating model

$$\dot{\phi} = K_\lambda \dot{m}_f \left(\Delta\lambda^{\triangleleft} - \mathcal{N}(\phi)\right) \tag{1}$$

in which ϕ is the oxygen storage state of the catalyst, $\Delta\lambda^{\triangleleft}$ is the pre-catalyst air fuel ratio deviation from stoichiometric, $\dot{m}_f$ is the fuel mass flow rate to the engine, and $\mathcal{N}(\phi)$ determines the oxygen chemisorption rate. Here oxygen storage and release rates depend on the difference between the forcing function $\Delta\lambda^{\triangleleft}$, which promotes adsorption under lean air fuel ratio conditions and desorption under rich air fuel ratio conditions, and the catalyst oxygen capacity $\mathcal{N}(\phi)$. This model describes stored oxygen relative to the equilibrium level when the pre-catalyst exhaust gas is stoichiometric, $\Delta\lambda^{\triangleleft} = 0$.

3.1. Catalyst Oxygen Capacity

In general, the catalyst oxygen capacity $\mathcal{N}(\phi)$ has a nonlinear spring characteristic. As ϕ increases from zero, it becomes progressively harder to store oxygen on the catalyst and as ϕ decreases from zero, it becomes progressively harder to remove oxygen from the catalyst. In this work, $\mathcal{N}(\phi)$ is parameterized using a quintic polynomial expansion of ϕ.

$$\mathcal{N}(\phi) = a_1\phi + a_2\phi^2 + a_3\phi^3 + a_4\phi^4 + a_5\phi^5 \tag{2}$$

We note that the stored oxygen level does not hit hard saturation/depletion limits as in limited integrator models, such as presented in [2] and [3], but instead attains an operating condition which is dependent on the steady-state pre-catalyst air fuel ratio. The oxygen saturation/depletion limits are approached asymptotically as expected from theory.

The post-catalyst, or final tail-pipe, air fuel ratio deviation from stoichiometric $\Delta\lambda^{\triangleright}$ is determined from the oxygen capacity function by the relationship

$$\Delta\lambda^{\triangleright} = \Psi(\phi, \Delta\lambda^{\triangleleft}) \tag{3}$$

$$\Psi(\phi, \Delta\lambda^{\triangleleft}) = \begin{cases} 0, & (\Delta\lambda^{\triangleleft} < 0) \text{ and } (\phi > 0) \\ \mathcal{N}(\phi), & \text{otherwise} \end{cases}$$

where the pre-catalyst air fuel ratio, $\Delta\lambda^{\triangleleft}$, is denoted using a left-facing triangle and the post-catalyst air fuel ratio is denoted using a right-facing triangle. In the case of rich pre-catalyst exhaust gas, $\Delta\lambda^{\triangleleft} < 0$, and an oxidized catalyst, $\phi > 0$, where stored oxygen is available for reducing the rich feedgas, the system is demand limited because the oxygen release is limited by the feedgas oxygen demand resulting in stoichiometric post-catalyst exhaust, $\Delta\lambda^{\triangleright} = 0$.

3.2. Post-catalyst EGO Sensor Distortion

Post-catalyst EGO sensor distortion under rich air fuel ratio operating conditions can be related to the

hydrogen generated by catalytic promotion of the water gas shift reaction.

$$CO + H_2O \rightleftharpoons CO_2 + H_2$$

When the reaction proceeds in the forward direction, the ability of hydrogen to diffuse faster than the other gas components results in a sensor output that appears richer than the true value. In the same manner, reduced levels of post-catalyst hydrogen due to a progressive inhibition of the water gas shift reaction result in a sensor output that appears leaner that the true value. Because hydrogen has a strong effect on the post-catalyst EGO sensor output, any changes from the concentration used to calibrate the sensor can result in a gas-composition dependent error in the post-catalyst measurement [4]. The rise toward leaner values of the measured post-catalyst air fuel ratio stoichiometric deviation shown in Figure 1 is due to this effect. The pre-catalyst sensor dependency on composition can be largely removed by calibration with representative engine-out exhaust. Because the post-catalyst composition varies dynamically according to the reactions taking place, static calibration is insufficient to correct for these errors which can propagate to the estimation algorithm.

