

An empirical study of moving horizon closed-loop demand response scheduling

Morgan T. Kelley¹, Ross Baldick², Michael Baldea^{1*}

¹*McKetta Department of Chemical Engineering,* ²*Department of Electrical and Computer Engineering*
The University of Texas at Austin, Austin, Texas 78712, United States[†]

Abstract

The potential of electricity-intensive chemical plants to engage in demand response (DR) initiatives in support of power grid operations has been the subject of many conceptual studies. In this work, using an industrially-relevant model of an air separation unit, we undertake an extensive simulation study of moving horizon (MH) rescheduling approaches, where we “close the scheduling loop” based on updated information regarding disturbances such as changes in electricity prices, ambient conditions and chemical product demand. Our study produces several unexpected findings regarding the non-periodic nature of rescheduling solutions, the impact of the accuracy of disturbance forecasts on the economics of DR scheduling, and on the interplay of simultaneously dealing with fluctuations on both the supply side (i.e., electricity prices) and the product demand side of the plant. We posit that the latter fluctuations pose significant limitations to the potential of a chemical plant to engage in DR.

keywords: Closed-loop scheduling; moving horizon; rescheduling; air separation unit

1 Introduction

The growing contribution of renewable-based electricity generation presents significant environmental benefits, but has led to an increased supply-side variability and uncertainty for the power grid [1, 2]. Coupled with existing time-of-day and seasonal variability of electricity *demand*, this phenomenon makes it increasingly challenging to balance electricity supply and demand in grid operations. An attractive approach for mitigating this imbalance is demand-side management, a set of initiatives

that focus on reshaping electricity use patterns rather than controlling power generation. Price-based demand-side management, or demand response (DR), relies on time-of-use electricity pricing to influence user-level load shifting decisions. Industrial, commercial, and, to a lower extent, residential users can choose (or are mandated to use) time-varying electricity prices that reflect power demand fluctuations. Broadly speaking, daily electricity prices evolve in “antiphase” with grid demand levels, reaching their lowest values during the low demand hours early in the morning, and peaking during peak demand times late in the afternoon.

Industrial users are particularly appealing DR participants, accounting for over 30% of annual electricity consumption in the US [3]. While this figure is lower than the value posted by residential buildings [3], manufacturing plants offer distinct DR benefits: they are large, localized loads that can be coordinated by a single decision-maker (the operator of the facility), and typically exhibit a lower endogenous fluctuation in electricity use. In contrast, residential users are a heterogeneous group of small loads that experience significant electricity consumption fluctuations dictated by endogenous factors such as occupancy and variations in human preferences (e.g., temperature setpoints of heating, ventilation and air conditioning (HVAC) equipment).

Chemical production and petroleum refining make up a significant portion of industrial electricity use (46% in 2018 [4]), suggesting that they are natural candidates for DR. Demand-side management participation of industrial plants includes, e.g., consideration of bidding algorithms for load reduction [5], increasing plant agility [6], and optimal scheduling [7, 8, 9]. In the latter case, which we consider here, DR participation involves a production scheduling strategy comprising two complementary events. Production rates are increased during off-peak electricity demand periods. This means using more electricity but at cheaper rates, resulting in overproduction of chemical products. Products made in excess of demand are stored and used to supply customers during heavy grid load periods, when the production rate (and

*Corresponding Author: mbaldea@che.utexas.edu

[†]Morgan Kelley and Michael Baldea are with McKetta Department of Chemical Engineering, The University of Texas at Austin, Austin, Texas 78712, United States. Ross Baldick is with the Department of Electrical and Computer Engineering, The University of Texas at Austin, Austin, Texas 78712, United States.

thereby electricity demand) of the chemical plant is decreased. Thus, DR engagement, as described above, can be regarded as a way to store electricity in the form of chemical products. In order to participate in DR programs, a chemical plant must be able to operate above its nominal product demand level and store any excess product safely and without significant quality degradation. More importantly, engaging in DR programs requires that a plant have the dynamic agility to change its production rate on a time scale comparable to that of the frequency of electricity price changes (hourly or less) established by electricity markets. Given that in many practical situations, the dominant plant dynamics evolve over the same (or longer) time scales as electricity prices, DR scheduling calculations must account for process dynamics. It is also beneficial for scheduling calculations to account for relatively long (in comparison with the process time constant) time horizons. In practice, this amounts to scheduling production over a few days. This can be problematic for several reasons: electricity prices are not known accurately for time spans exceeding 24 hours, product demand can fluctuate, and plants are naturally subjected to uncertainty regarding operating conditions (changes in ambient temperature are a good example). Forecasts of these disturbance variables are typically used. Predictions can be quite good for time instants in the near future relative to the time when the forecast is made, but their accuracy declines as longer horizons (typically, more than 24 hours) are considered.

There are two broad classes of mechanisms for accounting for such uncertainty. The first is schedule optimization under uncertainty, which requires that some quantitative description of the uncertainty be available, and, depending on the approach taken, can provide a scheduling solution with a known degree of conservativeness. The second is implementing feedback (in the sense of updating the scheduling solution as new information becomes available), which represents a natural way of dealing with exogenous factors whose values cannot be predicted easily for future time instants, but can be measured accurately at the current time instant. The latter approach naturally leads to moving horizon scheduling formulations, which rely on periodically recomputing the scheduling solution as new information concerning the uncertain variables becomes available. The time horizon for the scheduling calculation remains constant, and is “shifted” forward at each rescheduling point.

Several works have utilized moving horizon structures to mitigate uncertainty (both endogenous and exogenous) that arises in scheduling problems. Gupta and Maravelias [10] addressed closed-loop task scheduling using moving horizon scheduling formulations subject to endogenous uncertainty. Shyamal and Swartz [11] scheduled electric arc furnaces with a moving horizon implementation to periodically re-evaluate fractional energy inputs between chemical (e.g. natural gas) and electri-

cal sources based on time-varying electricity prices. He and Petit [12] solved a grid-side DR scheduling problem subject to uncertainties in renewable energy sources and consumption with a two-stage moving horizon framework. Coelho et al. [13] performed real-time byproduct gas scheduling for iron and steel making applications on a moving horizon and found that a reduced control horizon (resulting in frequent rescheduling) for implementation on a moving horizon led to increased system stability. Mathur et al. [14] performed moving horizon online scheduling of cascaded hydropower systems subject to uncertainty to aid in resolving scheduling “nervousness,” (a phenomenon where significant variation in schedules can result from even slight fluctuations in inputs, such as electricity prices). In our previous work [15], we introduced a framework for moving-horizon, closed-loop DR scheduling with a focus on the problem formulation, using dynamic process models. In this paper, using the model of an air separation unit (ASU), we present an extensive discussion, focusing on the practical circumstances that may be encountered in the implementation of such a strategy for chemical plants. The computational efficiency of our models enables us to consider a significant number of scenarios compared to other works and can therefore provide a comprehensive picture of the way accounting for uncertainty affects scheduling solutions and operating cost.

