2005 American Control Conference
June 8-10, 2005. Portland, OR, USA

Model-based control methodologies for catalytic surface reactions

Ethan A. Mastny, James B. Rawlings, and Yannis G. Kevrekidis

Abstract— **Kinetic Monte Carlo (kMC) simulations can accurately portray the dynamics of both the microscopic and macroscopic behavior of reactions on surfaces. Unfortunately, optimizing an input trajectory for real-time control based on kMC models is still computationally intractable. This paper discusses how kMC models of CO oxidation can be used as the basis for the design of three different MPC controllers that are implementable in real-time. These models are: (i) a mean field model that closes in the species coverages, (ii) a lookup table of linear state-space space models generated from realizations of kMC simulations, (iii) a low-order model (checkerboard mean field) conceptualized from observation of kMC simulation microscopic states. Closed-loop simulations of a kMC plant with these three MPC controllers show that all the controllers can meet overall control objectives, but with varying levels of performance. The checkerboard mean field controller demonstrates the best performance.**

I. INTRODUCTION

Reactions that occur at the gas-solid interface involve a variety of molecular level events (adsorption, diffusion, reaction, surface reconstruction, and desorption). Given a description of how these molecular level events occur, rate equations can be derived based on states that are short-range correlations of the coverages (such as total or pair coverages). In general these closures do not contain sufficient structure to accurately model the dynamics of the surface coverages, due to the simplifying assumptions used to achieve closure [1]. Alternatively, molecular level information can be used directly in kinetic Monte Carlo (kMC) simulations. KMC simulations track the evolution of all the system states (position of every molecule on the surface), as sampled from the chemical master equation. From the microscopic evolution, an accurate macroscopic (total coverage) evolution can be determined, with significant computational expense.

Using kMC models to perform systems engineering tasks is a challenging and significant research area. Systems engineering includes tasks such as bifurcation analysis, parameter estimation, model-based control, or other tasks that involve optimization of macroscopic observables. KMC simulations track single realizations of the evolution of every adsorbed species and therefore require significant computational effort for each simulation. KMC simulation results are samples from a probability distribution. Therefore, optimization of kMC models is usually performed on the expectation of the macroscopic observables. Also, efficient performance of systems tasks requires the partial derivatives of the expectation of the macroscopic states with respect to adjustable parameters (sensitivities). In general, sensitivities are not explicitly available from kMC simulations. Given noisy output from kMC simulations, sensitivity evaluation requires large parameter perturbations and significant averaging. Given limited computational resources, modelers and engineers are left to balance accuracy of results against and the time available to complete the calculations.

Since there are many applications in which kMC models constitute the best available description of system evolution, there have been many efforts to perform systems engineering tasks using kMC models. The following is a sample of some recent studies in this area. Raimondeau et al. used a kMC model for parameter estimation of CO oxidation on a Pt catalyst, using a response surface method [2]. Gallivan et al. [3] applied a model reduction technique to the chemical master equation, to determine an optimal open-loop temperature profile for epitaxial thin film growth. In another thin film application, Lou et al. [4] used a kMC model to predict output dynamics in between actual system measurements. These predicted outputs were used by a multivariate PI controller to determine the system input. Additionally, "coarse timesteppers" or "equation-free" methods have been applied to a variety systems engineering applications, such as stability and bifurcation analysis [5], [6] and model identification for LQ and pole placement controllers of lattice-gas kMC models [7], [6].

Most of these applications involved off-line use of computationally intensive kMC models. Lou et al. [4] used a kMC model in an online control application, but did not use the kMC model to find an optimal trajectory of inputs, as models are used in Model Predictive Control (MPC). Due to the computational expense of evaluating kMC models, noisy outputs, and lack of sensitivities, direct application of MPC with kMC models is still out of reach.

The objective of this paper is to demonstrate methods for model-based control of catalytic surfaces when a kMC model is the best available representation of the system behavior. The kMC simulation used to represent the plant is a 100x100 lattice. This lattice size maintains low fluctuations of macroscopic observables (i.e. coverages) while maintaining reasonable simulation times.

