

Identification and Online Updating of Dynamic Models for Demand Response of an Industrial Air Separation Unit

Calvin Tsay* Yanan Cao** Yajun Wang**
Jesus Flores-Cerrillo** Michael Baldea*

* *McKetta Department of Chemical Engineering, The University of Texas at Austin, Austin, TX 78712 USA (e-mail: calvint@utexas.edu, mbaldea@che.utexas.edu).*

** *Smart Operations, Center of Excellence (COE), Linde, Tonawanda, NY 14150 USA (e-mail: yanan.cao@linde.com, yajun.wang@linde.com, jesus.flores-cerrillo@linde.com).*

Abstract: Demand-response operation of air separation units often involves frequent changes in process setpoints (e.g., production rate), and therefore process dynamics should be considered in scheduling calculations to ensure feasibility. To this end, scale-bridging models (SBMs) approximate the scheduling-relevant dynamics of a closed-loop process in a low-order representation. In contrast to previous works that have employed nonlinear SBMs, this paper proposes linear SBMs, developed using time-series analysis, to facilitate online scheduling computations. Using a year-long industrial dataset, we find that compact linear SBMs are suitable approximations over typical scheduling horizons, but that their accuracies are unpredictable over time. We introduce a strategy for online updating of the SBMs, based on Kalman filtering schemes for online parameter estimation. The approach greatly improves the accuracy of SBM predictions and will enable the use of linear SBM-based demand-response scheduling in the future.

Keywords: Demand-side management, industrial big data, production scheduling and control

1. INTRODUCTION

Cryogenic air separation units (ASUs) are strong candidates for demand-response operation. The primary cost of ASU operation is associated with electricity usage, as compressors driven by large electric motors are used to bring feed air to the requisite operating pressures. As a result, process economics can be greatly improved by proactively scheduling production in response to (predicted) fluctuations in electricity price—which may become more extreme with increasing adoption of renewable-based generation. Moreover, the products of an ASU comprise purified components of air, which can be liquefied and stored. The above factors motivate the development of production schedules focused on “load shifting”: increasing production when electricity prices are low and storing excess products in liquefied form, which can later be used to satisfy product demand(s) when electricity prices are high.

There are several modeling considerations for demand-response scheduling of ASUs. To fully exploit time-varying electricity prices, which change hourly (or more frequently), production schedules must consider hourly (or shorter) scheduling intervals. Using such scheduling intervals requires that ASU dynamics be considered explicitly to ensure that production schedules are both dynamically feasible and economically optimal. Embedding a representation of process dynamics in production scheduling models requires a multi-scale approach that bridges the (faster) time scales of process dynamics/control with those

of production scheduling (Jamaludin and Swartz, 2017). However, such models must be used to optimize process operation over relatively long time horizons, yet still accurately represent complex process behavior over all relevant time scales. Several reduced-order modeling approaches (Cao et al., 2016; Caspari et al., 2020a; Schäfer et al., 2019) have been proposed for ASUs, mostly derived using physical insight and engineering expertise. Data-driven, system identification-based methods have also been described (Pattison et al., 2016; Tsay and Baldea, 2019).

In particular, our recent works (Tsay et al., 2019; Tsay and Baldea, 2020) demonstrated that *nonlinear* scale-bridging models (SBMs), or low-order representations of closed-loop process dynamics, can be identified from historical ASU operating data. Production scheduling is then cast as a dynamic optimization problem using SBMs to represent the dynamics of the ASU and its control system. Nevertheless, developing nonlinear SBMs requires significant system identification effort. Furthermore, SBM approximations may not remain accurate over extended periods of time, and the parameters may periodically require re-estimation. This degradation of accuracy can be attributed to the SBM not capturing all process/controller dynamics, to system identification only being performed on certain regions of the input space, and/or to changes in the process dynamics and control system over time.

Motivated by these challenges, in this paper we introduce an approach for online updating of SBMs used for produc-




Fig. 1. Scale-bridging model (SBM) concept. The shaded area reflects the system whose input-output dynamics are represented by the SBM.

tion scheduling. Using data recorded over a year of routine operation of an industrial cryogenic ASU, we show that the proposed Kalman filter-based strategy enables the use of low-order, *linear* SBMs (in the form of simple time-series models) over typical scheduling time horizons. We show that the online update strategy maintains model accuracy over the relevant time scales, which would otherwise degrade quickly over time.

