CLOSED-LOOP SCHEDULING WITH PROCESS FAULTS: FRAMWORK AND AN AIR SEPARATION UNIT EXAMPLE

Richard Pattison and Michael Baldea
McKetta Dept. of Chemical Engineering, The University of Texas at Austin
Austin, TX 78712

Abstract

The strong increase of the contribution of renewable sources (wind, solar) to the power generation portfolio has increased uncertainty in the operation of the grid and has motivated efforts in demand-side management. Of particular interest is the demand response (DR) operation of large industrial electricity consumers, which is encouraged via (real-)time-varying electricity pricing schemes. For industrial users, taking advantage of such price structures calls for far more frequent changes in the production schedule than ever before, and requires a close coordination between scheduling decisions, the actions of the control system, and the dynamic capabilities of a plant. Crucial to this end are, i) embedding representations of the process dynamics in the scheduling calculations and, ii) developing a framework for recomputing the scheduling decisions upon the occurrence of process events and faults. In this work, we address the latter challenge by proposing a novel closed-loop, moving-horizon approach to production scheduling, whereby process dynamics are represented using previously-developed scale-bridging models, and schedule updates are triggered by both scheduling-relevant market events and by process events that do not have an explicit impact on the schedule execution. We use the DR operation of an industrial-scale air separation unit model to demonstrate the successful implementation of these concepts.

Keywords

Integrated scheduling and control, Air separation, Closed-loop scheduling, Demand response, Fault detection

Introduction

In the modern economy, market conditions can be highly variable and can change frequently. For example, the electricity market is subject to high variability on both the supply side (owing to the rapidly growing incorporation of renewables like solar photovoltaics and wind generation), and the demand side, where variability –particularly in the residential sector- causes significant swings between peak and off-peak grid loads. The need to balance power supply and demand under such transient conditions has motivated grid operators to initiate demand-side management programs, whereby users are incentivized to reduce their peak time consumption via, e.g., time-varying electricity price structures, ranging from dual peak/off-peak tariffs to pricing that follows the real-time deregulated market (Miller et al., 2008; Zhang et al., 2015).

In the case of energy-intensive industrial users, of which air separation is a prototypical example, demand response (DR) operation (and taking advantage of time-dependent tariffs) entails increasing production beyond the product demand rate during off-peak hours and storing products, followed by reducing the production rate and using stored products to meet demand at peak times. In this context, optimizing the production schedules to maximize DR benefits typically entails making frequent changes in production rate and/or product grade, over time horizons that overlap with the time scale of the dynamic response of the plant. Under these circumstances, explicitly accounting for process dynamics and dynamic constraints (along with the behavior of the process control system) in DR scheduling calculations becomes an imperative necessity (Baldea and Harjunkoski, 2014).

The vast majority of current approaches to production scheduling (loosely defined as the selection of a sequence of production rate targets and product types over a future time horizon, typically spanning a few days to a few weeks) make use of steady-state process representations. These models statically correlate the operating level of the production system with a set of economic indicators (production rate, energy use, etc.), and capture process dynamics in the form of tabulated transition times between a (finite) set of system states. A natural approach to explicitly accounting for process dynamics entails using a (first-principles) dynamic model of the process and its control system in the scheduling problem formulation. However, capturing the entire spectrum of process dynamics, as relevant to both product quality and process and safety constraints, calls for the use of intricate, detailed dynamic models which are inevitably nonlinear and high-dimensional, rendering the integrated scheduling/control problem impossible to solve in a practical amount of time.

In order to alleviate this difficulty, in our previous work we introduced the concept of scale-bridging model –SBM-

(Du et al., 2015) as a low-dimensional representation of the scheduling-relevant dynamics of a process and its control system. Initially derived from first-principles arguments (Du et al., 2015), we later showed that SBMs can be successfully obtained from historical data collected during routine process operations (Pattison et al., 2016a). Subsequently, we extended these ideas in the development of a moving-horizon scheduling framework that allows the schedule to be recomputed when new *scheduling-relevant information* (e.g., updated product demand or energy price forecasts) becomes available, (Pattison et al., 2016b).

In this work, we consider another crucial aspect of closed-loop scheduling, that is, dealing with *process-level* faults and disturbances, which we define as non-critical, detectable events (i.e., with no safety implications, and not requiring plant shut-down) that occur at the process level and do not have an explicit connection with market conditions as was the case above. Rather, their impact on process economics is manifest in limiting the plant's ability to meet a subset of production rates or product grades originally part of the product wheel.

We begin with an overview of the technical background on scheduling under dynamic constraints, followed by introducing the proposed framework. We then illustrate the results with a case study where we consider the DR operation of an air separation unit subject to faults.

