Optimal Operation of Process Plants Under Partial Shutdown Conditions

Zhiwen Chong and Christopher L. E. Swartz
Dept. of Chemical Engineering, McMaster University, 1280 Main Street West, Hamilton, Ontario, Canada, L8S 4L7

DOI 10.1002/aic.14154
Published online June 25, 2013 in Wiley Online Library (wileyonlinelibrary.com)

A systematic strategy for optimal plant operation during partial shutdowns is proposed. We consider the situation where one or more process units are shut down due to failure or maintenance but where the remaining units are able to continue operation to some degree. The goal of the strategy is to manipulate the plant degrees-of-freedom—during and after the shutdown—such that production is restored in a cost-optimal fashion while meeting safety and operational constraints. Optimal control trajectories are obtained through the solution of a dynamic optimization problem. A novel multitiered optimization approach allows the prioritization of multiple competing objectives and the specification of tradeoffs between them. Uncertainty in the downtime estimate, a crucial parameter in shutdown optimization, is addressed through reoptimization. We employ a transient predictive control algorithm for implementing the computed control policy under feedback. © 2013 American Institute of Chemical Engineers *AIChE J*, 59: 4151–4168, 2013
Keywords: dynamic optimization, partial shutdowns, model predictive control, multi-unit systems, abnormal situation management

Introduction

In a typical chemical plant, raw materials are transported through a number of processing units to be fashioned into some end-product. Process units are shut down from time to time either for maintenance or due to equipment failure. Shutdowns in intermediate units constitute a disruption in the processing chain, which can adversely affect the plant's ability to continue operation. They also generally have an adverse effect on the operating economics of a plant due to the attendant loss of production capacity.

From an operations perspective, plant shutdowns can be classified as either critical or partial, with the former being those that lead to the shutdown of the entire plant and the latter, those that do not. In the case of critical shutdowns, the entire system is invariably forced to shut down, and under such circumstances the usefulness of any optimal control policy is limited. Under partial shutdown scenarios, however, it is frequently possible for an operator to pursue certain courses of action that will permit the unaffected units to continue operating to some degree. Possible remedial actions for handling partial shutdowns include reconfiguring the process pathways, rerouting material streams, slowing down production, and so forth.

Many of the approaches for combating the effects of partial shutdowns rely on the use of redundancies in the plant, and buffer tanks provide such a redundancy. There are significant benefits to having judiciously placed buffer capacities in a plant, especially, for the mitigation of process variation propagation along a production line. Buffer capacities serve to decouple various segments of a plant. They are not only able to dampen the effects of short term fluctuations, but are also able to deal with larger processing disturbances. In the case of partial shutdowns, they can effectively circumscribe the impact of a unit shutdown by decoupling the offline unit(s) from the rest of the plant.

When a process unit downstream of a buffer tank fails, the tank can hold and accumulate material for a period of time until the process unit is brought back online. Likewise, when a process unit upstream of a buffer tank fails, the material held in the buffer tank can be slowly discharged to the downstream units in order that downstream processing may continue.

The successful management of buffer inventories, recycles, and production rates during the shutdown period is one of the keys to minimizing the losses associated with partial shutdown scenarios. The end-product quality during the shutdown period may violate normal specifications, but in some cases, the material can be recycled and reprocessed.

Pettersson[1] developed a scheme for coordinating production in a pulp and paper mill, which included a systematic approach for managing the buffer tanks. A system comprising nine processing units and ten buffer tanks was considered for optimization, resulting in a production scheme for coordinating the plant in such a way that capacity restrictions were not violated. An example with a shutdown in the evaporator system was shown. In a later publication, Pettersson[2] considered the problem of producing an optimal plantwide production scheme that could account for maintenance shutdowns, limited buffer capacity and steam restrictions. The resulting optimal control problem was solved using Pontryagin's Maximum Principle[3].

Correspondence concerning this article should be addressed to C. L. E. Swartz at swartzc@mcmaster.ca.

Lee and Reklaitis[4] proposed a method for systematically utilizing buffer capacities to decouple upstream equipment failures from downstream processes and vice versa. Using Fourier series constructions, they derived a set of analytical expressions for determining the minimum volume of intermediate storage required as a function of frequency of failure and of failure durations. In a related article, Lee and Reklaitis[5] noted that overspecified intermediate storage led to the introduction of delays in product change-overs, thus there is considerable economic benefit to specifying only as much capacity as is absolutely required.

Allison[6] used an analytical approach for determining a policy for averaging the loads on a set of surge tanks in series in the presence of flow disturbances. In this scheme, the impact of a surge on a single tank is distributed across the plant to avoid upper and lower level constraints on buffer capacities. Optimal and suboptimal control strategies were explored.

Campo and Morari[7] presented a model predictive control formulation of the discrete-time optimal average level control problem. The primary goal of the controller is flow filtering; that is, the attenuation of flow disturbances propagating to downstream units. The control algorithm minimizes the rate of change of the outlet flows from surge tanks, while providing integral action and handling of tank level and outlet flowrate constraints. Endpoint contraints are imposed at the end of the prediction horizon to drive the levels back to their nominal ranges. The horizon length p is a tuning parameter that trades-off good flow filtering (requiring large p) and rapid integral time or settling (requiring small p).

Huang et al.[8] advocated the idea of using dynamic optimization within a fault accommodation framework, as opposed to the manual table-lookup approach typically adopted by operators. In a subsequent publication,[9] the issue of unit shutdowns—which the authors consider to be a type of fault—in a dynamic optimization context was briefly discussed. With regard to optimizing dynamic models with unit failure representations, the authors suggested first removing all equations related to the unit, followed by the activation of a set of discrete transition equations triggered using integer variables. The authors emphasized that the integer variables only act as transition conditions, and are directly prescribed by a fault detection module, hence the final model to be solved does not contain integer variables. A detailed formulation for addressing unit shutdowns was not provided.

Dubé[10] investigated a buffer storage operation strategy that minimizes time away from normal operation and seeks to prevent plant section shutdowns. This strategy was applied to a highly integrated Kraft pulp mill with the view of maximizing production. The author addressed the issue of determining the longest feasible shutdown time. The effect of preparation time on the production was also studied and the coordination of buffer capacities for handling planned and unplanned shutdowns was illustrated. A numerical optimization procedure was used in the solution process.