2. PROCESS AND DATA DESCRIPTION

In a cryogenic ASU, the feed stream passes through a feed air compressor (FAC) before being sent to the heat exchanger and separation columns. Further details regarding ASU processes and their control systems can be found in, e.g., Caspari et al. (2020b); Tsay et al. (2019). The FAC accounts for the majority of the electricity consumed by the ASU, and demand-response scheduling calculations therefore rely on accurate predictions of the FAC power consumption, denoted as \dot{W}_{FAC} , as a function of production rate and other factors. Given its importance, in this work we focus on the modeling of \dot{W}_{FAC} . Note that the techniques described here can be applied to all scheduling-relevant variables needed to formulate a complete production scheduling optimization problem (Pattison et al., 2016; Tsay and Baldea, 2020).

Introduced by Du et al. (2015), SBMs aim to represent the closed-loop, input-output response of process variables, with the inputs being controller setpoints and/or measured disturbance variables (Figure 1). As a result, system identification is performed on the *closed-loop* response of the process, in contrast to conventional *open-loop* system identification (where the input data are process inputs). The ASU operates under multivariable model predictive control (MPC), with eight operator setpoints available. Ambient temperature (T) acts as a measured disturbance variable. Data for setpoints, T , and \dot{W}_{FAC} were recorded at one-minute intervals over one year of routine operation.

An unexpected feature of the dataset was that data for six of the setpoints are strongly correlated. Principal component analysis (PCA) for these six setpoints reveals that 90% of the variation is explained by the first two principal components. Due to this correlation, the individual effects of these six setpoints cannot be isolated using the available data, and they are replaced with the first two principal components (ϕ_1, ϕ_2) to reduce model dimensionality. Of the remaining two setpoints, one is uncorrelated (<1%) with \dot{W}_{FAC} and is not treated as an SBM input. The other, denoted as SP_1 , is treated as an input. In total, four inputs were used to construct the SBM: $\phi_1, \phi_2, \text{SP}_1$, and T .




Fig. 2. Normalized AIC for varying model order N and number of days in D_{train} . Error bars denote one standard deviation.

All input and output variables are normalized and filtered to reduce high-frequency noise, and to protect the confidential industrial data. SBM input variables are filtered using a low-pass filter with a cutoff of 1 h^{-1} , while the output variable (\dot{W}_{FAC}) is treated using a Savitzky-Golay filter (Orfanidis, 1995) with a window of 15 samples.

3. SCALE-BRIDGING MODEL DEVELOPMENT

3.1 Time-Series Modeling

Our previous work (Tsay et al., 2019) showed that *non-linear* SBMs (i.e., Hammerstein-Wiener models) can accurately represent process dynamics over the span of several months. While this avoids frequent re-fitting of the SBM parameters, demand-response scheduling horizons are typically only several days in length, suggesting that even simpler model forms can be used. Therefore, in this work, we study *linear* SBMs, which simplify both the system identification step and potentially the later dynamic-optimization-based scheduling step. For example, Kelley et al. (2018) formulated the latter as a MILP by linearizing SBM dynamics, while integer variables were introduced to handle SBM nonlinearities. A linear SBM could eliminate the need for integer variables, making the scheduling problem an LP. We note that Dias et al. (2018) proposed embedding linear, *open-loop* dynamic models along with the process MPC in scheduling calculations.

We consider linear SBMs in the form of an autoregressive with extra inputs (ARX) time-series model:

$$y_t + \sum_{j=1}^N a_j y_{t-j} = \sum_{i=1}^M \sum_{j=1}^N b_{i,j} u_{i,t-j} + e_t \quad (1)$$

where y_t is the output value at time t , and $u_{i,t}$ is the value of the i^{th} input ($i = 1, \dots, M$) at time t . The coefficients a_j are autoregressive parameters, $b_{i,j}$ are the regressed coefficients for the i^{th} input, and e is the model error. Note that the linear time-series model (1) can easily be converted to a continuous-time form, such as a transfer function or state-space model. In practice, the