Specifically, the key contributions of this effort are:

- an extensive exploration of the impact of typical exogenous disturbances on DR scheduling, using a moving horizon approach enabled by computationally tractable reduced-order models representing the closed-loop nonlinear plant dynamics
- a discussion of methods and models of accounting for exogenous disturbances and uncertainties such as demand changes, price fluctuations and variations in ambient temperature
- a framework for optimization under uncertainty combined with moving horizon scheduling, which ensures the feasibility of the moving horizon DR scheduling problem in the case where the *duration* (rather than the *magnitude*) of a non-periodic, non-persistent disturbance (product demand) is not known or cannot be accurately predicted

2 DR Production scheduling

2.1 Problem formulation

Scheduling is part of a hierarchy of decisions involved in the operation of chemical processes (Figure 1). In the context of DR operation, the scheduling layer utilizes information concerning electricity prices, product demand, and any factors which may impact operating efficiency




Figure 1: Hierarchy of chemical process operation decisions, highlighting the interaction of scheduling, control and the physical process. u are the production rate setpoints/targets, w are the values of the state variables.

(such as ambient temperature) as an input, and determines the sequence of production rate targets/setpoints, u . The control layer translates these setpoint signals into control actions that are implemented in the process and the material storage system. Process state (or control) information, w , is provided to the control and scheduling layers, creating a feedback mechanism for control and rescheduling.

DR scheduling (from a demand-side perspective) consists of an optimization calculation that aims to minimize the operating cost associated with time-varying electricity prices, while maintaining process safety and service levels, the latter in terms of both product quality and meeting customer (product) demand. The fact that electricity prices change with typically hourly frequency and can thus be represented by a piecewise-constant signal, suggests the use of a slot-based scheduling formulation with fixed-length time slots. Given the different timescales involved in the DR scheduling of chemical processes (see Figure 1), it is required that the scheduling model embed dynamic information regarding the process and its control system [16, 17, 18]. The dynamic model of the process is typically discretized in time, with a discretization time step that is (much) shorter than the length of the scheduling time slot.

Thus, in order to capture the evolution of the process variables at the required resolution and over the time horizon of interest, we define three separate time grids for formulating and solving moving-horizon DR scheduling problems. As shown in Figure 2, the scheduling time-horizon of length N_D , is divided into days, represented by index d . This is motivated by the fact that electricity prices in the day-ahead market are typically known 24h in advance. Within each day, d , there exists a finer time-grid, referred to as the scheduling time grid, represented by N_I time slots of length T_I (denoted by index i). We specify that the length of each scheduling slot, i , coincide with the frequency that electricity prices change; later we will also impose that production rate changes can be made once per scheduling time slot. Within each scheduling time slot, we define a finer grid indicated by index j , where there are N_J time slots each of length T_J .




Figure 2: Representation of time for the DR scheduling problem, and the transient evolution of a hypothetical variable, y .

This grid is employed to represent the (discretized version) of the (higher-frequency/faster) process dynamics. The time step/sample time T_J is chosen based on the dominant time constant of the process [19, p. 319]. For convenience, we impose that T_I be an integer multiple of T_J , and that T_D be an integer multiple of T_I .

We will refer to the number of days considered in the scheduling model as the *scheduling window*, and the total number of days for which calculations are performed (on a moving horizon basis) as the *time horizon*.

Based on the above, we define the DR scheduling problem as follows. For clarity, we assume that the operating cost is influenced solely by the electricity prices (equivalently, that any other components of the operating cost are constant in time). Thus, the objective function of the problem represents a sum over the time horizon of the product of time-varying electricity prices (assumed to be constant in each scheduling time slot i), P_i , and the instantaneous power demand of the process at each time instant j in slot i , $P_{i,j}$. We assume that a process model is available, and that it represents the transient response of the process states, w , to changes in production targets, u , imposed in the scheduling calculation. It is assumed that the process operates under some form of feedback control and that the action of the control system is reflected in the process model. We also assume that a storage system is available and that it operates in tandem with the process to retain product generated in excess during periods of peak production/low energy prices. The process and the storage system are subject to measurable disturbances. Finally, we assume that process states and the production rate are subject to constraints, which may be variable in time and can be discretized according to the time grid described above. The DR scheduling problem thus takes the following form:

$$\begin{aligned}
 \min_{\bar{u}_i} J &= \sum_i \sum_j P_{i,j} P_i \\
 \text{s.t.} & \quad \text{Process model (differential-algebraic)} \\
 & \quad \text{Storage system model (differential-algebraic)} \\
 & \quad \text{Process constraints (algebraic)} \\
 & \quad \text{Demand satisfaction constraints (algebraic)}
 \end{aligned} \tag{1}$$




Figure 3: Block diagram of a scale-bridging model (SBM) for integration of scheduling and control. Adapted from Du et al. [26]

2.2 Process dynamics

Including an explicit representation of the dynamics of the process and storage system, along with the respective constraints, ensures that the derived DR schedule will be *dynamically feasible*, and its implementation will not lead to process or demand constraint violations. This is particularly important for DR scheduling, as it is expected that frequent changes in production targets will cause the process states to vary continuously in time. As we will discuss later, in the context of moving horizon DR scheduling, it is expected that the schedule will be recomputed frequently (either periodically or triggered by specific events). As a consequence, the dynamic representation of the process should be computationally efficient. First-principles models are far from meeting this requirement as they are typically large-scale, highly non-linear and ill-conditioned in practical cases.

In our previous work, we introduced the notion of a “scale-bridging model” (SBM) as a computationally efficient means of representing the process dynamics in DR scheduling calculations. SBMs bridge the time scales of the scheduling and control/process dynamics layers (Figure 1), and represent the closed-loop process dynamics between inputs u and states w (Figure 3). SBMs can be derived via system identification using either plant testing or routine historical operating data [20, 21, 8, 22] or by applying model reduction methods to (full-order, first principles) models [8, 23, 24, 25, 26].

The derivation of reduced-order SBMs is also guided by the observation (initially made empirically by Pattison et al. [8], and later confirmed using more rigorous machine-learning and data analysis approaches by Tsay and Baldea [27]), that only a small subset of the process state/output variables are relevant to DR scheduling. This result implies that the dimension of the (scale-bridging) dynamic process model used in scheduling calculations can be lower than the dimension of a rigorous first-principles model, with evident computational benefits.