CO oxidation with strong lateral CO-CO repulsions is the system used in this study. This paper describes how three deterministic models of CO oxidation were used as the basis for the design of three different MPC controllers and compares the controllers performance when implemented

E. A. Mastny is a PhD candidate and J. B. Rawlings is faculty with the Department of Chemical and Biological Engineering, University of Wisconsin-Madison, 1415 Engineering Drive, Madison, WI 53706, USA eamastny@wisc.edu, jbraw@bevo.che.wisc.edu
Y. G. Kevrekidis is faculty with the the Department of Chemical Engineering, Princeton University, Princeton, NJ 08544-5263, USA yannis@princeton.edu

0-7803-9098-9/05/$25.00 ©2005 AACC

TABLE I
REACTION MODEL PARAMETER VALUES

Parameter	Value		Parameter	Value
α	1.6 s^{-1}		β	Varies
γ	0.001 s^{-1}		k_r	1.0 s^{-1}
d	1.0 s^{-1}		ϵ	-2.0 kcal/mol
T	500 K			

on a kMC plant. The first model is a mean field (MF) model that closes in the surface coverage of each species. The second model is a scheduled linear state-space model located from a look-up table (generated off-line by kMC simulations) parameterized by the CO coverage set-point. The third model has four states and was inspired by visualization of kMC simulation microscopic states; we refer to it as the checkerboard mean field model (CMF).

II. CO OXIDATION

The system studied in this work is the surface oxidation of carbon monoxide, $2\mathrm{CO} + \mathrm{O_2} \longrightarrow 2\mathrm{CO_2}$. CO oxidation has become a model system for surface reaction studies because its simple mechanism exhibits many interesting features, such as: unstable steady states, surface patterns, and sustained oscillations.

In this study the reaction is assumed to proceed according to the following mechanism (where g refers to a gas phase species and $*, \mathrm{i}$ refers to an adsorbed species):

$$\mathrm{CO(g)} + *_\mathrm{i} \xrightarrow{\alpha} \mathrm{CO}_{*,\mathrm{i}} \quad (1)$$

$$\mathrm{O_2(g)} + *_\mathrm{i} + *_\mathrm{j} \xrightarrow{\beta} \mathrm{O}_{*,\mathrm{i}} + \mathrm{O}_{*,\mathrm{j}} \quad (2)$$

$$\mathrm{CO}_{*,\mathrm{i}} \xrightarrow{\gamma} \mathrm{CO(g)} + *_\mathrm{i} \quad (3)$$

$$\mathrm{CO}_{*,\mathrm{i}} + \mathrm{O}_{*,\mathrm{j}} \xrightarrow{k_r} \mathrm{CO_2(g)} + *_\mathrm{i} + *_\mathrm{j} \quad (4)$$

$$\mathrm{CO}_{*,\mathrm{i}} + *_\mathrm{j} \xrightarrow{d_r} *_\mathrm{i} + \mathrm{CO}_{*,\mathrm{j}} \quad (5)$$

$$\mathrm{O}_{*,\mathrm{i}} + *_\mathrm{j} \xrightarrow{d_r} *_\mathrm{i} + \mathrm{O}_{*,\mathrm{j}}. \quad (6)$$

Where i and j are adjacent sites. Additionally, adsorbed nearest neighbor CO molecules have repulsive interactions of energy ϵ. The system parameters are shown in Table I.

It is assumed that a 100x100 kMC simulation of the above mechanism accurately represents the dynamics of CO oxidation, therefore the kMC model is used as the plant.

III. MODEL PREDICTIVE CONTROL STRUCTURE

The objective of the MPC controllers in this study is to take the measurement of CO coverage (θ_{CO}) from a kMC simulation (the plant) to the set-point, by manipulating the gas concentration of $\mathrm{O_2}$ (β).

A. Controller Model

The details of the controller models are different but all the models can be represented by the structure shown in Eq. 7 - 9, where k is sample time, and x is the state vector.

$$x_{k+1} = f(x_k, u_k + X_u d_k) \quad (7)$$

$$d_{k+1} = d_k \quad (8)$$

$$y_k = g(x_k) + X_y d_k \quad (9)$$

The input (u) is proportional to the gas phase concentration of $\mathrm{O_2}$, and the measurements (y) are the surface coverages of CO and O. The disturbance model (d) is used to match plant dynamics with model predictions. The conditions represented by the states (x) depend upon the particular model being implemented.

B. Target Calculation

Given a θ_{CO} set-point, and disturbance value, the above model is used to determine the state and input targets (x_t, u_t) that correspond to a steady state in the model states. In these calculations, the target (x_t, u_t) is assumed constant over the horizon of N time steps.