This behavior suggests that the sensor distortion can be used as a measure of the reversible catalyst deactivation effect [5]. If the degree of water gas shift reaction inhibition is assumed proportional to the deactivated fraction of the catalyst surface, ψ, then the apparent post-catalyst AFR, $\Delta\lambda_a^{\triangleright}$, can be related to the true post-catalyst AFR, $\Delta\lambda^{\triangleright}$, as follows

$$\Delta\lambda_a^{\triangleright} = \Delta\lambda^{\triangleright} + K_\psi \psi \tag{4}$$

where the constant K_ψ represents the combined effects of the sensor sensitivity to hydrogen concentration and the inhibition of the water gas shift reaction due to reversible catalyst deactivation.

Assuming the rate of deactivation is proportional to the post-catalyst oxygen deficiency, $-\Delta\lambda^{\triangleright}$, and the fraction of the surface that is already deactivated, ψ,

$$\dot{\psi}_d = \dot{m}_f K_d(-\Delta\lambda^{\triangleright} - \psi)$$

where K_d is the deactivation constant of proportionality. The presence of pre-catalyst free oxygen, $\Delta\lambda^{\triangleleft} > 0$, reverses the deactivation process at a rate proportional to the supply of pre-catalyst oxygen until the catalyst is reactivated, $\psi = 0$,

$$\dot{\psi}_r = -\dot{m}_f K_r \Delta\lambda^{\triangleleft}$$

where K_r represents the reactivation constant. The resulting reversible deactivation model becomes

$$\dot{\psi} = \begin{cases} \dot{\psi}_d, & (\Delta\lambda^{\triangleleft} < 0) \text{ and } (\Delta\lambda^{\triangleright} < 0) \\ \dot{\psi}_r, & (\Delta\lambda^{\triangleleft} > 0) \text{ and } (\psi > 0) \\ 0, & \text{otherwise} \end{cases} \tag{5}$$

where $\Delta\lambda^{\triangleright}$ is the true post-catalyst AFR determined from the oxygen storage model in Eq. 3 and the sensor distortion of the measured post-catalyst AFR can be determined from Eq. 4.

3.3. Combined Model

The two states of the combined catalyst model are represented by Eqs. 1 and 5. Figure 2 presents a comparison between the measured post-catalyst air fuel ratio and the model prediction for the a series of more representative step transitions to the pre-catalyst air fuel ratio. As shown in this figure, the deviations at the end of each rich cycle are accounted for in the combined model. In particular, the combined model now characterizes the distortion that occurs in the post-catalyst sensor under rich air fuel ratio engine operation. The close agreement between the predicted and actual measurements, especially during the periods of rich operation, demonstrates that the reversible catalyst deactivation dynamics are well described by this model.

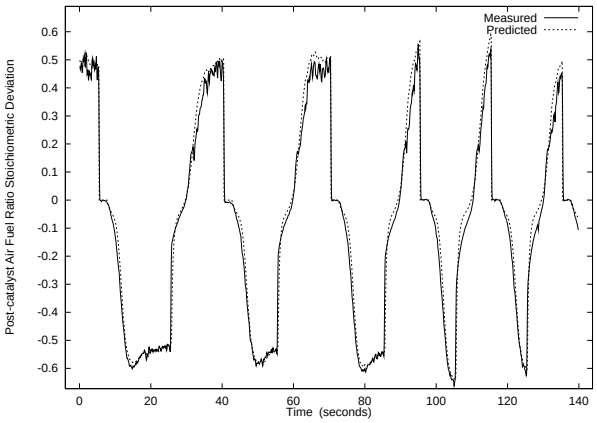

Figure 2: Post-catalyst air fuel ratio prediction.