The data-driven path is particularly attractive for deriving SBMs in practical situations. Two SBM structures are relevant in this case: finite step response (FSR) and Hammerstein-Wiener (HW). Both FSR and HW model forms expressed in this work are single-input single-output,



Figure 4: Hammerstein-Wiener (HW) model structure.

but multiple-input multiple-output structures are possible [28, 19]. In addition to being low-dimensional, these are linear or can be exactly linearized [22, 7], which presents the additional benefits of allowing the formulation of DR scheduling problems as mixed-integer linear programs (MILPs). In turn, this enables the use of powerful commercial solvers to derive solutions with well-defined optimality properties. We briefly describe the HW and FSR model structures below, where the model outputs are denoted with $w_{i,j}$, and inputs are denoted as u_i to reflect the fact that the evolution of the process variables are tracked over the finest time grid j while inputs only change once per time slot i .

Finite step response (FSR) models FSR models are advantageous for capturing complex process dynamics featuring (unknown) time delays [19]. In the case where setpoint changes are made once per scheduling time-slot and dynamics are sufficiently fast (i.e., for variables w that reach steady state within one scheduling slot), FSR models can be reduced to:

$$w_{i,j} = w_{i-1,N_j} + S_j(u_i - u_{i-1}) \quad (2)$$

The current state, $w_{i,j}$, depends on the final state of the previous scheduling time block (w_{i-1,N_j}), a model parameter, S_j , and the change in input between the current scheduling time slot (u_i) and the previous one (u_{i-1}).

Hammerstein-Wiener models Hammerstein-Wiener (HW) SBMs are used to capture variables with slower dynamics. A HW model comprises a linear dynamic model (here represented in state-space form), flanked by input ($H(u_i)$) and output ($W(y_{i,j})$) static nonlinear functions [28] (Figure 4). Similar to the FSR model, the input to the HW model is the target u_i , and the output is the value of the time-varying scheduling-relevant variable, $w_{i,j}$.

$$h_i = H(u_i) \quad (3a)$$

$$\mathbf{x}_{i,j+1} = A\mathbf{x}_{i,j} + Bh_i \quad (3b)$$

$$\mathbf{x}_{i,1} = \mathbf{x}_{i-1,N_j} \quad (3c)$$

$$y_{i,j} = C\mathbf{x}_{i,j} \quad (3d)$$

$$w_{i,j} = W(y_{i,j}) \quad (3e)$$

Exact linearization of HW models is possible when the static nonlinearities are piecewise linear as described in Kelley et al. [22].

2.3 Storage system dynamics

Storage system dynamics are modeled as an integrator with in/out flows, which will be discussed in more detail

later on as they apply to the case study.

3 Moving horizon DR scheduling

In principle, the DR schedule could be calculated once by solving a problem of the type outlined in (1), and implemented for the entire scheduling window. This “open loop” approach is optimal if the model is accurate, the predictions for the measured disturbances are perfect and no unmeasured disturbances occur. This is evidently not the case in practice, providing the impetus for developing a feedback mechanism where the schedule is updated as new information becomes available [8, 7].

In our study, we analyze two types of disturbances: periodic and non-periodic. Of the former category, we consider price changes and ambient temperature changes. While both exhibit periodicity over short time scales (in the order of days) relevant for DR scheduling, the accuracy of forecasts differs. Electricity prices are accurately known for a 24h time horizon, with new values becoming available at the end of the 24h interval. Weather predictions become more inaccurate as a longer time horizon is considered, but can be updated with an arbitrary frequency (a 6h update will be considered here). Changes in product demand are typically non-periodic. We will model them as square pulse signals of known magnitude but potentially unknown duration. Based on the disturbances under consideration, we define the following sets of moving horizon scheduling scenarios:

3.1 Electricity price updates

Three electricity price update scenarios are considered, denoted as problems **PP1-PP3**:

- **PP1** considers the situation where electricity prices are known with certainty for the entire scheduling time horizon and the DR scheduling problem is solved once for the full time horizon. This is in effect the “open loop” case where no moving horizon is used.
- In **PP2**, rescheduling occurs daily and a three-day scheduling window is used. Prices for the first day in the scheduling window are known with certainty, while the electricity prices for the two subsequent days are assumed equal to those in the first day.
- In **PP3**, electricity prices in the scheduling window are assumed to be known with certainty, but the problem is solved daily with a three-day moving scheduling window.

The schematic in Figure 5 illustrates these scenarios for a six-day time horizon. The selection of these scenarios can be explained as follows: **PP1** represents an ideal base case with perfect knowledge of the disturbance variable.




Figure 5: Electricity price update scenarios.

PP3 reformulates this scenario in a moving horizon context, and is thus expected to provide insights into the impact of using a moving horizon, feedback-based/closed-loop approach on the economics of the problem, while still assuming perfect knowledge of the exogenous variable. **PP2** represents a likely practical implementation of the moving horizon concept, whereby a forecast of the electricity prices will be necessary. The assumption made above, i.e. that prices for the day ahead can be used to forecast prices for the next two days, reflects a “minimal effort” approach to price forecasting, and can likely be improved by more detailed modeling of electricity price signals. Nevertheless, we believe that this simplistic choice provides useful insight on the impact of imperfect electricity price knowledge on DR scheduling.

3.2 Ambient temperature updates

Ambient temperature represents a measurable disturbance that affects the efficiency and/or operating conditions of many processes. It is assumed that its impact on process performance can be appropriately modeled (this is often dealt with by including a feed-forward component in the control system). The following scenarios are considered (we do not define a scenario **PT1**, rather, we use **PP1** as the reference problem where no temperature changes are taken into account):

- **PT2** is the case where ambient temperatures are assumed to be known accurately for the entire horizon, and the DR scheduling problem is solved once for the horizon
- **PT3** considers updates to the projected ambient temperature every 6 hours. A shrinking horizon is considered: the length of the scheduling window changes every time the scheduling problem is solved; 6 hours are subtracted from the window for each solve up until the end of the day, where another day is added to the scheduling window to account for an additional day of electricity prices becoming available.
- **PT4** is similar to **PT3**, but a constant length moving window is used.

Both **PT3** and **PT4** assume that temperatures beyond the first 6 hours in the scheduling window are at the same




Figure 6: Ambient temperature update scenarios.

level as the values at the corresponding time instants in the previous day. In solving **PT2-PT3**, we have assumed that electricity prices are known with certainty to better isolate the effect of temperature. The schematic in Figure 6 illustrates the scenarios discussed above.

3.3 Disturbances in the demand for the product of the chemical plant

Unlike the exogenous factors discussed above, demand changes cannot always be forecasted. As such, an event-driven rescheduling approach (rather than periodic rescheduling as above) is more appropriate. We consider several potential scenarios for demand disturbances. We relate these to events occurring at a customer facility located downstream. We will assume that disturbances are not persistent and they can be described as deviations from the nominal value of the demand in the form of square pulse signals of finite amplitude and finite (but unknown) duration, shown in Figure 7. Our treatment of the disturbances assumes some knowledge of the operation of the downstream customers (e.g. timing and/or duration of the maintenance event/outage).