C. Regulator

Given an estimate of the current state (x_0) the MPC regulator solves the following optimization problem to find the optimal trajectory of inputs.

$$\min_{x,u} \sum_{k=1}^{N-1} [(x_k - x_t)'Q(x_k - x_t) + (\Delta u_k)'S(\Delta u_k)] + (x_N - x_t)'Q_N(x_N - x_t) \quad (10)$$

$$\text{subject to:} \quad x_{k+1} = f(x_k, u_k + d_k) \quad (11)$$

Where $\Delta u_k = u_k - u_{k-1}$ and Q and S are penalty matrices. Additional constraints are imposed on the states and inputs (to ensure gas concentrations are greater than zero and surface coverages are between zero and one). The nonlinear MPC optimization was solved using the NMPC solver available in the Octave numerical computing language. This NMPC solver uses an SQP method with a trust region, and maintains the feasibility of every iterate [8].

D. State Estimator

Given the inputs and measurements over a horizon of the previous T steps, the state estimator reconstructs the model states (x_k) and the disturbances (d_k), over the horizon by solving the following optimization.

$$\min_{x_0,w,v} \sum_{k=0}^{T-1} \left[v_k' R^{-1} v_k + w_k' Q^{-1} w_k\right] + (x_0 - \overline{x}_0)'\Pi^{-1}(x_0 - \overline{x}_0) \quad (12)$$

$$\text{subject to:} \quad x_{k+1} = f(x_k, u_k + d_k) + w_k \quad (13)$$

$$y_k = g(x_k) + v_k \quad (14)$$

Where v_k is the measurement fitting error at time k, w_k are the state disturbances at time k, and $\overline{x}_0$ is the prior state estimate. R, Q, and Π are tuning matrices. Physically motivated constraints were imposed on the state estimates. The software for this optimization problem is also available in the Octave NMPC toolbox.

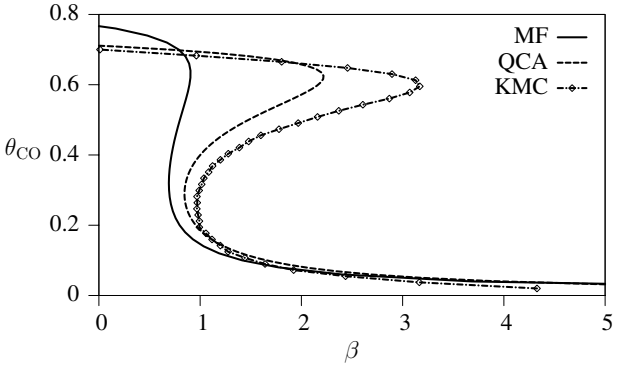

Fig. 1. MF and kMC steady-state curves, CO surface coverage as a function of O_2 gas concentration

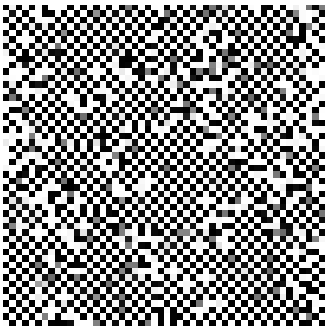

Fig. 2. Snapshot of a 50x50 lattice from a kMC simulation: CO black, O gray, empty site white

IV. Controller Models

A. Mean Field Model

Mean field (MF) theory assumes that all atoms have the same environment and are acted upon by some average field. In this application, MF theory is applied to CO oxidation by assuming that the species are uniformly distributed on a square lattice. Given the MF assumption and the proposed reaction mechanism, a model of the overall species coverages is given by 15 - 17.

$$\frac{d\theta_{\mathrm{CO}}}{dt} = \alpha\theta_* - \gamma\theta_{\mathrm{CO}} \exp\left(\frac{-4\epsilon\theta_{\mathrm{CO}}}{RT}\right) - 4k_r\theta_{\mathrm{CO}}\theta_{\mathrm{O}} \exp\left(\frac{-3\epsilon\theta_{\mathrm{CO}}}{RT}\right) \quad (15)$$

$$\frac{d\theta_{\mathrm{O}}}{dt} = 4\beta\theta_*^2 - 4k_r\theta_{\mathrm{CO}}\theta_{\mathrm{O}} \exp\left(\frac{-3\epsilon\theta_{\mathrm{CO}}}{RT}\right) \quad (16)$$

$$1 = \theta_* + \theta_{\mathrm{CO}} + \theta_{\mathrm{O}} \quad (17)$$

The states of the MF model (surface coverages: $\theta_{\mathrm{CO}}, \theta_{\mathrm{O}}, \theta_*$) are convenient for control because the surface coverages are also the measured variables.