Continuing along the lines of Dubé's work, Balthazaar[11] considered both a pre-emptive and a reactive response to shutdowns in a Kraft paper mill. In the pre-emptive case, the control problem is solved assuming knowledge of the shutdown ahead of time (as is the case in a scheduled maintenance scenario), thus allowing the plant to take preparatory action in anticipation of the shutdown. In the reactive case, the plant shutdown occurs without warning, and the control system is expected to respond immediately. Balthazaar also examined the problem of determining the location and capacity of additional buffer tanks, as well as steady-state buffer tank levels that are optimal for handling shutdowns, using probability distributions on failure type and duration. An economics-based objective function was employed in all the case studies.

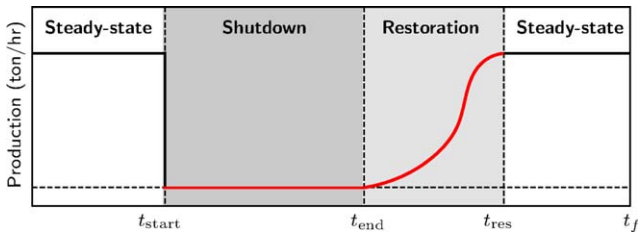

Figure 1. Operating phases of a single process unit, under partial shutdown conditions.

[Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com].

Key limitations in the aforementioned approaches are that the issue of uncertain downtimes is not treated in the process operations problem, as well as the handling of uncertainties due to plant-model mismatch and disturbances in the process. In our work, we address these issues using a multi-tiered economic model predictive control (MPC)-type framework. Here, we confine our attention to adjustment of the operation through management of material flowrates. However, the framework we have adopted is general and may be extended to the other types of remedial actions listed above.

This article is organized as follows. In the next section, we present a formulation for posing the partial shutdown problem within a dynamic optimization framework. Following that is a description of the MPC framework used to implement the solution trajectories obtained via dynamic optimization. Several simulation studies demonstrate the application of the described framework to the fiber line in a pulp-and-paper plant.

Dynamic Optimization Problem Formulation

Figure 1 illustrates the different phases of the shutdown process, as considered in this study. A process unit initially operates at a certain steady-state point, until a disruption causes it to switch into shutdown mode. The shutdown phase begins at t_{start}, when all input and output flowrates to the process unit (and unbuffered process units adjacent to it) are forced to 0. The shutdown phase proceeds until t_{end}. In this work, we assume that all necessary measures are taken to repair and restart the unit between t_{start} and t_{end}. At time t_{end}, the unit is deemed to be in a state that is ready for operation, and the restoration phase begins. In the restoration phase, control actions are prescribed to the plant to return it to its original steady-state operating point (or to some nearby nominal operating point). The restoration phase terminates at time t_{res}. At this juncture, the plant has been successfully restored to normal operation. The partial shutdown problem entails finding a set of control trajectories that will guide the operation of the plant and restore it to a nominal operating state in a cost-optimal fashion. This can be cast as a dynamic optimization problem of the following general form:

$$\max \Phi_{\text{econ}}(x(t), z(t), u(t)) \tag{1}$$

$$\text{s.t.} \quad \dot{x}(t) = f(x(t), z(t), u(t), \theta), \qquad x(0) = x_0 \tag{2}$$

$$h(x(t), z(t), u(t), \theta) = 0 \tag{3}$$

$$g(x(t), z(t), u(t), \theta) \leq 0 \tag{4}$$

for $t \in [0, t_f]$, where Φ_{econ} = economic objective functional, $x(t) \in \mathbb{R}^{n_x}$ = differential state vector, x_0 = initial state parameter vector, $z(t) \in \mathbb{R}^{n_z}$ = algebraic state vector, $u(t) \in \mathbb{R}^{n_u}$ = control input vector, $\theta \in \mathbb{R}^{n_\theta}$ = parameter vector, t_f = final time in prediction horizon. f and h are functions that define a semi-explicit differential-algebraic equation (DAE) system, and g is a function that defines the inequality constraints. We assume that the DAE system described by Eqs. 2 and 3 is of index-1; that is, g is solvable for z (equivalently, we require that g is at least once differentiable, and that its Jacobian with respect to z, or g_z, is nonsingular).

In the sequel, we will consider the discretized problem directly, with time t partitioned and mapped into N intervals. Piecewise-constant inputs are applied, and are represented as $u_k = u(t_k)$, for $k = 0, \ldots, N$ where $N = \lfloor t_f / \Delta t \rfloor$ and Δt is the discretized interval length.

The above DAE optimization problem can be solved in a variety of ways. In this work, we pursue the approach where we discretize the states $x(t)$ and $z(t)$ using the method of orthogonal collocation on finite elements.[12] This method effectively converts the above DAE system into a finite system of algebraic equations. The resulting equations are then posed as equality constraints in a mathematical programming problem. The mathematical modeling in this work was accomplished using an in-house modeling system, A Modeling Language for Dynamic Optimization (MLDO),[13] which accepts a DAE model as input and generates the requisite discretization constraints in the AMPL mathematical programming language.[14]

Shutdown constraints (α_{shut} formulation)

In engineering flowsheet models, each process unit p (where $p \in \mathcal{P}$ = set of all process units) is often modeled individually. Connection equations are used to connect the inputs and outputs of each process unit model. In a chemical engineering context, these connections are often used to couple the material and energy streams of different process units.

In order to represent shutdown behavior in individual process units within the flowsheet, we impose constraints that shut off the material flow F through a particular unit p. While our framework is readily extended to handle pre-emptive responses to planned shutdowns, in this work, we elect to confine our attention to reactive responses to shutdowns that occur instantaneously and without warning (for instance, shutdowns that are precipitated by equipment failure).