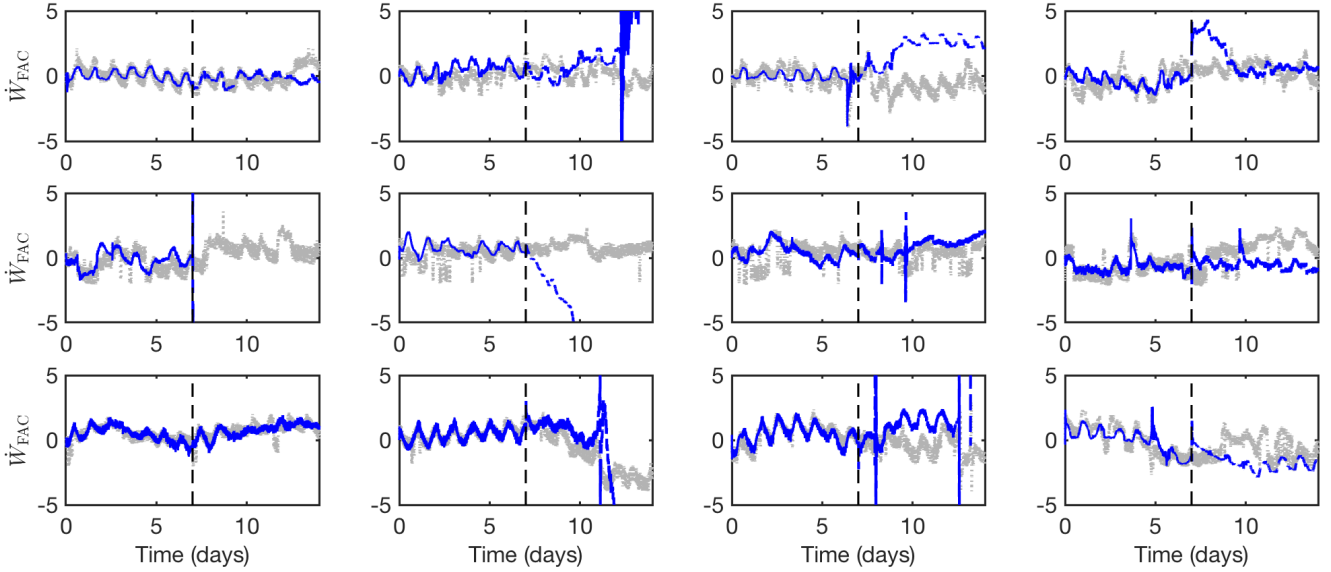


Fig. 3. SBM predictions for each month. The left-hand side of each plot (left of the dashed line) shows predictions for training data, while the right-hand side shows predictions over the next seven days. Recorded data are shown in gray.




Fig. 4. Scaled SBM parameter values. Each plot/ring represents one SBM (one for each of the 12 months).

number of autoregressive terms (number of poles) and the number of coefficients for each input (number of zeros) are independent; however, for simplicity we treat both of these as a single model order, denoted as N .

We consider each month of the one-year dataset separately, yielding 12 datasets in total. SBM models of the form (1) are trained using the first D_{train} days of data for each month. Figure 2 shows the normalized Akaike Information Criterion (nAIC) for several different model orders N and values of D_{train} . Models with $N=2$ and $N=3$ have similar nAIC, particularly as D_{train} increases. Thus, the rest of this work considers SBMs of the form (1) with $N=3$, which provide a good tradeoff between model size and accuracy.

3.2 System Identification Results

Figure 3 shows the predictions of the trained SBMs on seven days of training data, as well as the next seven days. The mean squared error (MSE) for predictions on training data (left side of each plot) for the twelve datasets is 0.47 ± 0.27 (mean \pm standard deviation), while the MSE for predictions on new data (right side of each plot) is 51.67 ± 94.78 . The simple, third-order ARX model can

accurately represent the closed-loop process dynamics over one week of training data, but its accuracy is unpredictable over time; some SBMs in Figure 3 make good predictions on new data, while others quickly lose accuracy.

The fitted parameter values for the trained SBMs are shown in Figure 4. The plotted values are scaled by the respective mean over the 12 SBMs. While the values of some parameters (e.g., $b_{3,j}; j = 1, 2, 3$) varied significantly among SBM models, the values of a_j and $b_{4,j}$ were nearly constant. Interestingly, this data-driven approach reveals that the system poles (i.e., closed-loop time constants) do not change significantly over the 12 monthly datasets. The dependence of \dot{W}_{FAC} on u_4 (ambient temperature T) also remains relatively stable. These observations confirm the physical intuition that the response of the process to the controller setpoints $\{u_1, u_2, u_3\} = \{\phi_1, \phi_2, \text{SP}_1\}$ is nonlinear and/or time-varying (e.g., as the control system is manipulated by process operators), and cannot be represented accurately using a linear dynamic model with fixed parameters. On the other hand, the poles and the response to changes in ambient temperature are characteristic to the process itself and do not change over the time span of the dataset.