The steady-state curve for the MF model of θ_{CO} vs. β, (controlled variable vs. input) is shown in Fig. 1. Additionally, the steady-state curve of the $\langle\theta_{\mathrm{CO}}\rangle$ vs. β for the kMC model is shown Fig. 1. The kMC steady-state curve was identified using "equation free" techniques, and is similar to the one previously identified for this system [5].

The large difference between the steady-state coverages of the MF and kMC models shown in Fig. 1 demonstrates that the MF model is a poor predictor of the actual steady-state behavior of the system, in the presence of strong lateral interactions.

Although not included in the control study, the steady-state diagram of the quasi-chemical approximation (QCA) model is also shown in Fig. 1. The QCA model is a well known model that assumes a random distribution of pairs on the lattice surface [9], [10]. The QCA model provides an accurate description of the kMC steady-state curve at intermediate interaction strengths ($\epsilon > -1.5$ kcal/mol) [5]. However, Fig. 1 shows that the QCA model does not capture the steady-state behavior of the overall coverages of the kMC model for strong lateral interactions ($\epsilon > -2.0$ kcal/mol).

B. Scheduled Linear Models

Model scheduling is a convenient method for storing information about the local dynamics of non-linear systems. For this application, a look-up table stores the A, B, F matrices, for models of the form,

$$x_{k+1} = Ax_k + Bu_k + F_k. \quad (18)$$

The states in the tabulated models are θ_{CO}, and θ_{O}. A, B, and F were identified off-line by evaluating the dynamics of $\langle\theta_{\mathrm{CO}}\rangle$, and $\langle\theta_{\mathrm{O}}\rangle$ from kMC simulations, at points along the steady-state curve of the kMC model ($\langle\theta_{\mathrm{CO}_s}\rangle, \langle\theta_{\mathrm{O}_s}\rangle, \beta_s$). A unique point along the kMC steady-state curve can be identified based on a θ_{CO} set-point, therefore the table is parameterized by the set-point of θ_{CO}. When a target calculation is performed using a controller with model scheduling, first a model is located in the look-up table (based on the θ_{CO} set-point), then the model is used to find (x_t, u_t). The identified model is then used for both regulation and state estimation.

C. Checkerboard Mean Field Model

Fig. 2 is a lattice snapshot from a kMC simulation at a relatively high coverage of CO ($\theta_{\mathrm{CO}} = 0.45$). This snapshot shows that spatial correlations develop between adsorbed species, giving the visual impression that species arrange themselves in checkerboard-like patterns and thus violate the MF assumption (random distribution of species on the lattice).

A low-order model that incorporates a checkerboard-like pattern can be developed by assuming two sub-lattices exist on the surface. If the lattice were represented by a perfect checkerboard, sub-lattice A would be the squares of one color and sub-lattice B would be the squares of the other color. Additionally, a random distribution of species on each sub-lattice is assumed.

The reaction mechanism has to be changed to account for the different sub-lattices. For example adsorbed oxygen

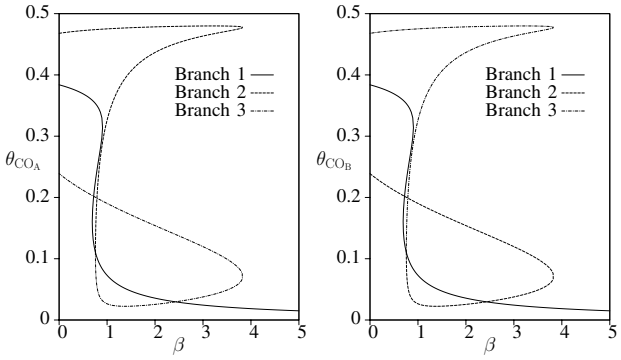

Fig. 3. Bifurcation diagrams of coverage of CO on the A and B sub-lattices, where the bifurcation parameter is the rate constant proportional to the concentration of oxygen above the surface