Let $\mathcal{I}_p$ denote the index set of all inlet flow variables associated with unit p, and $\mathcal{L}_p$, the index set of all variables associated with unit p that float when the flowrate is 0. Floating variables include concentrations, compositions, intermediate variables, and other quantities whose values are not defined when the material flow is 0 (see Remark 3 below). We model the shutdown of a process unit p through the following logical conditions:

$$\begin{cases} F_{i,k} \in \left[F^p_{\text{shut}}, F^U_i\right], & \text{if} \quad \alpha^p_{\text{shut},k} = 0 \\ F_{i,k} = 0, & \text{if} \quad \alpha^p_{\text{shut},k} = 1 \end{cases}, \forall i \in \mathcal{I}_p \tag{5}$$

$$\text{and,} \quad \begin{cases} y_{l,k} \in \left[y^L_l, y^U_l\right], & \text{if} \quad \alpha^p_{\text{shut},k} = 0 \\ y_{l,k} = y^M_l, & \text{if} \quad \alpha^p_{\text{shut},k} = 1 \end{cases}, \forall l \in \mathcal{L}_p \tag{6}$$

where $\alpha^p_{\text{shut},k} \in \{0, 1\}$ = indicator variable representing status of process unit p at time k (1 = shutdown, 0 = normal operation); $y_{l,k}$ = floating variable at time k; y^L_l, y^U_l = lower and upper bound parameters corresponding to y_l; y^M_l = nominal value parameter to which $y_{l,k}$ is pinned to when $\alpha^p_{\text{shut},k} = 1$; $F_{i,k}$ = inlet flowrate variable at time k; F^U_i = upper bound of $F_{i,k}$ for all k; F^p_{shut} = a positive nonzero parameter representing the minimum flowrate threshold before a unit is deemed to have shut down.

These conditions can be formulated through the following constraints:

$$F_{i,k} \geq \left(1 - \alpha^p_{\text{shut},k}\right) F^p_{\text{shut}} \tag{7}$$

$$F_{i,k} \leq \left(1 - \alpha^p_{\text{shut},k}\right) F^U_i \tag{8}$$

$$y_{l,k} \geq \left(1 - \alpha^p_{\text{shut},k}\right) y^L_l + \alpha^p_{\text{shut},k} y^M_l \tag{9}$$

$$y_{l,k} \leq \left(1 - \alpha^p_{\text{shut},k}\right) y^U_l + \alpha^p_{\text{shut},k} y^M_l \tag{10}$$

for $k = 0, \ldots, N-1$, $i \in \mathcal{I}_p$, $l \in \mathcal{L}_p$, $p \in \mathcal{P}$. We refer to the above as the α_{shut}-formulation.

One of the consequences of constraints 7–10 is that when $F_{j,k} < F^p_{\text{shut}}$ (that is, when the flowrate falls below a threshold), the unit is deemed to have shut down. For instance, suppose we wish to manually simulate a shutdown in a particular unit p'. This would entail shutdowns in the adjacent units of p' that have no buffer tanks surrounding them. Suppose, we denote the set of all shutdown units adjacent to p' as $\mathcal{P}_{\text{shut}}$ (where $\mathcal{P}_{\text{shut}} \subset \mathcal{P}$). We can then explicitly impose the following constraint where the values of $\alpha^p_{\text{shut},k}$ are fixed as follows:

$$\alpha^p_{\text{shut},k} = \begin{cases} 1, & p \in \mathcal{P}_{\text{shut}}, k = k_{\text{start}}, \ldots, k_{\text{end}} \\ 0, & \text{elsewhere} \end{cases} \tag{11}$$

where $k_{\text{start}} = \lfloor t_{\text{start}} / \Delta t \rfloor$ = shutdown start time; $k_{\text{end}} = \lceil t_{\text{end}} / \Delta t \rceil$ = shutdown end time (see Figure 1 for the temporal locations of t_{start} and t_{end}). Even though $\alpha^p_{\text{shut},k}$ is a discrete quantity, this explicit fixing of variables (where $\alpha^p_{\text{shut},k}$ are specified parameters in this case) results in the optimization problem being a continuous (in its decisions) rather than a discrete one. That said, this α_{shut} formulation lends itself quite naturally to discontinuous mixed-integer formulations of state-based shutdown triggers by a simple redeclaration of α_{shut} as binary decision variable vector.

Remark 1. In practice, the value of α_{shut} (as well as an estimate of the value of the expected endpoint of the shutdown, k_{end}) at a given time may be obtained from a plant fault detection module, or set manually by an operator based on past knowledge of similar shutdowns. The variable α_{shut} can be thought of as a flag or a discrete trigger that captures the operational state of the system and allows us to activate the relevant shutdown actions at the right time.

Remark 2. The F^p_{shut} quantity is directly related to a quantity familiar to industrial practitioners, namely the turndown ratio, R^p_t, for a unit p:

$$R^p_t = \frac{\text{Normal maximum flow}}{\text{Minimum controllable flow}} = \frac{F^p_U}{F^p_{\text{shut}}} \tag{12}$$

For instance, if due to design reasons a unit is only able to operate at a minimum of 25% of its maximum capacity, the

turndown ratio $R_t = 100/25 = 4$. In this context, if the unit is operating below a threshold F^p_{shut} and exceeding the turndown ratio, the unit will need to be shut off.

Remark 3. Floating variables frequently occur as a result of bilinear formulations. Consider the following stream equations example (where F_P, F_W, F_{DS} are component flows, and x_p is a composition variable):

$$F = F_P + F_W + F_{DS} \tag{13}$$

$$F_P = x_P F \tag{14}$$

When component flows $F_P = 0$, $F_W = 0$, and $F_{DS} = 0$, it follows that $F = 0$. We can then deduce from Eq. 14 that x_P is allowed to take any value within its given bounds. Physically, x_P is meaningless when $F = 0$ because the composition of a stream with zero flowrate is undefined. Allowing x_p to float in an optimization problem can induce nonunique solutions. This can be avoided by pinning down the value of x_p when $F = 0$, accomplished through constraints (9,10).