- **PD2** considers *planned maintenance* at the downstream facility, where full knowledge of product demand level is available for the entire horizon, and the DR scheduling problem is solved once at the start of the horizon.
- **PD3** considers *unplanned maintenance* at the downstream facility, where the length of the demand disturbance is known once the disturbance begins, but the start time of the demand disturbance is unknown.




Figure 7: Demand profiles (**PD5**). The assumed end-points of the duration of the demand event are shown as equally spaced in time; this is not necessarily the case.

- **PD4** considers *random failure* at the downstream facility, where the start time of the demand disturbance and length of the demand disturbance are unknown. A conservative approach in this case is to assume that the demand disturbance is persistent, i.e., that it will continue until the end of the scheduling window. If/when the disturbance ends (or a *new* disturbance occurs), another rescheduling event is triggered. We note that this is similar to the manner in which the feedforward component of MPC is typically defined.

It is important to point out that such random failures are difficult to deal with in a cost-effective manner, and their occurrence may render the DR scheduling problem infeasible (examples of such circumstances are, i) exceeding the capacity of the storage facility if demand has dropped to a low level and production has reached its minimum bound, and ii) the complete depletion of stored material in case of demand increases - we note that both outcomes are possible *regardless* of whether moving horizon scheduling is used, but they represent an obstacle to obtaining a feasible solution to the DR scheduling problem).

Motivated by potential shortcomings of **PD4**, **PD5** deals with the end-time of random disturbances in a probabilistic way, as described below.

A probabilistic method of accounting for random failure (PD5) In dealing with random failures, we will assume that the magnitude of the corresponding disturbance is known, while its length is considered a random variable. When a demand disturbance occurs due to random failure, a set of N_r possible demand vectors

are generated (Figure 7), where the end time of the demand disturbance (t_{end}) is different for each vector and is drawn from a uniform distribution: $t_{end} \sim \mathcal{U}[t_{event}, T_m]$ bounded between the start of the demand disturbance (t_{event}) and the end of the time horizon, T_m .

We then reformulate the demand satisfaction constraint:

$$F_{i,j}^p - D_i = f_{s_{i,j}}^{in} - f_{s_{i,j}}^{out} \quad (4)$$

as a chance constraint in (5). $F_{i,j}^p$ represents the production rate at time j in scheduling slot i , D_i is the product demand in slot i , and $f_{s,i,j}^{in}$ and $f_{s,i,j}^{out}$ are, respectively, the flow rates of material entering and leaving the storage system.

Chance constraints state that the *probability* of meeting a constraint is above a specified tolerance, α . In a chance-constrained implementation, (4) can be transformed into an inequality, with a binary indicator variable, z_r , being used register when the constraint is met for a particular demand vector, $D_{i,r}$. The constant, M , is large enough such that if the constraint is met, $z_r = 1$, and $z_r = 0$ otherwise.

$$F_{i,j}^p - D_{i,r} \geq f_{s_{i,j}}^{in} - f_{s_{i,j}}^{out} - M(1 - z_r) \quad \forall r \quad (5)$$

The use of the binary indicator variable enables calculation of the probability that the demand constraint is met, and allows for imposing that this probability be above the tolerance, α , $0 \leq \alpha \leq 1$.

$$\sum_r z_r \geq \alpha N_r \quad (6)$$

Remark. The transformation of equality constraint (4) to an inequality constraint is valid and will not impact the solution of the scheduling problem. As we will show below, we are solving a (linear) cost minimization problem, and the optimal solution is to not produce any excess product (i.e. product that is neither sent to the storage system ($f_{s,i,j}^{\text{in}}$) nor used to meet customer demand ($f_{s,i,j}^{\text{out}}$)).

4 ASU case study

We utilize a single-column air separation unit (ASU) to demonstrate our moving-horizon scheduling method. Cryogenic air separation is an electricity intensive process in the industrial gas sector, which accounted for 2.62% of the yearly industrial electricity consumption in the US in 2014 [29]. The products of air separation (nitrogen, oxygen, and argon) are utility streams in various industries such as steel production and microelectronics. Cryogenic ASUs are attractive for DR participation because their only feedstocks are air and electricity. Furthermore, the liquefied cryogenic products can be stored safely and efficiently.




Figure 8: Diagram showing the single-column ASU model based on Johansson [21].

4.1 Process model

The single-column ASU has a single product: high-purity nitrogen gas (Figure 8). For a full description of the model, we refer the reader to the works of Johansson [21] and Pattison et al. [8]. Inlet feed air is compressed to 6.8 bar before entering the primary multistream heat exchanger (PHX), where it condenses. The condensed air is then fed to the cryogenic column at the bottom, where it is distilled. A high-purity nitrogen gas stream is obtained from the top. The product nitrogen stream passes through the PHX as a refrigeration stream. Beyond the PHX, the product stream is sent to customers to meet demand, and any excess is sent to storage.

In the storage system, gaseous N_2 produced in excess of demand, $f_{s_{i,j}}^{in}$, is liquefied in a liquefier. When production is lower than demand, product is removed from the storage tank, at a rate $f_{s_{i,j}}^{out}$, to meet demand. A material balance equation is used to model the storage holdup:

$$s_{i,j+1} = (f_{s_{i,j}}^{in} - f_{s_{i,j}}^{out})\Delta j + s_{i,j} \quad (7a)$$

$$f_{s_{i,j}}^{in} \geq 0 \quad (7b)$$

$$f_{s_{i,j}}^{out} \geq 0 \quad (7c)$$

The following continuity conditions are imposed:

$$f_{s_{i,1}}^{in} = f_{s_{i-1}, N_i}^{in} \quad (8a)$$

$$f_{s_{i,1}}^{out} = f_{s_{i-1},N_i}^{out} \quad (8b)$$

$$s_{i,1} = s_{i-1,N_j} \quad (8c)$$

along with the requirement that the total storage holdup is less than 50 kmol:

$$0 \leq s_{i,j} \leq 50\text{kmol} \quad (9)$$

and the storage holdup endpoint condition, where the piecewise constant parameter, μ , is dependent on the

time (t) of day the end of the scheduling window occurs at (11).

$$\mu s_{N_i, N_j} \geq s_{1,1} \quad (10)$$

$$\mu(t) = \begin{cases} 0.5 & t \leq 6 \\ 0.75 & 6 < t \leq 12 \\ 0.5 & 12 < t \leq 18 \\ 1 & 18 < t \leq 24 \end{cases} \quad (11)$$