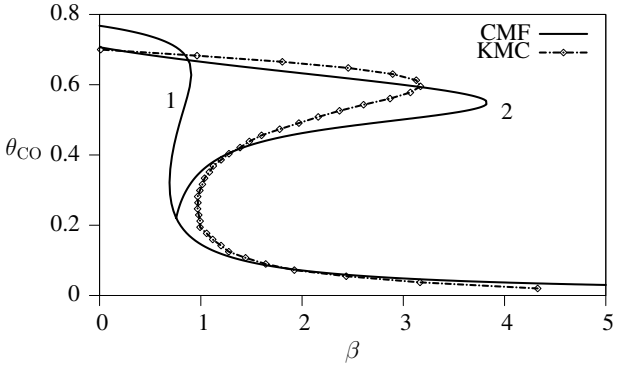

Fig. 4. CMF and kMC steady-state curves, CO surface coverage as a function of O_2 gas concentration

diffusion, as represented in Eq. 6, now comprises two events, to account for movement of adsorbed oxygen from the A sub-lattice to the B sub-lattice and vice versa.

$$\mathrm{O_A} + *_\mathrm{B} \xrightarrow{d_r} *_\mathrm{A} + \mathrm{O_B} \tag{19}$$

$$\mathrm{O_B} + *_\mathrm{A} \xrightarrow{d_r} *_\mathrm{B} + \mathrm{O_A} \tag{20}$$

The states of the Checkerboard Mean Field (CMF) model are the coverages of each species on each sub-lattice ($\theta_{\mathrm{CO_A}}, \theta_{\mathrm{CO_B}}, \theta_{\mathrm{O_A}}, \theta_{\mathrm{O_B}}, \theta_{*_\mathrm{A}}, \theta_{*_\mathrm{B}}$). The time evolution of these states are represented by 4 ODEs and 2 DAEs. For example, the evolution of $\theta_{\mathrm{O_A}}$ is given by:

$$\frac{d\theta_{\mathrm{O_A}}}{dt} = 8\beta\theta_{*_\mathrm{A}}\theta_{*_\mathrm{B}} - 8k_r\theta_{\mathrm{CO_B}}\theta_{\mathrm{O_A}}\exp\left(\frac{-6\epsilon\theta_{\mathrm{CO_A}}}{RT}\right) - 8d_r\theta_{\mathrm{O_A}}\theta_{*_\mathrm{B}} + 8d_r\theta_{\mathrm{O_B}}\theta_{*_\mathrm{A}} \tag{21}$$

Fig. 3 shows the bifurcation diagrams of the coverage of CO on the A and B sub-lattices, where the bifurcation parameter is the rate constant proportional to the concentration of oxygen above the surface. A pitchfork bifurcation occurs at $\beta = 0.76$ and $\theta_{\mathrm{CO_A}} = 0.11$. Branch 1 of these diagrams represents steady states where the species coverages on each lattice are the same ($\theta_{\mathrm{CO_A}} = \theta_{\mathrm{CO_B}}$). The Branch 1 solutions are stable below the low-$\theta_{\mathrm{CO_A}}$ saddle-node point, and unstable above that point. Branches 2 and 3 represent steady states where $\theta_{\mathrm{CO_A}} \neq \theta_{\mathrm{CO_B}}$. However, due to the symmetry of the model Branch 2 of $\theta_{\mathrm{CO_A}}$ is equivalent to Branch 3 of $\theta_{\mathrm{CO_B}}$. Branches 2 and 3 of $\theta_{\mathrm{CO_A}}$ have stable solutions on their upper most branches. Similar pitchforks and symmetry are occur in the $\theta_{\mathrm{O_A}}$ and $\theta_{\mathrm{O_B}}$ bifurcation diagrams.

Fig. 4 shows the steady-state curve for the CMF model of the total coverage of CO ($\theta_{\mathrm{CO_A}} + \theta_{\mathrm{CO_B}}$) versus β. Two regions of the CMF curve are labeled in Fig. 4. Region 1, corresponds to the Branch 1 solutions of Fig. 3 and is identical to the MF steady-state curve. Region 2 is the physically realistic region and corresponds to Branches 2 and 3 of Fig. 3. Each point along Region 2 corresponds to two steady-states that are related by symmetry.