4. Catalyst Control Strategy

From a control perspective, the post-catalyst air fuel ratio stoichiometric deviation should ideally remain at $\Delta\lambda^{\triangleright} = 0$ to maximize catalyst efficiency and minimize the vehicle tail-pipe emissions. When $\Delta\lambda^{\triangleright} < 0$, there is the possibility of hydrocarbon breakthrough. When $\Delta\lambda^{\triangleright} > 0$, there is the possibility of nitrogen oxide breakthrough. Therefore, the catalyst is typically operated such that the oxygen storage level is always either filling or emptying around the stoichiometric equilibrium storage level. A common practice to achieve this operation is to cycle the pre-catalyst air fuel ratio across stoichiometric at a frequency determined during engine calibration or obtained through relay feedback from a post-catalyst heated exhaust gas oxygen (HEGO) sensor. The output of such sensors switches sharply around stoichiometry resulting in a limit cycle behavior of the catalyst system [6].

A slightly more sophisticated model-based control strategy would aim to maintain the catalyst near the stoichiometric equilibrium state in order to maximize the time required to reach either rich or lean breakthrough conditions when subject to disturbances. The optimum value of ϕ depends on the relative oxidation verses reduction rates, the relative probability and magnitude of rich verses lean disturbances, and the penalties incurred following a rich verses lean post-catalyst breakthrough. A model-based control strategy of this type is presented in [7]. The strategy adopted in [8] allows ϕ to vary depending on engine demand while preventing post-catalyst breakthrough.

A problem with this control strategy for a model-based algorithm is that the desired engine operation results in significant periods of stoichiometric post-catalyst air fuel ratio. The oxygen storage state is then unobservable since changes in the feedgas have, at least to a first approximation, no effect on the post-catalyst air fuel ratio. Examination of Eq. 3 reveals that the catalyst state for the model presented here cannot be reconstructed from the pre- and post-catalyst air fuel ratio sensors when $\Delta\lambda^{\triangleleft} < 0$ and $\phi > 0$. Therefore, a model-based control strategy would essentially run open-loop at these conditions.

5. Oxygen Storage State Estimation

There are significant demands placed on oxygen storage state estimation in model-based catalyst control applications. The current oxygen storage state can only be estimated during a fraction of the lean switching cycle or during disturbances when the post-catalyst air fuel ratio moves off stoichiometry. The estimate must then be based on rather noisy air fuel ratio measurements. Although it is possible to reduce the switching frequency to provide additional post-catalyst sensor information, this practice would result in increased tail-pipe emissions and could not be applied in a regular fashion. There are also restrictions on the computation time allowed for implementation of the estimator because of the system dynamics and the restrictions inherent in an automotive engine control computer. In this work, a moving horizon least squares estimator is presented that addresses these issues.

The single estimated parameter ω_k is obtained at time t_k from the solution to the following nonlinear least squares problem

$$\begin{aligned} \min_{\omega_k} \int_{t_k - t_h}^{t_k} (\Delta\lambda_a^{\triangleright} - \Delta\lambda_m^{\triangleright})^2 \; + \; \alpha \; \omega_k^2 \\ \Delta\lambda_a^{\triangleright} \;\; = \;\; \Delta\lambda^{\triangleright} + K_\psi \psi \\ \text{Subject to:} \quad \Delta\lambda^{\triangleright} \;\; = \;\; \Psi(\phi, \Delta\lambda^{\triangleleft}) \\ \omega_{\min} \leq \;\; \omega_k \;\; \leq \omega_{\max} \end{aligned} \tag{6}$$

in which t_k is the time the kth estimate, t_h is the horizon length, $\Delta\lambda_a^{\triangleright}(t)$ is the model predicted post-catalyst air fuel ratio stoichiometric deviation, $\Delta\lambda_m^{\triangleright}(t)$ is the measured post-catalyst air fuel ratio deviation, and α is the regularization term weighting. This estimator is analogous to that presented in [9] except that we do not necessarily restrict ω_k to be the state estimate at the beginning of the horizon.