The liquefier (ℓ) is modeled after an ideal cycle, where N_2 is isothermally compressed and then expanded in an isentropic turbine. The liquefier is assumed to have an efficiency that fluctuates based on the ambient temperature (T_∞), regardless of the flow through the unit. The linear relationship below assumes the nominal liquefier efficiency of 80% that fluctuates $\pm 5\%$ over the range of historical temperatures studied from Fresno, CA in 2017 [30].

$$\eta_\ell(T_\infty) = 0.0055T_\infty + 0.6955 \quad (12)$$

The liquefier and the compressor (C) drive the electricity consumption of the process, the two turbines ($t1, t2$) partially offset this electricity consumption, but the impact is minimal. In calculating the power consumption of each unit, we assume a linear relationship, where the flow through each unit is multiplied by the polytropic head of the unit (Ω) to calculate the net work ($W_{i,j}$) of each unit.

$$\begin{aligned} W_{i,j}^C &= \Omega_C F_{i,j}^f \\ W_{i,j}^{t1} &= \Omega_{t1} F_{i,j}^f \\ W_{i,j}^{t2} &= \Omega_{t2} F_{i,j}^p \\ W_{i,j}^\ell &= \Omega_\ell f_{s_{i,j}}^{in} \end{aligned} \quad (13)$$

The variable $F_{i,j}^f$ is the plant air feed flowrate, which becomes an optimization variable for the operation of the process, as it dictates power demand. The superscripts C , $t1$, $t2$, and ℓ represent the compressor, turbine 1, turbine 2, and the liquefier, respectively. The time-varying power requirement for the plant is given below, and is taken to be the sum of the net work (W) for all relevant units described above.

$$\mathcal{P}_{i,j} = W_{i,j}^C + W_{i,j}^{t1} + W_{i,j}^{t2} + W_{i,j}^\ell \quad (14)$$

Power calculated in the above equation is used in the objective function (1) of our DR scheduling problem, which minimizes operating cost.

4.2 Scale-Bridging Models for the ASU

We utilize SBMs to represent scheduling-relevant variables, which are near their bounds during normal operation or are included in the objective function [8].

- **Production Rate and Feed Flowrate:** The production rate ($F_{i,j}^p$) and feed flowrate ($F_{i,j}^f$) determine the electricity consumption of the compressor and liquefier. As such, both variables directly affect the objective function 1. The production rate is also bounded: $16 \leq F_{i,j}^p \leq 24$, reflecting that it can vary $\pm 20\%$ from its nominal level of 20 mol/s.
- **Impurity Levels in the Nitrogen Product:** The maximum allowed value for the product impurity concentration is $I_{max} = 2000 \text{ ppm}$, where ppm is defined by the mole fraction of O_2 . Impurity level, $I_{i,j}^p$, is a *critical* quality variable, and therefore, is bounded using a “backoff” (i.e. more restrictive) constraint, i.e., $I_{i,j}^p \leq 1800 \text{ ppm}$.
- **Column Reboiler Level:** The reboiler level ($M_{i,j}^r$) is bounded such that $0 \leq M_{i,j}^r \leq 120 \text{ kmol}$ and such that the holdup at the end of the time horizon is greater than or equal to the holdup at the beginning of the time horizon ($M_{1,1}^r \leq M_{N_I, N_J}^r$) to prevent the depletion of liquid (and an artificial drop in power requirements).
- **Column Flooding:** We use the flooding fraction ($\delta_{i,j}^f$) to quantify column flooding. Based on Sinnott [31] and Johansson [21], we set the following bound: $\delta_{i,j}^f \leq 0.97$, and note that detailed velocity equations are given by Johansson [21].
- **Phase of the Air Stream:** We have constrained the ratio between the outlet pressure from zone 1 of the PHX and the dew pressure of the air outlet; $P_{i,j}^{zone1}/P_{i,j}^d \leq 0.96$ at the relevant temperature. This constraint ensures the streams in zone 1 of the PHX are in the gas phase.
- **Minimum Temperature Driving Force:** Imposing a minimum temperature driving force across the reboiler-condenser ensures that heat exchange occurs in the correct direction. Therefore, the temperature difference ($\Delta T_{i,j}$) is bounded such that it is less than the minimum approach temperature of $1.9^\circ C$.

The variables relevant to the constraints above were modeled using a combination of HW (3) and FSR (2) models, with one SBM for each scheduling-relevant variable. Closed-loop plant data for model identification were obtained by simulating a full-order first principles ASU model described in detail in Johansson [21]. The SBM used for each scheduling-relevant variable is given below (where the input for each model is the production rate setpoint, u_i):

- HW Models
 - Impurity, $I_{i,j}^p$
 - Reboiler Level, $M_{i,j}^r$




Figure 9: Price sets used (bottom) from [34] from October 16-21, 2017 and temperature sets used (top) from [30] for 16-21 October, 2017.

- Column Flooding, $\delta_{i,j}^f$
- Air Stream Phase, $P_{i,j}^d$
- Minimum Temperature Driving Force, $\Delta T_{i,j}$
- FSR Models
 - Production Rate, $F_{i,j}^p$
 - Feed Flowrate, $F_{i,j}^f$

The storage system model in (7) is linear and was used as-is.

5 Results

All problems were coded in GAMS 25.1.3 [32] and solved using CPLEX 12.8.0 [33] to a 0.1% optimality gap on a 64-bit Windows system with Intel Core-7-2600 CPU at 3.40 GHz and 16 Gb RAM.

Two six-day electricity price vectors are considered (Figure 9). In price set 1, the price profile for the first 24 hours is repeated in the subsequent five days, resulting in a moderate set of prices which are truly periodic. Price set 2 reflects historical day-ahead market prices from CAISO [34], which have high variation between days. We also use two six-day ambient temperature profiles (Figure 9). Temperature set 1 is structured in a similar way to price set 1: the temperature for the first day is replicated in the five subsequent days. Temperature set 2 comprises historical temperature data [30] corresponding to the six days represented in price set 2. We consider two different demand vectors in conjunction with price set 1 and price set 2. Both demand vectors feature two events/disturbances (Figure 10) to test the effects of a sudden increase/decrease in demand at different points in the time horizon.




Figure 10: Plots of demand vector 1 (bottom) and demand vector 2 (top) .

Table 1: Results for **PP1-PP3** for price set 1.