When the CMF model is implemented in an MPC controller, the target calculation finds a set of CMF model states and the corresponding input, given a θ_{CO} set-point. The state and input targets are identified uniquely by limiting the search to Region 2 if $\theta_{\mathrm{O}} > 0.22$, and constraining the solution to $\theta_{\mathrm{CO_A}} \geq \theta_{\mathrm{CO_B}}$.

Assuming that the only θ_{CO} and θ_{O} can be measured, the state estimator is used to reconstruct the CMF model states, as well as the disturbance model, from the dynamics of the measured variables.

V. Control Examples and Discussion

Two closed-loop control examples are discussed below. For each example, MPC controllers with each of the three models (MF, scheduled linear, and CMF) are used to control θ_{CO} of a kMC plant. These examples represent a single realization of the closed-loop response of a 100x100 kMC plant. Given a different microscopic initial condition or random number sequence, we have observed that with a lattice of this size the closed-loop response is not significantly different than those shown in the following examples.

A. Set-Point Change Example

In this example, the set-point is changed from a stable set-point with low θ_{CO} to an unstable set-point with high θ_{CO} and back to the original set-point. Fig. 5 shows the closed-loop response of θ_{CO} of the kMC plant, under each controller. Fig. 6 shows the control action for each controller.

All of the controllers successfully control to the low θ_{CO} set-point, and successfully execute the set-point changes. Under MF control, the θ_{CO} of the kMC plant shows strong oscillations while locating the high θ_{CO} set-point, which is accompanied by a severe control action.

The scheduled model controller has a much more sluggish response because the set-point had to be changed gradually, so the state of the system is never too far from the region of validity of the linear model. If the full set-point change is made immediately the scheduled linear controller becomes unstable, because the model is only valid near the set-point (i.e. high θ_{CO}), but the state may be far from

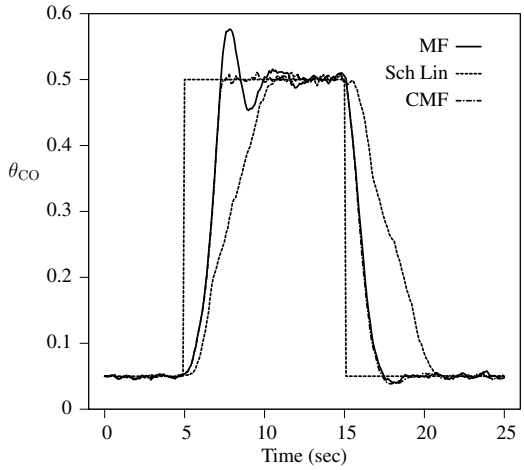

Fig. 5. θ_{CO} (controlled variable) closed-loop response, set-point tracking

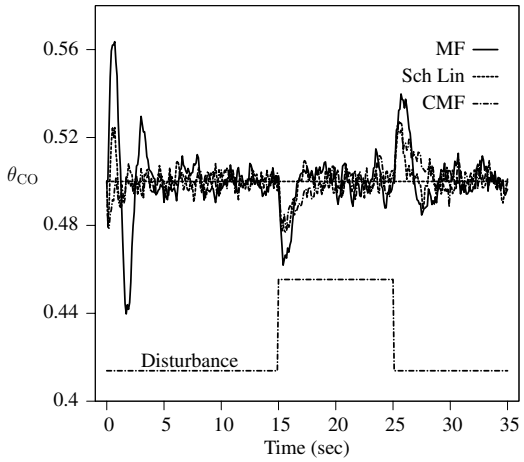

Fig. 7. θ_{CO} (controlled variable) closed-loop response, disturbance rejection

the set-point (i.e. low θ_{CO}). Full non-linear models do not suffer from these scheduling problems. The CMF model makes both set-point changes quickly and with virtually no overshoot.

B. Disturbance Rejection Example

In this example, the set-point remains constant, at an unstable and high θ_{CO}. An input disturbance is added to the plant for 10 seconds. Fig. 7 shows the closed-loop response of θ_{CO} of the kMC plant, under each controller. Fig. 8 shows the control action for each controller.

All three controllers are able to reject the disturbance and maintain the kMC plant θ_{CO} in the vicinity of the set-point before, during, and after the disturbance. Under MF control the overall CO coverage shows initial oscillations about the set-point, with similar behavior in the control action. These initial oscillations occur while the disturbance models are learning how to compensate for the plant model mismatch. The scheduled linear and the CMF models have comparatively small initial oscillations, due to less plant model mismatch. The plant with the MF controller has significantly larger offset as the disturbance starts and finishes, than the plants with the other controllers.