The most obvious choice for ω_k is a bias to the estimated oxygen storage state at time $t_k - t_h$.

$$\Delta\hat{\phi}(t_k - t_h) = \omega_k \tag{7}$$

The estimation algorithm in Eq. 6 then becomes the initial state moving horizon estimator in [10]. A second approach is to bias one of the model inputs. For the present model, either a bias to the measured pre-catalyst air fuel ratio stoichiometric deviation $\Delta\lambda_m^{\triangleleft}$

$$\Delta\hat{\lambda}^{\triangleleft}(t) = \Delta\lambda_m^{\triangleleft}(t) + \omega_k, \quad t \in (t_k - t_h, t_k] \tag{8}$$

or a multiplicative bias to the space velocity $K_\lambda \dot{m}_f$

$$K_\lambda \hat{m}_f(t) = (K_\lambda \dot{m}_f)\omega_k, \quad t \in (t_k - t_h, t_k] \tag{9}$$

can be applied. The choice of bias depends on how well it captures the dynamic nature of the true disturbances and how sensitive the model output is to the bias. Because a single parameter is estimated by this algorithm, the determination of the sequence $\{\omega_k\}$ can be carried out as a series of one-dimensional optimization or search problems that are not computationally demanding. Provided a feasible sequence is generated by the optimization procedure [11], an approximate solution is available if the computational time limit is reached before convergence.

This estimator is well motivated if the unmodeled disturbances are negligible and the post-catalyst measurement is corrupted by zero mean noise. If the system is unobservable over the estimation horizon t_h, the constraints and regularization penalty will prevent the estimation algorithm from introducing a large correction that will drive the catalyst model out of demand limited operation. This behavior would not be the case for a post-catalyst air fuel ratio stoichiometric deviation bias correction. For this reason, and the observer stability problems associated with output bias disturbance models for integrating systems, we do not consider an output bias in this work.

6. Example

We test the moving horizon estimator on the engine operating data presented in Figure 3. The pre-catalyst air fuel ratio, shown in the top half of the figure, is operated approximately at stoichiometry for the first eight seconds and then operated rich for the next twenty-five seconds in order to precondition the catalyst to an oxygen depleted state. After this point, it is cycled between lean and rich operation at a frequency of 1 Hz. As shown in the bottom half of the

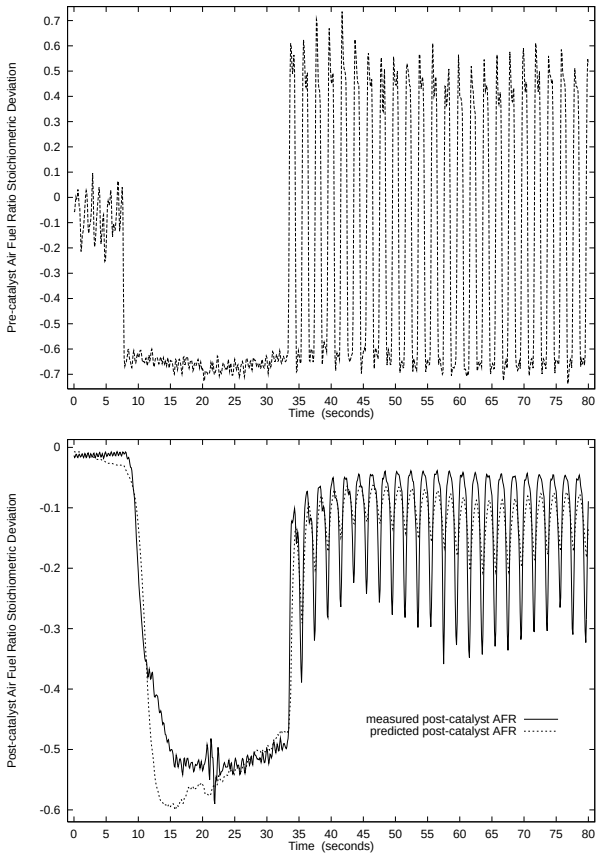

Figure 3: Example engine operation.

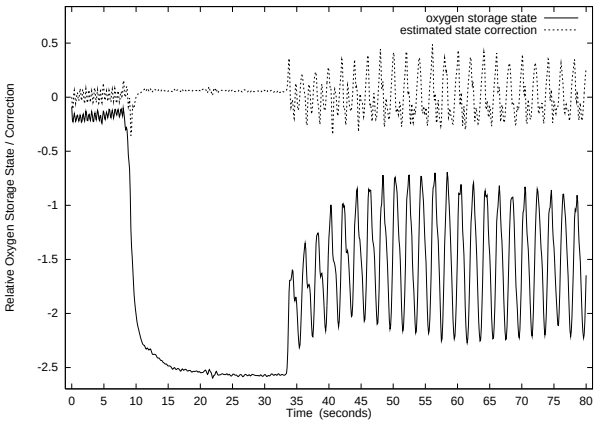

Figure 4: Estimated oxygen storage state and correction.