Background: scheduling under dynamic constraints

We focus on a class of continuous chemical processes having (for simplicity) a single product stream and product storage capability. The process and the hierarchy of operational decision making are depicted in Figure 1.
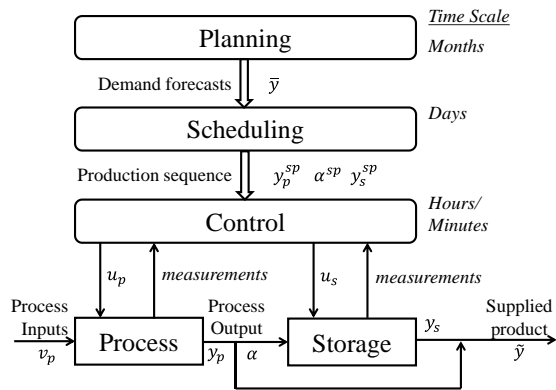

Figure 1. Hierarchy of operational decision making for a single product process with product storage capacity.

The planning layer establishes the long term demand forecasts ($\bar{y}$) based on contractual agreements and market dynamics, while the scheduling layer uses the demand forecasts as well as feedstock availability and price forecasts to determine an economically optimal production target sequence for the process (y_p^{sp}) over a daily to weekly horizon, as well as the utilization of product storage (α^{sp}, y_s^{sp}). The process control system is then tasked with meeting these production targets throughout by adjusting the manipulated variables for the process (u_p) and storage system (u_s) while ensuring that the plant operation is stable

and is meeting strict operational, product quality ($\tilde{y}$), and safety constraints.

To account for the dynamics of the (day-ahead) electricity market, the scheduling horizon must be at least two days, and the schedule must be recomputed whenever new market information becomes available. The corresponding optimization problem, using a dynamic model of the process as motivated above, can be expressed as:

$$\min_{y_p^{sp}(t)} \int_0^{T_m} \Phi(pr, x) dt \tag{1}$$

s.t. *Product storage model*
Process dynamic model
Process operating constraints
Product quality constraints

where the price (pr) and demand ($\bar{y}$) forecasts are assumed to be known for the entire scheduling horizon (T_m). The objective function (Φ) accounts for the operating costs, and the decision variable is the production target sequence (y_p^{sp}). We typically assume that the product storage system is described by a transient mass balance model, and x are the process states. Additionally, path constraints are included to ensure that the process operating limits and product quality constraints are satisfied throughout the horizon.

Problem (1) is an infinite-dimensional dynamic optimization, with a highly nonlinear and high-dimensional set of dynamic constraints given by the (typically differential-algebraic) process model. Obtaining a numerical solution entails performing a discretization of the decision variables (and process dynamics) over the time horizon of interest, resulting in a very large nonlinear optimization problem that must be solved in a short time frame to ensure that online implementation is possible. Several works have attempted this approach, with limited results for industrially relevant problems (Bansal, 2003).

Conversely, the aforementioned SBMs are designed to represent explicitly the closed-loop behavior of the process and its control system in an input-output form; that is, a SBM accepts the outputs of the scheduling layer as inputs (notably, production and product grade *targets*), and outputs the time evolution of scheduling-relevant variables, such as the production and the product grade, as well as the evolution of variables pertaining to operating and safety constraints.

In our recent work (Pattison et al., 2016a), we proposed an algorithm for selecting scheduling-relevant variables from the plethora of measurements usually available in modern chemical plants. We demonstrated that the number of scheduling-relevant variables is typically very small compared to the number of states in an industrial process. We also proposed using system identification techniques to obtain SBMS from historical, closed-loop process operating data, observing that (routine) production transitions made in the past provide a rich dataset resembling system identification experiments. Using Hammerstein-Wiener

forms, we showed that the derived data-driven SBMs are single-input, multi-output and sparse, therefore lending themselves very naturally for use in integrated scheduling and control formulations where reduced problem dimensions and fast execution are of essence.

With these models, the production scheduling optimization problem becomes:

$$\min_{y_p^{sp}(t)} \int_0^{T_m} \Phi(pr, w) dt \qquad (2)$$

s.t. *Product storage model*
$$SBMs \left(0 = \Psi\left(y_p^{sp}, w, \dot{w}, \ddot{w}, \dots\right)\right)$$
$$Operating\ constraints: \hat{h}_p(w, y_p^{sp}, t) \leq 0$$
$$Product\ constraints: \hat{h}_y(w, \bar{y}, t) \leq 0$$

where the low-order scale bridging models (Ψ) predict the evolution of the scheduling relevant process variables (w). The path constraints are expressed in terms of the SBM and the scheduling-relevant variables as $\hat{h}_p$ and $\hat{h}_y$.