Some observations from the closed-loop behavior of the kMC plant with various model controllers are:

- It is difficult for the MF controller to control θ_{CO} of the kMC plant at high values of θ_{CO}, because the MF model does not predict the short-time dynamics of the plant coverages accurately. At high θ_{CO} patterns form on the lattice and the mean field assumption is no longer valid. The MF model is able to maintain θ_{CO} of the kMC plant near set-point only with the help of disturbance models.
- The scheduled model controller performs well when the plants observables are near the point of the linear model identification. This limitation requires slow set-point changes, so the plant observables are always near

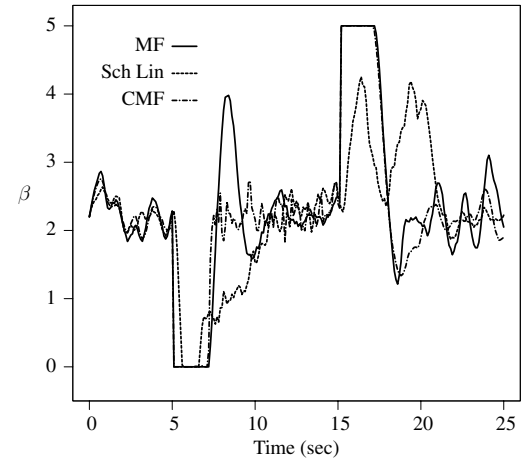

Fig. 6. Gas Concentration of O_2 (Input) closed-loop response, set-point tracking

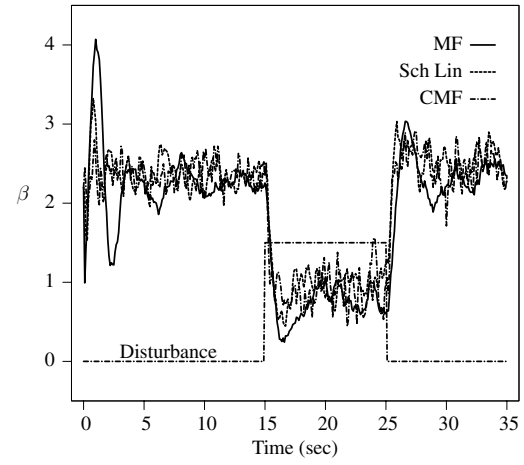

Fig. 8. Gas Concentration of O_2 (Input) closed-loop response, disturbance rejection

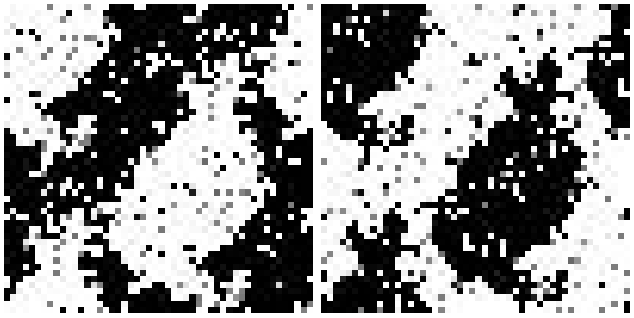

Fig. 9. Depiction of species distribution on the sub-lattices: Right sub-lattice A, Left sub-lattice B.

a steady-state point.

- The CMF controller is able to meet control objectives on the kMC plant, with fluctuations mostly induced by plant noise. The CMF model was built so that it could capture the dominant structures that appear on a kMC lattice, therefore, CMF can adequately model the short-time dynamics of the coverages of a kMC model.

VI. Future Work

Are the assumptions of the CMF model valid? We intend to answer this question in our further investigations. The CMF model assumes there is a random distribution of species within each sub-lattice. The validity of this assumption can be checked by the observing the distribution of species along the diagonals of the lattices generated from kMC simulations. Fig. 9 is a depiction of the species coverages on one set of diagonals of the lattice shown in Fig. 2. The picture on the right depicts one sub-lattice and the picture on the left depicts the other sub-lattice. These pictures were created by stacking alternating diagonals of the full kMC lattice into adjacent columns. These maps demonstrates that there are alternating regions of high and low CO coverage along lattice diagonals. The CMF model assumes a random distribution, along the diagonals and does not account for the grain boundaries observed on the kMC lattices. If the contribution of boundaries to the overall reaction rates is small compared to the contribution of the rest of the lattice, then we expect the CMF model to be an accurate model of the surface coverage evolution. Part of our future work is to determine when or if these grain boundaries make significant contributions to the reaction rates of adsorbed species.