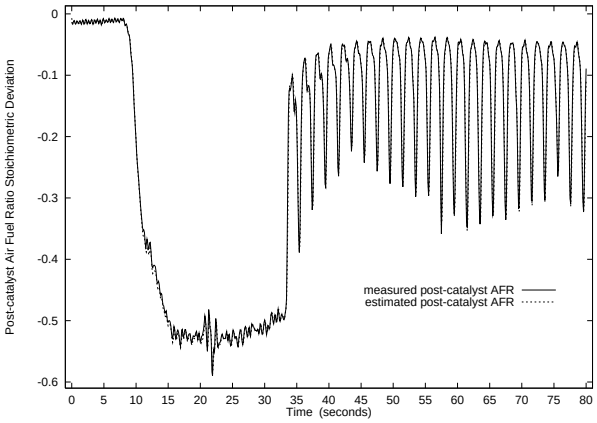

Figure 5: Post-catalyst estimated air fuel ratio.

figure, the post-catalyst air fuel ratio goes rich after the pre-catalyst step transition and then cycles in response to the pre-catalyst air fuel ratio. The model predicted response, using the same model parameters that produced the response in Figure 2, follows the initial transient and cycles at the same frequency as the measured value, but under-predicts the magnitude of the subsequent oscillations.

When the catalyst oxygen storage level is near depletion or saturation, the magnitude of the gain relating a change in the output $\Delta\lambda^{\triangleright}$ to a change in the catalyst state ϕ is large and the model prediction is generally good as shown in Figure 2. When the oxygen storage level is closer to its stoichiometric equilibrium level, as in the high frequency portion of the engine operating data in this example, the magnitude of this gain approaches zero and the output is much less sensitive to the catalyst state. Under these conditions, the effects of noise and unmodeled dynamics, when integrated by the model, can result in larger prediction errors. On-line estimation of the stored oxygen level is therefore useful to correct for these effects, especially since stored oxygen is often used as a control objective.

We first consider the initial state estimation scheme presented in Eq. 7 with a horizon of $t_h = 1$ sec. The estimated oxygen storage state $\hat{\phi}$ and the state correction $\Delta\hat{\phi}$ for the engine operation in Figure 3 is presented in Figure 4. As shown in this figure, the state corrections are essentially zero for the low frequency engine operation. During the high frequency operation, larger state corrections are necessary to accurately predict the magnitude of the post-catalyst air fuel ratio oscillations. The estimated post-catalyst air fuel ratio, shown in Figure 5, is essentially the same as the measured values when these state corrections are applied.

We next consider the model input bias schemes presented in Eqs. 8 and 9 with a horizon of $t_h = 1$ sec. The pre-catalyst air fuel ratio bias in Eq. 8 for the engine operation in Figure 3 is presented in Figure 6. The space velocity multiplicative bias in Eq. 9 for the same engine operation is presented in Figure 7. As shown in these figures, the bias values are essentially zero for the low frequency engine operation and then become unreasonably large during the high frequency operation. In the case of the space velocity bias, the constraints $0.01 \leq \omega_k \leq 5$ are active for a significant fraction of the operation. For these in-

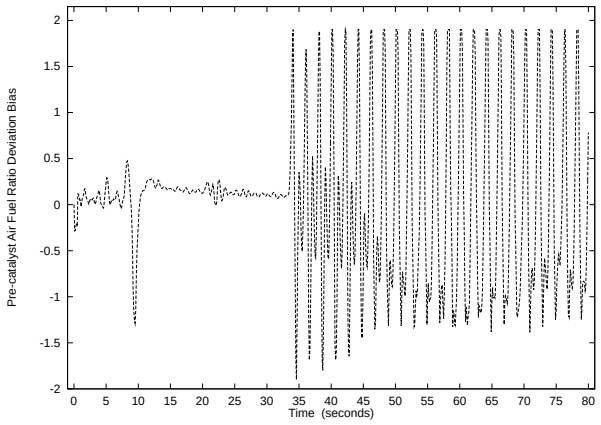

Figure 6: Estimated pre-catalyst AFR bias.