While Figure 3 shows the SBM predictions over a fixed horizon, scheduling applications rely on SBM predictions over a *moving* horizon. Consider a scheduling horizon $t_{\text{sched}} = 4$ days. At time $t = 0$ days the SBM would be used to make predictions for days 0–4, and at time $t = 1$, for days 1–5, etc. In other words, at any time t we are interested in the instantaneous accuracy of the SBM predictions over the interval $[t, t + t_{\text{sched}}]$, computed as:

$$\text{MSE}_t = \frac{1}{n_{\text{sched}}} \sum_{t'=t}^{t+n_{\text{sched}}} (y_{t'} - \hat{y}_{t'})^2 \quad (2)$$

where $y_{t'}$ is the output value at time t' (1), $\hat{y}_{t'}$ is its predicted value, and n_{sched} is the number of samples taken during the interval t_{sched} . Figure 5 shows MSE_t (2) for the




Fig. 5. SBM accuracy over a four-day moving horizon. Each gray line depicts one SBM (from one monthly dataset), while the blue line indicates the median values. The dashed vertical lines mark days 3–7.

same SBMs in Figure 3 evaluated each hour for $t_{\text{sched}} = 4$ days. The first seven days comprise the SBM training data, as in Figure 3. The vertical dashed lines mark times from $t = 3$ days, where the four-day intervals first begin to include new data, to $t = 7$ days, where the four-day intervals comprise completely new data.

For most of the SBMs, the four-day MSEs are relatively low for days 0–3. As data outside the training set are introduced (days 3–7), the four-day MSEs begin increasing. The MSEs are highest past day seven, when all data are from outside the training set. Like the predictions in Figure 3, the moving-horizon prediction quality for these linear SBMs degrades after a few days, showing that re-fitting is needed for accurate predictions in scheduling calculations.

4. ONLINE PARAMETER UPDATING

4.1 Kalman Filter for Parameter Values

While model parameters can be periodically re-fitted (e.g., moving-horizon estimation), this work investigates online updating methods. Along with least-squares methods, the Kalman filter is a well-established strategy for online parameter estimation (Ljung and Gunnarsson, 1990), and its behavior and stability in this context have been analyzed by, e.g., Cao and Schwartz (2004); Guo (1990). Other works have also employed nonlinear extensions of the Kalman filter for online parameter estimation, e.g., the extended Kalman filter (Graichen et al., 2006), or the unscented Kalman filter (Radecki and Hencze, 2012).

The underlying assumption is that the system state variables comprise the vector of model parameters, which vary via a random walk:

$$\boldsymbol{\theta}_t = \boldsymbol{\theta}_{t-1} + \boldsymbol{\omega}_t \quad (3)$$

$$\boldsymbol{\omega}_t \sim N(\mathbf{0}, \mathbf{Q}) \quad (4)$$

where $\boldsymbol{\theta}$ is the vector of parameters, and \mathbf{Q} is the system noise covariance matrix. For the ARX model (1), $\boldsymbol{\theta} = [-a_1, \dots, -a_N, b_{1,1}, \dots, b_{1,N}, \dots, b_{M,1}, \dots, b_{M,N}]^T$ and $\boldsymbol{\theta}_t$




Fig. 6. SBM accuracy over a four-day moving horizon with online parameter estimation. Each gray line depicts one SBM, while the blue line indicates the median values. The dashed vertical lines mark days 3–7.

is the estimate of $\boldsymbol{\theta}$ at time t . The model output is computed as a linear combination of the inputs:

$$y_t = \mathbf{u}_t^T \boldsymbol{\theta}_t + e_t \quad (5)$$

$$e_t \sim N(0, R) \quad (6)$$

where \mathbf{u}_t^T is the vector of inputs, and R is the residual variance. For the SBM, $\mathbf{u}_t^T = [\phi_{1,t-1}, \dots, \phi_{1,t-N}, \phi_{2,t-1}, \dots, \phi_{2,t-N}, \text{SP}_{1,t-1}, \dots, \text{SP}_{1,t-N}, T_{t-1}, \dots, T_{t-N}]$, making (1) and (5) equivalent. For this system, Ljung and Gunnarsson (1990) showed that a Kalman filter provides the best estimation of $\boldsymbol{\theta}_t$ when \mathbf{Q} and R are exactly known. Nevertheless, when \mathbf{Q} and R are unknown and/or the random walk (3) does not model the true parameter dynamics well, the Kalman filter can often still track time-varying parameters accurately (Cao and Schwartz, 2004; Niedzwiecki, 2000).