Additionally, we intend to determine what portions of parameter space (coverages, interaction strength) can be modeled accurately with the MF, QCA, and CMF models. This study will be performed by comparing steady-state diagrams and dynamic data of these deterministic models with those of the large kMC plant.

VII. Conclusions

Representative dynamics of a surface reaction can be accurately modeled by kinetic Monte Carlo (kMC) simulations. While many researchers have used kMC models to perform systems engineering tasks, direct use of kMC models in MPC controllers is still computationally intractable. KMC models can be used, however to identify other low-order models that can be used in real-time MPC controllers.

In this study three low-order models of CO oxidation with lateral CO-CO interactions were identified off-line and used in MPC controllers: mean field (MF), scheduled linear, and checkerboard mean field (CMF). The MF model assumes a uniform distribution of species on a surface and the model's states are the species surface coverages. The scheduled linear controller comprises a table of linear state-space models, identified off-line from kMC simulations. The CMF model, has four states and accounts for checkerboard-like patterns that form during kMC simulations. These three controllers were used to control the CO coverage from a kMC simulation of CO oxidation by manipulating the O_2 gas concentration.

Closed-loop simulations of a kMC plant with these three MPC controllers show that all the controllers can meet overall control objectives, with varying performance levels. The CMF controller demonstrated the best performance because of its ability to accurately predict the short-time dynamics of the kMC plant coverages. Improved performance of the CMF model over the other models suggest that observation of kMC simulations is useful for inspiring the development of accurate low-order models.

References

[1] H. C. Kang and W. H. Weinberg, "Modeling the kinetics of heterogeneous catalysis," *Chemical Review*, vol. 95, pp. 667–676, 1995.

[2] S. Raimondeau, P. Aghalayam, A. B. Mhadeshwar, and D. G. Vlachos, "Parameter optimization of molecular models: Application to surface kinetics," *Ind. Eng. Chem. Res.*, vol. 42, no. 6, pp. 1174–1183, March 2003.

[3] M. A. Gallivan and R. M. Murray, "Reduction and identification methods for markovian control systems, with application to thin film deposition," *Int. J. Robust and Nonlinear Control*, vol. 14, pp. 113–132, 2004.

[4] Y. Lou and P. D. Christofides, "Feedback control of growth rate and surface roughness in thin film growth," *AIChE J.*, vol. 49, no. 8, pp. 2099–2113, August 2003.

[5] A. G. Makeev, D. Maroudas, A. Z. Panagiotopoulos, and I. G. Kevrekidis, "Coasrse bifurcation analysis of kinetic Monte Carlo simulations: A lattice-gas model with lateral interactions," *J. Chem. Phys.*, vol. 117, no. 18, pp. 8229–8240, November 2002.

[6] C. I. Siettos, M. D. Graham, and I. G. Kevrekidis, "Coarse brownian dynamics for nematic liquid crystals: Bifurcation, projective integration, and control via stochastic simulation," *J. Chem. Phys.*, 2003.

[7] C. I. Siettos, A. Armaou, A. G. Makeev, and I. G. Kevrekidis, "Microscopic/stochastic timesteppers and "coarse" control: A KMC example," *AIChE J.*, vol. 49, no. 7, p. 1922, 2003.

[8] M. J. Tenny, S. J. Wright, and J. B. Rawlings, "Nonlinear model predictive control via feasibility-perturbed sequential quadratic programming," *Comp. Optim. Appl.*, vol. 28, no. 1, pp. 87–121, April 2004. [Online]. Available: http://www.kluweronline.com/issn/0926-6003/

[9] V. P. Zhdanov, "Lattice-gas model of bimolecular reaction of surface," *Surface Science*, 1981.

[10] S. Sundaresan and K. R. Kaza, "Non-random distribution of adsorbates on catalytic surfaces: The role of adsorbate mobilities on reaction rates," *Chem. Eng. Commun.*, 1985.