Shutdown modeling assumptions

We assume that the time to restart a failed unit is either: (1) short relative to the overall process dynamics, and is hence neglected, or (2) included in the shutdown duration. The implication of both these assumptions is that a unit is deemed to be immediately ready for operation at the end of a shutdown phase (i.e., at t_{end}; refer to Figure 1). On a similar note, the details of the shutdown procedure itself are not modeled; we assume that a manual or an automated procedure for start-ups and shutdowns is in place.[9]

We consider these assumptions to be reasonable given that we are essentially solving a macro-level inventory/production problem (that is, we are concerned with the overall ingress and egress of material from process units, as well as the buffer inventories), which is an abstraction level above the actual operating details of the offline unit.

Restoration constraints

Once the shutdown period is over, the system undergoes a restoration phase where it is brought back to some nominal operating condition. At the end of the restoration phase, we require the independent states and inputs to attain some nominal (steady-state) values. We can express this requirement in the optimization problem by way of restoration constraints, that is, hard terminal constraints imposed on the independent states and inputs post-restoration (during $t \in [t_{\text{res}}, t_f]$ in Figure 1). We assume that the number chosen for the horizon length N is high enough to permit a feasible restoration within its defined time constraints.

Allison[6] demonstrated that in the absence of restoration constraints, the optimal course of action in buffer inventory optimization problems is to drive the storage vessels empty. This is clearly an undesirable outcome in our case as it would deprive the plant of the ability to handle any further shutdowns past the end of the dynamic optimization horizon.

In our partial shutdown problem, we are interested in driving a plant toward a steady state after the shutdown. Ideally, at the end of the optimization horizon one would like every differential state x (including the nonindependent states) to return to their exact pre-shutdown values, that is, the original steady state.

In practice, the nonindependent states may not be returnable to their original steady-state values within the specified finite optimization horizon. Certain states, such as compositions for instance, may take a long time to reach a steady-state point. The time duration required for these states to reach steady state may lie outside of the time horizon of the dynamic optimizer. We note that one may, at one's discretion, elect to extend the horizon as needed; however, this will incur additional computational cost, therefore, a balanced approach is proposed.

Instead of attempting to return every state to its pre-shutdown value, we aim to return only the levels of the buffer inventories (independent) to their nominal states at the end of the given horizon, and prescribe a set of inputs that, when held constant, will allow the compositions (nonindependent) to eventually settle to their steady-state values. This is done to allow the plant to anticipate further shutdowns.

To steer the system toward a steady-state operating point at the end of the horizon, we employ restoration constraints such as the following in our control optimization, enforced at the last $n_f + 1$ discrete time points in the prediction horizon:

$$x^{i_x}_{ss}(1-\epsilon) \le x^{i_x}_k \le x^{i_x}_{ss}(1+\epsilon), \qquad k \in \mathcal{K}_f, i_x \in \mathcal{I}^x_{ss} \tag{15}$$

$$u^{i_u}_{ss}(1-\epsilon) \le u^{i_u}_k \le u^{i_u}_{ss}(1+\epsilon), \qquad k \in \mathcal{K}_f, i_u \in \mathcal{I}^u_{ss} \tag{16}$$

where $\mathcal{K}_f = \{N - n_f, \ldots, N\}$, $N =$ the horizon length of the controller, $x^{i_x}_k$ and $u^{i_u}_k$ are elements of the state x and input u vectors at time k, indexed by i_x, i_u respectively. The index sets are defined as $\mathcal{I}^x_{ss} = \{i | x_i \in \mathcal{X}_{ss}\}$ and $\mathcal{I}^u_{ss} = \{i | u_i \in \mathcal{U}_{ss}\}$, where $\mathcal{X}_{ss}$, $\mathcal{U}_{ss}$ are the sets of independent state and input variables necessary to specify a steady-state. $\epsilon \ge 0$ is a deviation tolerance parameter. $x^{i_x}_{ss}, u^{i_u}_{ss}$ are (constant) pre-shutdown values. The time point $(N - n_f)$ is related to restoration endpoint t_{res} as follows:

$$N - n_f = \lfloor t_{\text{res}} / \Delta t \rfloor + 1 \tag{17}$$

We wish to emphasize again that only independent states and inputs are to be selected for restoration (a detailed discussion follows in the sections below), for electing to restore all state and input variables to their original values may lead to an over-specified system and may cause convergence difficulties in the presence of plant-model mismatch. In most inventory management problems, $\mathcal{X}_{ss}$ typically comprises inventory level variables, and $\mathcal{U}_{ss}$ often contains a subset of the flowrates in the plant.

Determining the subset of states and inputs to specify an endpoint steady state

A system is said to be at steady state when its state variables do not change with time. In a dynamic model, this corresponds to the derivatives of its differential states being 0 (that is, $\dot{x} = 0$). A model may have a multiplicity of steady states.

In principle, one could impose hard constraints that require the derivatives of the state variables to be driven to zero at the prediction horizon endpoint. However, given that we are solving a fixed horizon problem (with a limited mandate), certain quantities may not reach steady state within the prescribed prediction horizon. One solution is to lengthen the prediction horizon, but this incurs added computational cost in the optimization problem. Our proposed method aims instead to fix all the independent state and input variables that are necessary for the system to eventually settle at a steady state (a point which may occur outside of the prediction horizon).

To illustrate the difference between independent and dependent states and inputs, we present the following example. Consider a dynamic model of the form:

$$\frac{dm}{dt} = F_5 - F_6, \qquad m(0) = m_0 \tag{18}$$

$$m\frac{dx^A}{dt} = F_5\left(x_5^A - x^A\right), \qquad x^A(0) = x_0^A \tag{19}$$

$$F_1 = F_2 + F_3 \tag{20}$$

$$F_2 = \alpha F_1 \tag{21}$$

$$F_1 x_1^A = F_2 x_2^A + F_3 x_3^A \tag{22}$$

$$F_3 + F_4 = F_5 \tag{23}$$

$$F_3 x_3^A + F_4 x_4^A = F_5 x_5^A \tag{24}$$

where $\alpha, x_1^A, x_2^A, x_4^A$ are constant parameters; m, x^A are differential state variables; F_1, F_4, F_6 are control input variables (decision variables for the optimizer); and $F_2, F_3, F_5, x_3^A, x_5^A$ are algebraic variables. There are 10 variables (including the inputs), and 7 equations. Therefore, three independent specifications are possible when the system is in transition (that is, nonsteady state), and we select the control input variables as the specifications to the system.