Closed-loop moving horizon scheduling

The solution of the scheduling problem (2) formulated above is optimal in the (idealized) case where the SBMs provide perfect predictions of the process dynamics over the horizon T_m and there are no disturbances. However, in practice, the process is subject to operational disturbances, plant-model mismatch, operational faults, etc. Additionally, it is unlikely that the schedule remains optimal when updated market condition forecasts (price and product demand) become available. As a consequence, we propose a novel moving-horizon scheme whereby we "close the scheduling loop" based on two sets of criteria, i) market-driven updates, which are either periodic (e.g., associated with regular changes in price and/or demand forecasts) or event-driven (related to unexpected changes in the aforementioned market conditions) (Pattison et al., 2016b) and, ii) the occurrence of process-driven faults, as reported by a process fault diagnosis system (Touretzky et al., submitted), which is assumed to be available. A block diagram of the infrastructure supporting the proposed framework is given in Figure 2.
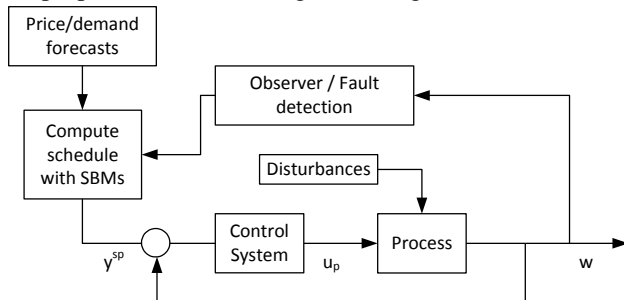

Figure 2. Closed-loop scheduling framework

The implementation proceeds as follows:
1) The production target sequence (schedule) is computed solving problem (2) for the horizon T_m for which accurate forecasts of the feedstock availability, price and product demand are available
2) The control system steers the process to the desired targets and handles high frequency disturbances
3) Measurements of the scheduling-relevant process variables are recorded and an observer is used to update the states of the SBMs
4) Market conditions are monitored and process measurements are used to detect and isolate process faults, triggering rescheduling as relevant
5) Recompute of the schedule (periodic or event-triggered), shifting the time horizon and updating the states of the SBMs
6) Return to step 1, repeat at the next update interval or scheduling or process event

In the case of process disturbances, Step 5 above typically entails updating the process constraints ($\hat{h}_p$) to reflect the impact of process faults (Touretzky and Baldea, 2016). Moving horizon scheduling also requires that concerns related to inventory stabilization be addressed; this can be done, e.g., by imposing end-point constraints on inventory levels (Pattison et al., 2016b). Process faults can be detected using one of many available. We recall that in this work we consider faults that are not critical in the sense of requiring a process shut-down; rather, they manifest themselves in limitations in the process operation, and can be captured via updating the process parameters and/or operating constraints.

As an additional remark, we note that when an event-triggered rescheduling is performed, the time horizon T_m can be shortened to reflect the period of time for which market data forecasts are available (since their update may not coincide with the time of occurrence of the event).

Case Study: Demand response operation of an air separation unit

We consider the operation of an air separation unit (ASU) producing purified nitrogen gas. The process makes use of a refrigeration cycle to cool and liquefy the inlet air stream and a cryogenic distillation column to separate nitrogen from oxygen and argon. The ASU is outfitted with a liquefier and a product storage tank to enable DR operation. A detailed model of the ASU is based on the work of (Cao et al., 2015), and scale-bridging models of the Hammerstein-Wiener form have been identified for eight selected scheduling-relevant variables (Pattison et al., 2016a).

Electricity prices over a six day window in July 2013 are obtained from, ERCOT, the Texas Independent System Operator. We assume that (perfectly accurate) 2-day forecasts of the electricity price and product demand are available and updated every six hours (see the top of Figure 3 – the vertical lines indicate the points at which new forecasts are available). Thus, the schedule is recomputed on a periodic basis every 6 hours. In this case study, we assume that the product demand is constant at 20 mol/s and impurities in the nitrogen product must remain below 1900ppm. Additionally, we assume that the process can operate within +/- 20% of the nominal capacity.

To illustrate the implementation of our moving horizon, closed-loop scheduling framework, we consider two scenarios. First, the nominal case where no fault is

present – the results are plotted by the solid black lines in Figures 3 & 4. The result, as expected, is overproduction and product storage accumulation during the evening hours when electricity prices are low, and a reduced production rate and consequent storage depletion during the daytime when electricity is expensive. The result is a 4.8% savings in electricity costs in comparison to a constant production rate profile (20 mol/s) over the 4 day time frame.

The second case (plotted in dashed red lines in Figures 3 & 4) considers the occurrence and detection of vibrations in the operation of the main air compressor of the ASU at hour 51, requiring that the upper bound on throughput be reduced from +20% to +5% of the nominal capacity. Detection of this fault triggers a rescheduling event (this occurs between the periodic re-execution points when new price and demand forecasts are available), and requires that the process operating constraints be updated to reflect the fault – specifically, the production target maximum is reduced to the nominal value +5%.