The Kalman filter for (3)–(6) can be written as:

$$\hat{\boldsymbol{\theta}}_t = \hat{\boldsymbol{\theta}}_{t-1} + \mathbf{K}_t(y_t - \hat{y}_t) \quad (7)$$

$$\hat{y}_t = \mathbf{u}_t^T \hat{\boldsymbol{\theta}}_{t-1} \quad (8)$$

$$\mathbf{K}_t = \mathbf{P}_{t-1} \mathbf{u}_k (R + \mathbf{u}_k^T \mathbf{P}_{t-1} \mathbf{u}_k)^{-1} \quad (9)$$

$$\mathbf{P}_t = [\mathbf{I} - \mathbf{K}_t \mathbf{u}_k^T] \mathbf{P}_{t-1} + \mathbf{Q} \quad (10)$$

where $\hat{\boldsymbol{\theta}}_t$ and \hat{y}_t are the estimates for the parameter vector and the output, respectively, at time t . The vector \mathbf{K}_t is the filter gain, and \mathbf{P}_t is the estimated covariance matrix, both at time t . To maintain both nonzero \mathbf{K}_t and stability, \mathbf{P}_t should satisfy the standard matrix inequalities $\alpha \mathbf{I} \leq \mathbf{P}_t \leq \beta \mathbf{I}, \forall t$, where α and β are positive scalars.

We fix the residual variance $R = 1$ for simplicity, leaving \mathbf{Q} and \mathbf{P}_0 as tuning parameters for the filter. Assuming that the initial parameter estimates $\boldsymbol{\theta}_0$ are good at $t = 0$, the initial covariance \mathbf{P}_0 should be small, and we set each parameter variance proportionally, such that $\mathbf{P}_0 = \text{diag}(\boldsymbol{\theta}_0) \times 0.1\%$. For the system noise covariance \mathbf{Q} , which has less physical intuition, the parameter covariance matrix across the 12 training datasets, $\boldsymbol{\Sigma}$, was first computed. We then set $\mathbf{Q} = \boldsymbol{\Sigma}/n_{\text{train}}$, where n_{train} is the number of

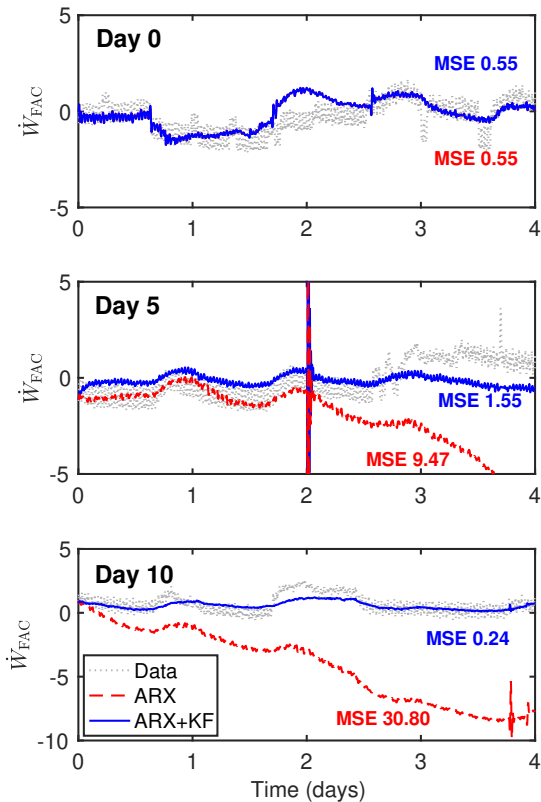


Fig. 7. Four-day predictions for ARX model with and without Kalman filter (KF) parameter updates at day zero (top), day five (middle), and day ten (bottom).

samples in the training dataset, such that the variance after n_{train} time steps is Σ , or $(\theta_t - \theta_{t-n_{\text{train}}}) \sim N(\mathbf{0}, \Sigma)$.