Problem	Operating cost (\$)	vs. PT0	Energy (MJ)	vs. PT0
PP0	1902.93	–	43.26	–
PP1	1898.26	-0.25%	43.90	+1.48%
PP2	1898.77	-0.22%	43.90	+1.48%
PP3	1898.77	-0.22%	43.90	+1.48%

5.1 MH scheduling for periodic changes in electricity prices

Price set 1 The results are shown in Table 1 and Figure 11. The solutions of **PP2** and **PP3** are equivalent, and the optimal schedules for **PP1-PP3** are very similar, exhibiting near-periodic operation. This is also evident in the maximum difference in inventory (2.2.kmol) between the solutions for **PP1** and **PP2-PP3** (shown in Figure 11). These results are to be expected given that all problems effectively have perfect price knowledge (owing to the repetitive nature of the price profile). The cost for **PP2-PP3** is slightly higher than **PP1** due to the constraint on inventory that is enforced at the end of each scheduling window (with the window being shorter for the former scenarios). Conversely, the longer time horizon in **PP1** makes the most effective use of stored material. The results are compared with a simulation case **PP0**, where the plant operates at a constant production rate but is still subject to the same price fluctuations. DR operation achieves some modicum of cost savings, but with the disadvantage of higher overall energy consumption.

Price set 2 Significant variation between the three MH scheduling approaches (Figure 12) was seen for price set 2. The operating cost for **PP2** has the highest value at the optimum due to imperfect price knowledge beyond the first day (Table 2). Inventory levels for **PP2** are more conservative than in **PP1** and **PP3**, with **PP1** exhibiting the lowest overall inventory levels as a result of considering the longest time horizon. The objective func-




Figure 11: Optimal schedule of production rate setpoints (top), storage holdup (middle), and the difference in stored inventory (**PP1** vs **PP3**) for price set 1.

Table 2: Results **PP1-PP3** for price set 2.

Problem	Operating cost (\$)	vs. PP0	Energy (MJ)	vs. PP0
PP0	1772.86	—	43.26	—
PP1	1754.93	-1.01%	44.17	+2.10%
PP2	1758.98	-0.79%	44.15	+2.06%
PP3	1754.99	-1.01%	44.20	+2.17%

tion values for **PP1** and **PP3** are very similar, demonstrating that MH scheduling remains an effective way of reducing production cost (assuming that accurate price predictions are available). The energy consumption for all three scenarios is similar. Overall, we note that all MH schemes achieve a cost reduction relative to case **PP0** (operating at constant production rate), but that energy use increases due to liquefier usage.

5.2 MH scheduling for periodic chagnes in ambient temperature

Price/temperature set 1 For the six-day time horizon, the objective function values at the optimum for **PT2-PT4** are very similar (Table 3). However, the scheduled production rate target and inventory levels are noticeably different for the three scenarios (Figure 13), particularly in the timing and value of the peaks in the production setpoint. Interestingly, the periodic behavior seen in Figure 11 is not observed to the same extent in the solution of **PT2-PT4**, even though the operating cost of **PP1-PP3** with price set 1 is nearly the same as that of **PT2-PT4** with price and temperature set 1. We ascribe this to solution degeneracy at the tolerance level specified.

Figure 13 reveals that inventories for scenario **PT3** are in general higher than **PT2** and **PT4**. This reflects a similar lack of flexibility as in **PP2** since the scheduling window shrinks at every schedule update (up until a new




Figure 12: Optimal schedule of production rate setpoints (top) and storage holdup (bottom) for price set 2.

Table 3: **PT2-PT4** results for price/temperature set 1.

Problem	Operating cost (\$)	vs. PT0	Energy (MJ)	vs. PT0
PT0	1902.93	—	43.26	—
PT2	1898.62	-0.23%	43.87	1.41%
PT3	1898.63	-0.23%	43.89	1.46%
PT4	1898.71	-0.22%	43.91	1.50%

day starts), while the endpoint constraint on inventory remains unchanged until the horizon is shifted in time. In contrast, the storage constraint in **PT1** is only enforced at the end of the time horizon and in **PT4** the storage constraint is updated at every rescheduling operation.

Price/temperature set 2 cause slightly more variation in the results of **PT2-PT4** (Table 4). As before, the optimal schedules are different (Figure 14), particularly in the timing and height of the peaks in the production rate set point. The inventories for **PT2-PT4** are less disparate than in the case of price/temperature set 1, potentially a consequence of a higher variability in electricity prices in price set 2 (Figure 9).

Table 4: **PT2-PT4** results for price/temperature set 2.

Problem	Operating cost (\$)	vs. PT0	Energy (MJ)	vs. PT0
PT0	1772.86	—	43.26	—
PT2	1754.62	-1.03 %	44.17	2.10%
PT3	1755.27	-0.99 %	44.18	2.13%
PT4	1755.37	-0.99 %	44.23	2.24%




Figure 13: Optimal schedule of production rate setpoints (top) and storage holdup (bottom) for price and temperature set 1




Figure 14: Optimal schedule of production rate setpoints (top) and storage holdup (bottom) for price and temperature set 2.

5.3 MH scheduling for product demand disturbances

PD2 assumes perfect knowledge of product demand and electricity prices for the entire horizon, leading to the lowest operating cost (Table 5). Energy demand is quite similar for all scenarios. The production schedules (Figure 15) further emphasize the benefit of this perfect knowledge, showing that the production rate and inventory levels for **PD2** are quite different from the other scenarios. **PD5** and **PD3** follow similar trajectories since both scenarios assume knowledge of the end-point of the pulse demand disturbances. In scenario **PD3**, inventory remains low post-demand disturbance. This is due to the endpoint storage constraint (10), which states that the holdup must be 50% of its level at the start of the rescheduling event that occurs at $t=72h$.

Results for **PD2-PD5** for price set 2 and demand

Table 5: Objective function values for **PD2-PD5** for price/demand set 1.

Problem	Operating cost (\$)	Energy (MJ)	Time horizon (days)
PD2	1905.00	43.75	6
PD3	1721.94	40.00	5.54
PD4	1913.66	43.77	6
PD5	1921.16	43.88	6




Figure 15: Optimal schedule of production rate setpoints (top) and storage holdup (bottom) for price set 1 and demand vector 1.

vector 1 show similar trends, with **PD2** preserving its low operating cost advantage as expected. Complete results and plots are not included in the interest of saving space.

Demand vector 2 The results here (Table 6 and Figure 16) are similar to the case of demand vector 1 in terms of operating cost, with **PD2** having the lowest operating cost.

The results for **PD2-PD5** with price/demand vector 2 are given in Table 7. In this case, the final rescheduling calculation (at the end of the second demand disturbance at $t=72h$) for **PD4** was infeasible. This can be interpreted as follows: the definition of **PD4** assumes that the disturbance will persist until the end of the time horizon, and as a consequence the solution of this problem dictates that the storage system be fully depleted. When the disturbance does in fact end, and product de-

Table 6: Results for **PD2-PD5** for price set 1 and demand vector 2.