Intuitively, following the detection of the fault and the corresponding schedule update, the rate at which the stored liquefied product accumulates is considerably reduced due to the limitation of the maximum throughput in the compressor (bottom of Figure 3). Nevertheless, Figure 4 shows that in both cases, the strict process operating and product quality constraints are met, thereby illustrating the benefit of including (and appropriately updating) a dynamic process model in the scheduling formulation.
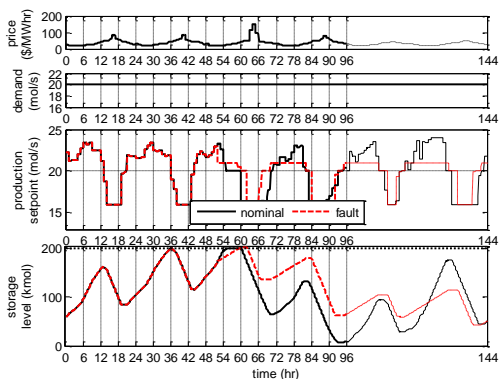

Figure 3. Simulation results for nominal and faulty operation cases. The fault is detected at hour 51.
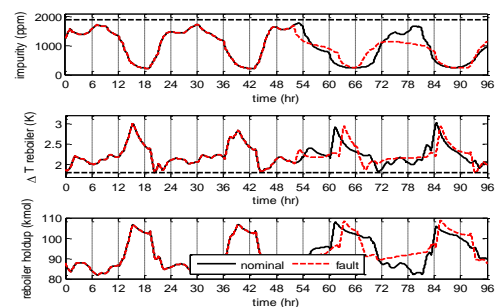

Figure 4. Operational and product quality constraints

Conclusions

In this paper, we propose a moving-horizon production scheduling framework with dynamic constraints for continuous processes operating under fast-changing market conditions. For computational efficiency, the process dynamics are representing using our previously-introduced sale-bridging models. We close the scheduling loop via, i) periodic rescheduling based on regular market condition updates and, ii) event-driven rescheduling in the presence of both scheduling-level disturbances and process disturbances, the latter characterized via a fault diagnosis mechanism. In practice, this framework is ideally suited for managing the DR operation of energy-intensive processes, whereby it enables the optimal exploitation of production capacity and product storage facilities to take advantage of time-sensitive electricity prices. We illustrate these developments on an industrial-scale air separation case study, showing clear economic benefits under realistic circumstances when process operations are subject to faults. We expect that these results are easily transferable to other systems that are highly dependent on the electric grid, as well as to any production facilities that operate under market circumstances that feature fast and significant fluctuations.

References

Baldea, M., Harjunkoski, I. (2014). Integrated production scheduling and process control: A systematic review. *Comput. & Chem. Eng.*, *71*, 377-390.

Bansal, V., Sakizlis, V., Ross, R., Perkins, J. D., & Pistikopoulos, E. N. (2003). New algorithms for mixed-integer dynamic optimization. *Comp. & Chem. Eng.*, *27*, 647-668.

Cao, Y., Swartz, C. L., Baldea, M., & Blouin, S. (2015). Optimization-based assessment of design limitations to air separation plant agility in demand response scenarios. *Journal of Process Control*, *33*, 37-48.

Du, J., Park, J., Harjunkoski, I., & Baldea, M. (2015). A time scale-bridging approach for integrating production scheduling and process control. *Comp. & Chem. Eng.*, *79*, 59-69.

Miller, J., Luyben, W., and Blouin, S. (2008). Economic incentive for intermittent operation of air separation plants with variable power costs. *Ind. & Eng. Chem. Res.*, *47*, 1132–1139.

Pattison, R.C., Touretzky, C.R., Johansson, T., Harjunkoski, I., 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*, 4562-4584.

Pattison, R. C., Touretzky, C. R., Harjunkoski, I., & Baldea, M. (2016). Moving horizon closed-loop production scheduling using dynamic process models. *AIChE J.*

Touretzky, C. R. & Baldea M. (2016). Closing the scheduling loop: incorporating dynamic models, fault based triggers, and moving horizon concepts. Submitted.

Wang, R.C., Edgar, T.F., Baldea, M., Nixon, M., Wojsznis, W., Dunia, R. (2015). Process fault detection using time-explicit Kiviat diagrams. *AIChE J.*, *61*, 4277-4293.

Zhang, Q., Grossmann, I.E., Heuberger, C.F., Sundaramoorthy, A., and Pinto, J.M. (2015). Air separation with cryogenic energy storage: optimal scheduling considering electric energy and reserve markets. *AIChE J.*, *61*, 1547-1558.