At steady state however, the conditions $dm/dt = 0$ and $dx^A/dt = 0$ impose additional restrictions on the system. They imply that $F_6 = F_5$, therefore F_6 ceases to be an independent input variable at steady state—its value depends on F_5—and therefore its value cannot be specified arbitrarily. As well, $x^A = x_5^A$, but it can be shown that this can be achieved with any choice of F_1 and F_4 that lie within some feasible operating window, therefore these are independent input variables. Note that the variable m does not appear in the steady-state equations, therefore, it may be independently specified.

Through this kind of reasoning[1], we are able to deduce a set of independent states and inputs at steady state comprising m, F_1, and F_4. If we specify the vectors $x = [m, x^A]^{\mathrm{T}}$ and $u = [F_1, F_4, F_6]^{\mathrm{T}}$, we obtain $\mathcal{X}_{\mathrm{ss}} = \{m\}$, $\mathcal{U}_{\mathrm{ss}} = \{F_1, F_4\}$, and $\mathcal{I}^x_{\mathrm{ss}} = \{1\}$, $\mathcal{I}^u_{\mathrm{ss}} = \{1, 2\}$.

If one attempts to simultaneously specify F_1, F_4 and F_6 at steady state, one may either end up with a model whose derivatives dm/dt and dx^A/dt are not exactly zero (as would be the case in the above model), or worse, an overspecified model that is numerically infeasible.

Multitiered optimization

In order to explain the motivation behind the multitiered optimization approach, we first furnish the reader with a brief description of the issues surrounding nonunique solutions in control optimization problems, and the deviation of plant product qualities from their targets during the shutdown and restoration phases.

Input regularization and Nonuniqueness. Situations can arise in which the computed optimal input trajectory $u_k \in \mathbb{R}^{n_u}$ (for $k = 0, \ldots, N-1$) is nonunique, that is, there exists more than one input trajectory that corresponds to an identical objective value at the optimum. In mathematical programming terms, we say there exist nonstrict local minimizers; that is, there are directions in which the objective function is flat. When this occurs, a numerical optimization procedure has no recourse for differentiating between the different trajectories, and is thus liable to return the first solution trajectory that satisfies some given specified termination condition. Some of these solution trajectories may exhibit a high-frequency chatter-like behavior, while others prescribe large changes in an input variable over a short period of time. Such trajectories, though optimal in a mathematical sense, are undesirable in an operational sense because they give rise to actuator/valve wear and are unappealing to plant operators due their opacity to interpretation.

Current practice prescribes the addition of a regularization term to the objective function to ensure strictness of local minimizers. This is typically done through a move suppression penalty on the input trajectory u:

$$\max \qquad \Phi_{\mathrm{econ}}(x, z, u) - \sum_{k=1}^{N-1} (u_k - u_{k-1})^{\mathrm{T}} R (u_k - u_{k-1}) \tag{25}$$

where $R = \mathrm{diag}\left[\rho_1, \ldots, \rho_{n_u}\right]$ is a user-specified penalty weight matrix, and ρ_i (for $i = 1, \ldots, n_u$) are positive-valued penalty weights that are sufficiently large to ensure strictness of local minimizers. The vectors $x = \left[x_0^{\mathrm{T}}, x_1^{\mathrm{T}}, \ldots, x_N^{\mathrm{T}}\right]^{\mathrm{T}}$ = differential states, $z = \left[z_0^{\mathrm{T}}, z_1^{\mathrm{T}}, \ldots, z_N^{\mathrm{T}}\right]^{\mathrm{T}}$ = algebraic states, and $u = \left[u_0^{\mathrm{T}}, u_1^{\mathrm{T}}, \ldots, u_{N-1}^{\mathrm{T}}\right]^{\mathrm{T}}$ = control inputs. The constraints in this optimization problem are assumed to satisfy the linear independence constraint qualification (LICQ), a regularity property.

We note that a concomitant benefit of having this particular regularization term in the objective is that one also obtains smooth trajectories. This technique is often successfully applied in control optimization problems where the object is either to track a setpoint or to regulate a system around an equilibrium point.

However, in our case we have an economic objective function that is denominated in dollar terms. Therefore, there is a certain degree of arbitrariness to the selection of the value of the penalty weights, since the penalty term is not denominated in dollar terms and does not directly represent an economic measure. An arbitrarily configured move suppression penalty is susceptible to trading-off the original economic objective in ways that are not easy to quantify. Therefore, in this work, we elect not to use the move suppression penalty approach.

Product Quality Targets During Shutdown and Restoration. Product qualities (e.g., percentage impurity, minimum concentration, etc.) are a set of variables that are closely regulated around some target value during nominal operation. These variables are monitored to ensure that the material at various stages of the plant have the desired properties for processing.

During the shutdown and restoration phases, some of these product quality variables may be rendered incapable of reaching their targets due to the disruption caused during partial shutdowns. In practice, this necessitates the relaxation or release of a subset of the specifications during the transition period. This entails allowing excursions of certain product qualities outside their normal allowable bands while maintaining the process as close to its nominal operating state as possible.

The task of choosing the most suitable set of product quality variables to relax or release often falls to the plant operator. This is a task that calls for either the use of engineering

[1]*In this work, we utilized both a computer algebra system and engineering principles to identify sets Xss; Uss. In principle, these independent sets could conceivably be arrived at through some a graph-theoretic algorithm; however this is outside the scope of our work.*

heuristics or personal experience. However, we believe a case can be made for using an optimization technique for selecting trajectories that attempt to meet the targets in the nonshutdown units as best as possible, without sacrificing economics.

Multitiered Optimization. In light of the above, we propose an alternative approach to regularizing the trajectories, which takes the form of a hierarchical optimization strategy. In this section, we put forward a multitiered optimization approach, which extends the "two-tiered optimization" approach due to Balthazaar and Swartz.[15] This approach directly exploits the nonuniqueness property of the dynamic optimization problem by systematically selecting a solution—within the space of optimal solutions—that meets certain desirable criteria.