Problem	Operating cost (\$)	Energy (MJ)	Time horizon (days)
PD2	1896.67	43.68	6
PD3	1713.36	40.00	5.54
PD4	1912.49	43.80	6
PD5	1944.19	44.36	6




Figure 16: Optimal schedule of production rate setpoints (top) and storage holdup (bottom) for price set 1 and demand vector 2.

Table 7: **PD2-PD5** results for price/demand vector 2.

Problem	Operating cost (\$)	Energy MJ	Time horizon (days)
PD2	1752.82	43.84	6
PD3	1612.32	40.20	5.54
PD4	Infeasible	—	6
PD5	1798.92	44.63	6

mand is restored (increased) to the nominal level, there is no stored product available to meet demand while the plant production is ramped up.

Conversely, the solution of **PD5**, which uses chance constraints to estimate the time instant when the demand disturbance ends, does reserve some storage holdup for the eventual increase in demand and the problem is feasible across all scheduling windows.

6 Discussion and conclusions

Our empirical results reveal that accounting for fluctuations in operating circumstances (electricity prices, ambient conditions) in production scheduling can yield economic benefits compared to steady-state operation, the latter being the accepted (and likely preferred) approach of plant operators. These benefits are reflected in the operating cost of the plant, rather than in a reduction in overall energy use. While the power demand is reduced during peak hours (an important benefit for the power grid), overall energy use increases during DR operation since material must be liquefied to be stored.

For the different scenarios considered, the economics of using a moving horizon rescheduling strategy did not differ significantly from the best case of a scheduling calculation based on perfect knowledge of all disturbances and considering the entire time horizon. This observation is true in the case where moving horizon scheduling itself




Figure 17: Optimal schedule of production rate setpoints (top) and storage holdup (bottom) for price set 2 and demand vector 2.

possesses perfect knowledge of the disturbances for the scheduling window considered. However, even simplistic forecasting strategies for periodic disturbances (e.g., using information from the preceding day) yield economic gains. This finding should provide sufficient impetus for the adoption of such strategies in industry. Our results also indicate that, while the *economic* outcomes of different rescheduling strategies are not significantly different, the *production schedules* may be quite dissimilar. This finding is important in that it suggests that different strategies may lead to, e.g., different patterns of equipment use and wear and tear/fatigue, a factor that was not explicitly considered in this study but should be taken into account in future work and in practical implementations.

Interestingly, even when both disturbances are perfectly cyclical (i.e., profiles from the first day are repeated in all subsequent days), the optimal production schedules do not exhibit a perfectly cyclical pattern. We ascribe this to the fact that the time horizons considered here were finite and relatively short (at most six times the “period of oscillation” for price and temperature disturbances). As practical implementations cannot consider infinite horizons, we posit that these results reveal the need for using a rigorous, optimization-based rescheduling strategy for engaging in demand response initiatives (rather than relying solely on human operators, who would likely follow a periodic operation strategy).

Finally, we note that dealing with disturbances on the chemical product demand side is much more difficult than dealing with changes in electricity prices and ambient conditions. Demand changes are in many cases random in magnitude, time of occurrence, and duration. Assuming (in typical MPC fashion) that a disturbance will persist until the end of the prediction window can render scheduling calculations infeasible. Our approach

here was to assume that the disturbance will *eventually* end at some time point *within* the scheduling window. Of course, this assumption may be violated in practice, where such disturbances may persist beyond the end of the scheduling window or may in effect materialize in frequent (e.g., hourly or faster) changes in product demand. Further investigation in dealing with such circumstances is warranted, and we posit that the need to deal with fluctuations in chemical product demand may severely limit the potential of an electricity-intensive chemical plant to engage in demand response or other initiatives aimed at supporting the operation of the power grid. The ASU herein was used to demonstrate how a chemical process might react to DR operation. We posit that the concepts and outcomes could be extended to other chemical processes operating under DR.

7 Acknowledgments

This material is based upon work supported by the Department of Energy under Award Number(s) DE-OE0000841. Partial financial support from the National Science Foundation (NSF) through the CAREER Award 1454433 and Award CBET-1512379 is acknowledged with gratitude. M.T.K. was supported through the Department of Energy Computational Science Graduate Fellowship through grant number DE-FG02-97ER25308.