4.2 Online Parameter Estimation Results

The Kalman filter (7)–(10) was applied to the same SBM models using the one-minute process sampling interval. The initial parameter values θ_0 for each SBM were set to their respective estimated values from Section 3. Figure 6 shows the same hourly MSE_t (2), over a scheduling horizon $t_{\text{sched}} = 4$ days, for the SBMs as in Figure 5, but with the SBM parameters updated at each one-minute sample via Kalman filter. For most of the SBMs, the four-day MSEs are again relatively low for days 0–3. While the four-day MSEs for some SBMs similarly begin increasing at day 3, the MSEs remain lower in general, owing to the updated parameter estimates. Unlike the SBMs with no Kalman filter in Figure 5, the median MSE remains relatively flat overall, showing that the online parameter estimation strategy effectively removes the need to periodically re-fit SBMs for use in scheduling calculations.

The four-day predictions for the SBMs with and without parameter updates are shown in Figure 7 for month five. These predictions correspond to the time points at zero, five, and ten days in Figures 5 and 6. Month five is shown because the SBM quickly loses accuracy without parameter updates (Figure 3). At day zero, the the Kalman filter has not updated the model parameters, and the two models are the same, with an MSE over the four-day




Fig. 8. Scaled SBM parameter estimates updated by Kalman filter over time. The three lines in each series depict three time-lagged coefficients.

window of 0.55. At day five, the time window comprises two days in the training dataset and two new days. Without parameter updates, the model is inaccurate for the two new days (Figure 7, middle), giving an overall MSE of 9.47. With parameter updates, the model maintains an MSE of 1.55. At day ten, when the four-day window does not include any training data, the updating strategy decreases the prediction MSE from 30.80 to 0.24. Note that the predictions at days five and ten differ from those in Figure 3 because the initial conditions are updated when the models are evaluated in a moving horizon.

4.3 Parameter Dynamics

The values of the estimated parameter values may change significantly over time as they are updated by the Kalman filter (7)–(10). These parameter “dynamics” are shown in Figure 8 for month five. As expected, the values of autoregressive parameters a_j remain relatively constant in time. Since the historical data (Σ) were used to tune the Kalman filter (Q), the low (co)variances of a_j in Figure 4 result in low variations in their filtered dynamics. Similarly, the values of $b_{4,j}$ only increase by approximately a factor of 1.5, which is similar to their variation across monthly datasets (Figure 4). The values of $b_{1,j}$ and $b_{2,j}$ increase relatively quickly, and future study should investigate the stability of SBM parameters with online estimation and/or alternative Kalman filter tuning(s). However, note that the four-fold increase is of a similar magnitude as the variations in $b_{1,j}$ and $b_{2,j}$ shown in Figure 4. We verified that the parameters remained bounded for this process when the Kalman filter is applied over the entire year-long dataset.

While some parameters may approach new steady-state values (perhaps indicating a transition in the process dynamics), others may continue to fluctuate. For example, if the relationship between y and u_i is nonlinear, estimates of the parameters $b_{i,j}$ will be updated to approximate the local input-output response. Given the above, the dynamics of the SBM parameters could be monitored to reveal changes in the process/controller dynamics, or potentially aid in fault diagnoses. The dynamics of the

SBM parameters may also reveal differences between ASU plants. For instance, in a transfer learning approach, the SBM parameters from one ASU could be used as the initial estimates θ_0 for another facility that has a similar design. Online parameter estimation can “learn” the SBM parameters of the new plant over time, revealing differences (or confirming similarities) between the plants.

5. CONCLUSIONS

Optimal demand-response scheduling of industrial ASUs should account for process dynamics and control to ensure that schedules are dynamically feasible. Here, scale-bridging models (SBMs) can represent closed-loop process dynamics in production scheduling using a low-order approximation. This work presents a data-driven strategy to create linear SBMs for an industrial process from routine operational data. We find that compact time-series models can accurately represent process dynamics over several days, but that their accuracy can suffer over time.

To mitigate this issue, we apply Kalman filtering for online updating of the SBM parameters, including a tuning that incorporates the variations on the parameters in historical data. The strategy enables the SBM to adapt to nonlinear behavior and/or changes in the process dynamics, thereby greatly improving the accuracy of SBM predictions over four-day scheduling horizons. Furthermore, the dynamics of the parameter estimates can reveal changes in the process dynamics and potentially be used in monitoring applications. The Kalman filtering strategy can be used to adapt SBM to changes over extended periods of time, or to transfer an SBM from one plant to another.