The multitiered optimization approach allows us to prioritize several competing objectives and specify the trade-offs between them. Separate optimization problems are solved one after the other, with the preceding tiers providing information to the latter tiers. One configuration of a multitiered optimization problem is as follows, where in all the tiers described below, we assume that the constraint set includes the dynamic model:

Tier 1: Obtain trajectory for optimal economics. We solve the dynamic optimization problem that maximizes an economic objective function, and record the optimal value of the objective. This value represents an upper-bound on economics achievable by the system.

$$\Phi^*_{\text{econ}} = \max \Phi_{\text{econ}}(x, z, u) \tag{26}$$

where $\Phi_{\text{econ}}(x, z, u) =$ economic objective function, $x = \left[x_0^{\mathrm{T}}, x_1^{\mathrm{T}}, \ldots, x_N^{\mathrm{T}}\right]^{\mathrm{T}} =$ differential states, $z = \left[z_0^{\mathrm{T}}, z_1^{\mathrm{T}}, \ldots, z_N^{\mathrm{T}}\right]^{\mathrm{T}} =$ algebraic states, and $u = \left[u_0^{\mathrm{T}}, u_1^{\mathrm{T}}, \ldots, u_{N-1}^{\mathrm{T}}\right]^{\mathrm{T}} =$ control inputs. The optimal value of the economic objective (a scalar value) is stored as $\bar{\Phi}_{\text{econ}} := \Phi^*_{\text{econ}}(x^*, z^*, u^*)$. We note that in this tier, the product quality targets are allowed to vary within their admissible bounds. In the cases considered in this work, the quality targets constitute (intermediate) operational targets that are noncritical to the final product quality. In other applications, these targets may be essential, in which case Tiers 1 and 2 may be interchanged and their trade-offs appropriately reassigned.

Tier 2: Minimize deviation from product quality targets. In this tier, we solve an optimization problem to deduce the minimum tolerable deviation of product qualities from their target values during the shutdown and restoration phases, while still remaining feasible. Let s denote the vector of product quality variables in the plant, and $\bar{s}$ their nominal targets.

$$\Phi^*_{\text{s}} = \min \sum_{k=0}^{N} (s_k - \bar{s})^{\mathrm{T}} (s_k - \bar{s}) \tag{27}$$

$$\text{s.t.} \quad \Phi_{\text{econ}}(x, z, u) \geq \bar{\Phi}_{\text{econ}}(1 - \varepsilon_e) \tag{28}$$

where $\varepsilon_e \geq 0$ is a tolerance that may also be used as an economic trade-off parameter, representing the maximum fraction of achievable economics (i.e. $\bar{\Phi}_{\text{econ}}$) to trade-off. The squared deviation of the product quality targets at the optimal solution $\left(\bar{\Phi}_{\text{s}} := \Phi^*_{\text{s}}\right)$ is recorded and enforced in the next tier.

Tier 3: Minimize control effort. The objective function is replaced with another objective that minimizes control effort, subject to a minimum constraint on the economic performance and minimum deviation from product quality targets.

$$\min \sum_{k=1}^{N-1} (u_k - u_{k-1})^{\mathrm{T}} (u_k - u_{k-1}) \tag{29}$$

$$\text{s.t.} \quad \Phi_{\text{econ}}(x, z, u) \geq \bar{\Phi}_{\text{econ}}(1 - \varepsilon_e) \tag{30}$$

$$\sum_{k=0}^{N} (s_k - \bar{s})^{\mathrm{T}} (s_k - \bar{s}) \leq \bar{\Phi}_s (1 + \varepsilon_s) \tag{31}$$

where $\epsilon_{\text{s}} \geq 0$ is a tolerance or the product quality deviation trade-off parameter, which represents the fraction of the attainable target deviation to trade-off in order to achieve minimum control effort. Choosing $\varepsilon_{\text{e}} = 0, \varepsilon_{\text{s}} = 0$ may be feasible, but higher values should deliver smoother control profiles, which may be more desirable from an operations point of view. Positive values of ε_{e} and ε_{s} are also favorable from a numerical point of view, as they behave like numerical relaxation factors.

We wish to emphasize that the above setup represents only one possible configuration of a multitiered dynamic optimization problem–many more are possible. Multitiered optimization can be thought of as a general technique for systematically prioritizing control objectives (an idea explored in Swartz[16]) without resorting to using arbitrary penalty weights; instead it allows one to specify the trade-offs between one objective and the next. The first tier represents the most important objective and last tier represents the least important. The relaxation factors ε are used to specify the trade-offs in each tier. Therefore, the order of the individual optimization problems can be switched around according to their importance in a particular application. Indeed, even the optimization problems in each tier can be replaced or removed as needed. Figure 2 shows a simple schematic of the multitiered optimization approach that we have adopted for the present application. We also note that an attractive property of this approach is that if the first tier optimization problem is feasible, feasibility is guaranteed for all successive tiers.

Another attractive feature of this particular multitiered optimization configuration is that unlike the move-suppression penalty approach seen in Eq. 25, it permits the user to decide, for example, the exact amount of the economics to trade-off (in dollar terms as opposed to some arbitrarily weighted quantity) in order to obtain smaller product quality deviations or smoother input trajectories. As well, one can specify the percentage of attainable product quality deviation to trade-off in order to obtain smoother trajectories with the minimum input effort.

Nonlinear MPC for Partial Shutdowns

In previous sections, we have presented dynamic optimization formulations for addressing partial shutdowns in a plant. The dynamic optimization problem can be solved offline and its solution used as an advisory tool for making plant operation decisions. Plant operators can use dynamic optimization as a tool for assessing what-if scenarios, based on different hypothetical operating strategies. This is a common way of using the results from dynamic optimization problems.