References

- [1] U. EPA. EPA: Overview of Greenhouse Gases, 2018. URL <https://www.epa.gov/ghgemissions/overview-greenhouse-gases>.
- [2] S. Greenlee and A. Gonzales. Summer Loads and Resources Assessment. Technical report, California Independent System Operator (CAISO), Folsom, CA, 2018. URL <http://www.caiso.com/Documents/2018SummerLoadsandResourcesAssessment.pdf>.
- [3] US EIA. Monthly Energy Review March 2019, 2019. URL <https://www.eia.gov/totalenergy/data/monthly/pdf/mer.pdf>.
- [4] US EIA. Energy use in Industry, 2019. URL https://www.eia.gov/energyexplained/index.php?page=us_{_}energy_{_}industry{#\}tab1.
- [5] P. Schäfer, H. G. Westerholt, A. M. Schweidtmann, S. Ilieva, and A. Mitsos. Model-based bidding strategies on the primary balancing market for energy-intense processes. *Computers & Chemical Engineering*, 120:4–14, jan 2019. ISSN 00981354. doi: 10.1016/j.compchemeng.2018.09.026. URL <https://linkinghub.elsevier.com/retrieve/pii/S009813541830749X>.
- [6] P. Schäfer, L. F. Bering, A. Caspari, A. Mhamdi, and A. Mitsos. Nonlinear Dynamic Optimization for Improved Load-Shifting Agility of Cryogenic Air Separation Plants. *Computer Aided Chemical Engineering*, 44:547–552, 2018. doi: 10.1016/B978-0-444-64241-7.50086-0. URL <https://linkinghub.elsevier.com/retrieve/pii/B9780444642417500860>.
- [7] M. T. Kelley, R. C. Pattison, R. Baldick, and M. Baldea. An MILP framework for optimizing demand response operation of air separation units. *Applied Energy*, 222:951–966, jul 2018. ISSN 03062619. doi: 10.1016/j.apenergy.2017.12.127. URL <http://linkinghub.elsevier.com/retrieve/pii/S0306261917318524>.
- [8] R. C. Pattison, C. R. Touretzky, T. Johansson, I. Harjunkoski, and M. Baldea. Optimal Process Operations in Fast-Changing Electricity Markets: Framework for Scheduling with Low-Order Dynamic Models and an Air Separation Application. *Industrial & Engineering Chemistry Research*, 55(16):4562–4584, apr 2016. ISSN 0888-5885. doi: 10.1021/acs.iecr.5b03499.
- [9] A. Allman and Q. Zhang. Distributed cooperative industrial demand response. *Journal of Process Control*, 86:81–93, 2020. ISSN 09591524. doi: 10.1016/j.jprocont.2019.12.011. URL <https://doi.org/10.1016/j.jprocont.2019.12.011>.
- [10] D. Gupta and C. T. Maravelias. Framework for studying online production scheduling under endogenous uncertainty. *Computers and Chemical Engineering*, 135:106670, 2020. ISSN 0098-1354. doi: 10.1016/j.compchemeng.2019.106670. URL <https://doi.org/10.1016/j.compchemeng.2019.106670>.
- [11] S. Shyamal and C. L. Swartz. Real-time energy management for electric arc furnace operation. *Journal of Process Control*, 74:50–62, 2019. ISSN 09591524. doi: 10.1016/j.jprocont.2018.03.002. URL <https://doi.org/10.1016/j.jprocont.2018.03.002>.
- [12] Y. He and M. Petit. Demand response scheduling to support distribution networks operation using rolling multi-period optimization. *Journal of Process Control*, 74:13–22, 2019. ISSN 0959-1524. doi: 10.1016/j.jprocont.2018.07.005. URL <https://doi.org/10.1016/j.jprocont.2018.07.005>.
- [13] J. G. Coelho, V. B. D. Oliveira, and J. L. Félix. Optimal scheduling of a by-product gas supply system in the iron and steel-making process under uncertainties. *Computers and Chemical Engineering*, 125:351–364, 2019. ISSN 0098-1354. doi: 10.1016/j.compchemeng.2019.01.025. URL <https://doi.org/10.1016/j.compchemeng.2019.01.025>.
- [14] P. Mathur, C. L. E. Swartz, D. Zyngier, and F. Welt. Uncertainty management via online scheduling for optimal short-term operation of cascaded hydropower systems. *Computers and Chemical Engineering*, 134:106677, 2020. ISSN 0098-1354. doi: 10.1016/j.compchemeng.2019.106677. URL <https://doi.org/10.1016/j.compchemeng.2019.106677>.
- [15] R. C. Pattison, C. R. Touretzky, I. Harjunkoski, and M. Baldea. Moving horizon closed-loop production scheduling using dynamic process models. *AIChE Journal*, 63(2):639–651, feb 2017. ISSN 00011541. doi: 10.1002/aic.15408.
- [16] C. T. Maravelias and C. Sung. Integration of production planning and scheduling: Overview, challenges and opportunities. *Computers and Chemical Engineering*, 33(12):1919–1930, 2009. ISSN 00981354. doi: 10.1016/j.compchemeng.2009.06.007.
- [17] I. Grossmann. Enterprise-wide optimization: A new frontier in process systems engineering. *AIChE Journal*

- nal*, 51(7):1846–1857, jul 2005. ISSN 0001-1541. doi: 10.1002/aic.10617.
- [18] M. Baldea and I. Harjunoski. Integrated production scheduling and process control: A systematic review. *Computers & Chemical Engineering*, 71:377–390, dec 2014. ISSN 00981354. doi: 10.1016/j.compchemeng.2014.09.002.
 - [19] D. E. Seborg, D. A. Mellichamp, T. F. Edgar, and F. J. Doyle. *Process Dynamics and Control*. John Wiley & Sons, Inc., Hoboken, NJ, USA, 2010. ISBN 9780470128671. doi: 10.1007/s13398-014-0173-7.2.
 - [20] C. Tsay, M. Baldea, J. Shi, A. Kumar, and J. Flores-Cerrillo. Data-Driven Models and Algorithms for Demand Response Scheduling of Air Separation Units. In *Process Systems Engineering (PSE)*, pages 1273–1278. San Diego, CA, 2018. doi: 10.1016/B978-0-444-64241-7.50207-X. URL <https://linkinghub.elsevier.com/retrieve/pii/B978044464241750207X>.
 - [21] T. Johansson. *Integrated Scheduling and control of Air Separation Unit Subject to Time-Varying Electricity Price*. PhD thesis, KTH Royal Institute of Technology, Stockholm, Sweden, 2015.
 - [22] M. T. Kelley, R. C. Pattison, R. Baldick, and M. Baldea. An efficient MILP framework for integrating nonlinear process dynamics and control in optimal production scheduling calculations. *Computers & Chemical Engineering*, 110:35–52, feb 2018. ISSN 00981354. doi: 10.1016/j.compchemeng.2017.11.021. URL <http://linkinghub.elsevier.com/retrieve/pii/S0098135417304222>.
 - [23] M. Baldea and P. Daoutidis. *Dynamics and Nonlinear Control of Integrated Process Systems*. Cambridge University Press, New York City, 1st edition, 2012. ISBN 978-0-521-19170-8.
 - [24] S. S. Jogwar, M. Baldea, and P. Daoutidis. Dynamics and Control of Process Networks with Large Energy Recycle. *Industrial & Engineering Chemistry Research*, 48(13):6087–6097, jul 2009. ISSN 0888-5885. doi: 10.1021/ie801050b.
 - [25] C. R. Touretzky and M. Baldea. Integrating scheduling and control for economic MPC of buildings with energy storage. *Journal of Process Control*, 24(8):1292–1300, 2014. ISSN 09591524. doi: 10.1016/j.jprocont.2014.04.015.
 - [26] J. Du, J. Park, I. Harjunoski, and M. Baldea. A time scale-bridging approach for integrating production scheduling and process control. *Computers & Chemical Engineering*, 79:59–69, aug 2015. ISSN 00981354. doi: 10.1016/j.compchemeng.2015.04.026.
 - [27] C. Tsay and M. Baldea. Integrating production scheduling and process control using latent variable dynamic models. *Control Engineering Practice*, 94:104201, jan 2020. ISSN 09670661. doi: 10.1016/j.conengprac.2019.104201. URL <https://linkinghub.elsevier.com/retrieve/pii/S0967066119301789>.
 - [28] S. A. Billings. Nonlinear system identification : NARMAX methods in the time, frequency, and spatio-temporal domains. 2013.
 - [29] US EIA. Manufacturing Energy Consumption Survey 2014, 2017. URL <https://www.eia.gov/consumption/manufacturing/data/2014/>.
 - [30] Fresno Yosemite International Airport, CA, 2017. URL <https://www.wunderground.com/history/daily/us/ca/fresno/KFAT/date/2017-10-20>.
 - [31] R. K. Sinnott. Chemical Engineering Design. *ELSEVIER - Coulson & Richardson’s Chemical Engineering series*, 6(4):440–445, 2005.
 - [32] GAMS. General Algebraic Modeling System (GAMS) Release 25.1.3, 2018.
 - [33] IBM. CPLEX 12.8.0, 2019.
 - [34] CAISO. California Independent System Operator, 2017. URL <http://www.caiso.com/Pages/default.aspx>.