ACKNOWLEDGEMENTS

This work was funded by CESMII Project *Smart Manufacturing for Chemical Processing*, and National Science Foundation CAREER Award 1454433. This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

REFERENCES

Cao, L. and Schwartz, H.M. (2004). Analysis of the Kalman filter based estimation algorithm: an orthogonal decomposition approach. *Automatica*, 40(1), 5–19.
 Cao, Y., Swartz, C.L., Flores-Cerrillo, J., and Ma, J. (2016). Dynamic modeling and collocation-based model reduction of cryogenic air separation units. *AIChE J.*, 62(5), 1602–1615.

Caspari, A., Offermanns, C., Ecker, A.M., Pottmann, M., Zapp, G., Mhamdi, A., and Mitsos, A. (2020a). A wave propagation approach for reduced dynamic modeling of distillation columns: Optimization and control. *J. Process Control*, 91, 12–24.
 Caspari, A., Tsay, C., Mhamdi, A., Baldea, M., and Mitsos, A. (2020b). The integration of scheduling and control: Top-down vs. bottom-up. *J. Process Control*, 91, 50–62.
 Dias, L.S., Pattison, R.C., Tsay, C., Baldea, M., and Ierapetritou, M.G. (2018). A simulation-based optimization framework for integrating scheduling and model predictive control, and its application to air separation units. *Comput. Chem. Eng.*, 113, 139–151.
 Du, J., Park, J., Harjunkoski, I., and Baldea, M. (2015). A time scale-bridging approach for integrating production scheduling and process control. *Comput. Chem. Eng.*, 79, 59–69.
 Graichen, K., Hagenmeyer, V., and Zeitz, M. (2006). Feedforward control with online parameter estimation applied to the Chylla–Haase reactor benchmark. *J. Process Control*, 16(7), 733–745.
 Guo, L. (1990). Estimating time-varying parameters by the Kalman filter based algorithm: stability and convergence. *IEEE Trans. Autom. Control*, 35(2), 141–147.
 Jamaludin, M.Z. and Swartz, C.L. (2017). Approximation of closed-loop prediction for dynamic real-time optimization calculations. *Comput. Chem. Eng.*, 103, 23–38.
 Kelley, M.T., Pattison, R.C., Baldick, R., and Baldea, M. (2018). An MILP framework for optimizing demand response operation of air separation units. *Appl. Energy*, 222, 951–966.
 Ljung, L. and Gunnarsson, S. (1990). Adaptation and tracking in system identification—a survey. *Automatica*, 26(1), 7–21.
 Niedzwiecki, M. (2000). *Identification of time-varying processes*. Wiley, Hoboken, NJ.
 Orfanidis, S.J. (1995). *Introduction to signal processing*. Prentice Hall, Upper Saddle River, NJ.
 Pattison, R.C., Touretzky, C.R., Johansson, T., Harjunkoski, I., and Baldea, M. (2016). Optimal process operations in fast-changing electricity markets: framework for scheduling with low-order dynamic models and an air separation application. *Ind. Eng. Chem. Res.*, 55(16), 4562–4584.
 Radecki, P. and Hency, B. (2012). Online building thermal parameter estimation via unscented Kalman filtering. In *American Control Conference*, 3056–3062.
 Schäfer, P., Caspari, A., Kleinhans, K., Mhamdi, A., and Mitsos, A. (2019). Reduced dynamic modeling approach for rectification columns based on compartmentalization and artificial neural networks. *AIChE J.*, 65(5), e16568.
 Tsay, C. and Baldea, M. (2019). 110th anniversary: Using data to bridge the time and length scales of process systems. *Ind. Eng. Chem. Res.*, 58(36), 16696–16708.
 Tsay, C. and Baldea, M. (2020). Integrating production scheduling and process control using latent variable dynamic models. *Control Eng. Pract.*, 94, 104201.
 Tsay, C., Kumar, A., Flores-Cerrillo, J., and Baldea, M. (2019). Optimal demand response scheduling of an industrial air separation unit using data-driven dynamic models. *Comput. Chem. Eng.*, 126, 22–34.