This usage case is akin to the use of a process simulator to simulate different scenarios, except dynamic optimization offers the additional facility of being able to compute optimal solutions based on different objective functions. In addition, while process simulation software often includes the

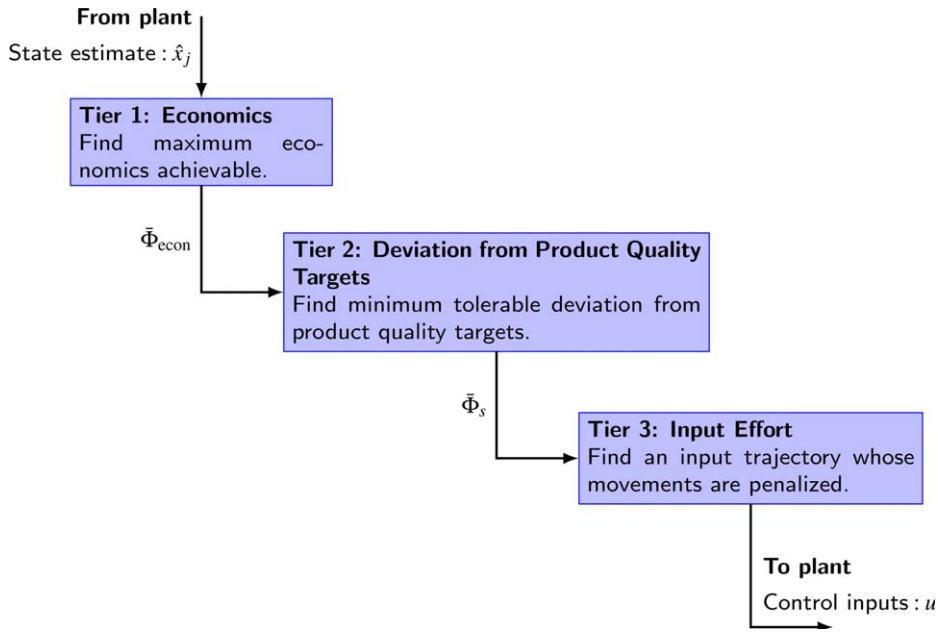

Figure 2. Multitiered optimization approach.

[Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com].

ability to optimize certain time-invariant plant parameters, dynamic optimization techniques offer the ability to optimize time-varying control input trajectories that can be made to obey arbitrarily complex constraints (subject to feasibility). As such, open-loop dynamic optimization is a tremendously useful tool for generating plant operating policies and heuristics.

In principle, it may be possible to directly implement the dynamic optimization solution trajectories on the plant. We note that in reality, process disturbances, noise, and plant-model mismatch may invalidate or render infeasible the calculated trajectories. For the implementation of the solution trajectories to be practicable, they need to reflect the current realities in the plant.

One could conceivably update the dynamic optimization problem with the latest information available from the plant and resolve the problem. This could be done every time an operational decision has to be made. However, if we take the idea one step further and mechanize the procedure, it leads us naturally to the notion of feedback control, where a control algorithm takes an operational decision based on updated plant information at every sampling iteration.

MPC[17] is a framework that is particularly well-suited to handle this case. Toward this end, we propose a nonlinear MPC-based shutdown controller (with an embedded dynamic optimization problem) as a means of implementing abnormal-situation control on the plant. We selected MPC for its centralized multivariate structure and constraint-handling capabilities. Given that our primary objective is economics-based, the predictive controller here is distinct from a conventional MPC controller in that its objective is not to track given setpoints, but instead to implement manipulated variable adjustments to optimize an economic objective (as well as satisfy lower tier objectives like minimizing deviation of product qualities from their targets, etc.).

In our scheme, the MPC-based shutdown controller is activated at the beginning of a shutdown (t_{start}), and takes over from the nominal control system. We provide a description of the steps carried out by the MPC algorithm below. Figure 3 summarizes the shutdown controller algorithm. Note that it is important to distinguish between the model's discrete time counter k, and the MPC iteration counter j.

1. **Downtime Estimate** (d^p_{est}). At the beginning of the shutdown, a value for the downtime estimate, $d^p_{est} = t^{est}_{end} - t_{start}$ (where t^{est}_{end} is the estimated shutdown end time) is supplied to the shutdown controller. This downtime estimate parameter can be manually updated during any control interval as new information about the downtime becomes available. This is an instance of manual parameter feedback. The values of α^p_{shut} for all process units $p \in \mathcal{P}$ are determined from this parameter.

2. **Open-loop Dynamic Optimization**. The MPC dynamic optimization problem (in the "MPC Shutdown Controller" block) is solved using a multitiered optimization approach.

3. **Implement Control Inputs on Plant**. Only the current step (time j) of the calculated inputs $\left(u_{k=j}\right)$ is implemented on the plant.

4. **State Measurements**. Plant state measurements at the next sample-time are taken. In a simulation environment, a 1-step integration is performed from t_j to t_{j+1} (in the "Plant DAE (One-Step Integration)" block); where x^p, z^p = plant model variables; θ^p = plant model parameters; f^p, h^p = plant DAE model functions. The internal states of the plant, $\bar{x}^p_j$ are the initial values for the DAE. The vector y^m_{j+1} is conveyed back to the predictive controller, where it is used as initial values for the next iteration. In this work, we assume full-state feedback. If full state measurements are not available, a state observer is employed.

5. **Iterate**. The counter j is incremented. This process is repeated from Step 1 until the system is restored to its nominal steady state operation (at time t_{res}, see Figure 1), after which control is handed back to the nominal control system (when $j = N$).