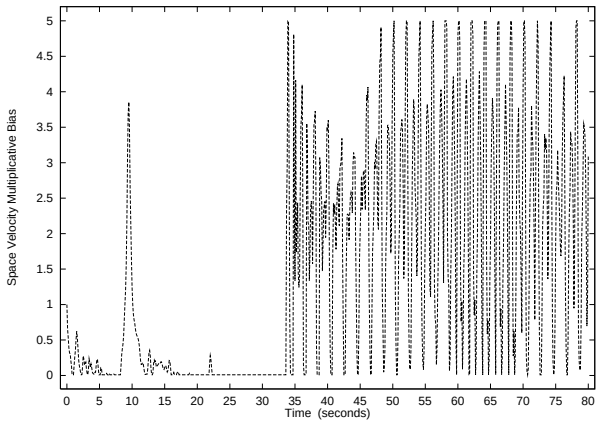

Figure 7: Estimated fuel mass flow rate bias.

put disturbances, the bias values required to minimize the prediction error to the level shown in Figure 5 result in large future oxygen storage state disturbances. These state disturbances must then be corrected by unreasonable bias estimates in the future. This behavior can be eliminated by increasing the regularization weight α, but with a significant increase in the post-catalyst air fuel ratio prediction error. These results suggest that input disturbances are not appropriate for the integrating oxygen storage model considered in this work.

7. Conclusions

We have presented an optimization-based estimation scheme for the oxygen storage state of a three-way automotive catalyst. The scheme is based on the minimization of the post-catalyst air fuel ratio squared prediction error over a past prediction horizon by the selection of a single estimated parameter. A single parameter is chosen in this application in order to obtain an algorithm that can reasonably be implemented in an automotive engine control computer. Based on the results of this study, the initial state correction approach provided the best performance with physically reasonable estimated parameter values. Future work includes on-line implementation of this algorithm and a combination estimation/parameter identification algorithm for on-line adaptation of the catalyst model parameters.

Acknowledgments

Support from the National Science Foundation under grant CTS-0215920 is gratefully acknowledged.

References

[1] J. Peyton Jones, R. Jackson, J. Roberts, and P. Bernard. A simplified model for the dynamics of a three-way catalytic converter. SAE Paper 2000-01-0652, SAE World Congress, 2000.

[2] E. Brant, Y. Wang, and J. Grizzle. Dynamic modeling of a three-way catalyst for SI engine exhaustemission control. *IEEE Trans. Contrl. Sys. Tech.*, 8:767–776, 2000.

[3] M. Balenovic, A. Backx, and J. Hoebink. On a model-based control of a three-way catalytic converter. SAE Paper 2001-01-0937, SAE World Congress, 2001.

[4] J. Peyton Jones and R. Jackson. Potential and pitfalls in the use of dual EGO sensors for three-way catalyst monitoring and control. *J. Automobile Engineering*, in press, 2003.

[5] J. Peyton Jones. Modeling combined oxygen storage and reversible deactivation dynamics for improved emissions predictions. SAE Paper 2003-01-0999, SAE World Congress, 2003.

[6] P. Eastwood. *Exhaust Gas Aftertreatment*. Research Studies Press, Baldock, England, 2000.

[7] M. Balenovic, T. Backx, and T. de Bie. Development of a model-based controller for a three-way catalytic converter. SAE Paper 2002-01-0475, SAE World Congress, 2002.

[8] K. Muske and J. Peyton Jones. Feedforward/Feedback air fuel ratio control for an automotive catalyst. In *Proceedings of the 2003 ACC*, pages 1386–1391, 2003.

[9] H. Michalska and D. Mayne. Moving horizon observers and observer-based control. *IEEE Trans. Auto. Contrl.*, 40:995–1006, 1995.

[10] K. Muske and T. Edgar. Nonlinear state estimation. In *Nonlinear Process Control*. Prentice–Hall, 1997.

[11] J. Nocedal and S. J. Wright. *Numerical Optimization*. Springer–Verlag, New York, 1999.