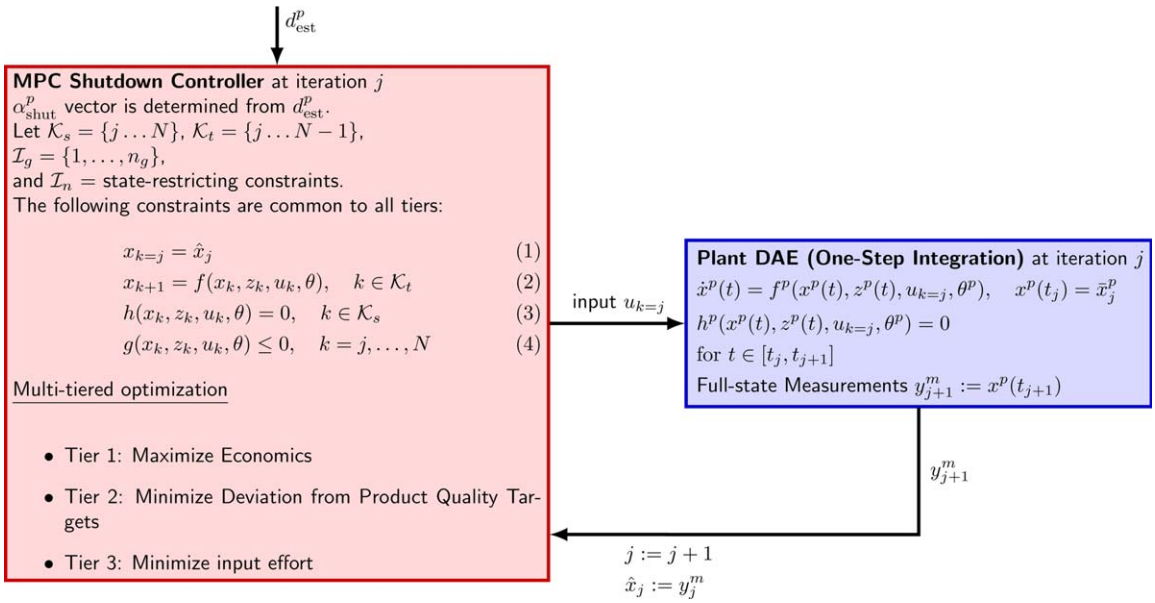

Figure 3. MPC simulation workflow.

Iterations terminate when $j = N$. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com].

The transient processes we are considering (namely partial shutdowns) are finite in duration, therefore, the prediction/control horizon length in the controller is finite. Instead of receding prediction and control horizons found in conventional MPC, shrinking prediction and control horizons are used. (The prediction and control horizon lengths decrease as the controller advances toward the end of the time horizon.) Here, a shrinking horizon is selected for reasons of computational efficiency and ease of problem initialization, but in practice, a receding horizon is also admissible.

The horizon length N needs to correspond to the duration of the shutdown and restoration combined. If a shorter length is chosen, suboptimal control may result from the controller not having an adequate picture of the full transient process. Also, in our predictive controller, the control horizon length equals the prediction horizon length. Economics optimization is integrated directly into the predictive controller, which constitutes a 1-layer MPC with economics approach (in contrast to a layered Dynamic RTO approach).

We note that the equality constraints are enforced from the initial point in order that the values of the algebraic states may be computed.

Remark. The estimated duration of the shutdown $\left(d^p_{\text{est}} = t^{\text{est}}_{\text{end}} - t_{\text{start}}\right)$ is a important parameter in calculation of the optimal trajectories because it influences the shape of the resulting trajectories. The downtime estimate would typically be provided by the operator to the control system, based on past operational experience or direct information about the prognosis of the shutdown. In practice, this estimate will not correspond exactly to the actual downtime (d_{act}) for various reasons, not least due to the fact that it may be very difficult in most cases to make an accurate prediction of a unit's downtime. One approach to addressing this issue is to treat the downtime estimate as a parametric feedback variable. In contrast to other feedback variables, the downtime is a single parameter manually supplied by maintenance personnel. By adopting this approach, the operator is allowed to revise any downtime estimates dynamically.

Application Study

Process description

The model used for this partial shutdown study is the fiber line of a Kraft pulp and paper plant. This study was inspired by an actual industrial problem in a plant in Ontario, Canada. The Kraft fiber line is composed of various production departments that are separated by buffer capacities. The major departments considered in this work are digestion, knotting, washing, screening, and delignification.

The model consists primarily of material component balances. Three components are tracked: pulp (P), dissolved solids (DS), and water (W). We assume quasi-steady-state operation for all operating units except the buffering tanks, which we justify on the basis of the large inventory capacities typical of such plants. The level of modeling detail was chosen to ensure that the macroeffects germane to our application were adequately captured (that is, inventories and compositions). The mathematical equations of the model can be found in Chong.[18] This model is based on work done by Balthazaar[11] and Dubé,[10] with some extensions based on various literature sources.

A simplified schematic of the plant in shown in Figure 4. The variables F denote flowrates (ton/h), while H are levels (%) and x are compositions (%). There are two major buffering tanks, the blowtank and the storage tank located before the delignification department. There are also buffering tanks (sealtanks) in the washing department which are used for internal recycles. Each variable name is tagged with a superscript indicating the unit it is associated with. The mapping between the tag names and their corresponding unit names is found in Table 1.

A digester vessel is filled with wood chips and white liquor (a concentrated solution of Na_2S and NaOH) and cooked to form pulp. The cooked pulp is then subjected to various physical and chemical separation processes designed to remove unprocessable wood, residual black liquor, and lignin.

The top section of our digester model is based on the model developed by Wisnewski and Doyle.[19] As lignin is

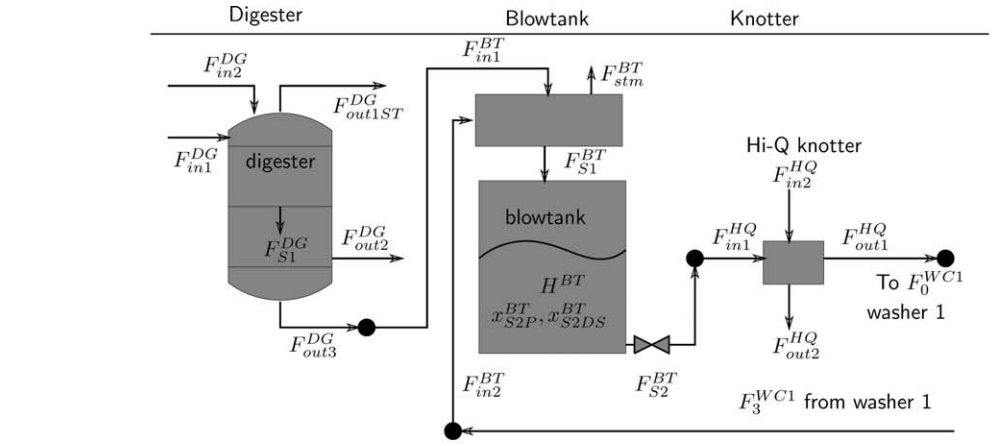

(a) Digestion and knotting

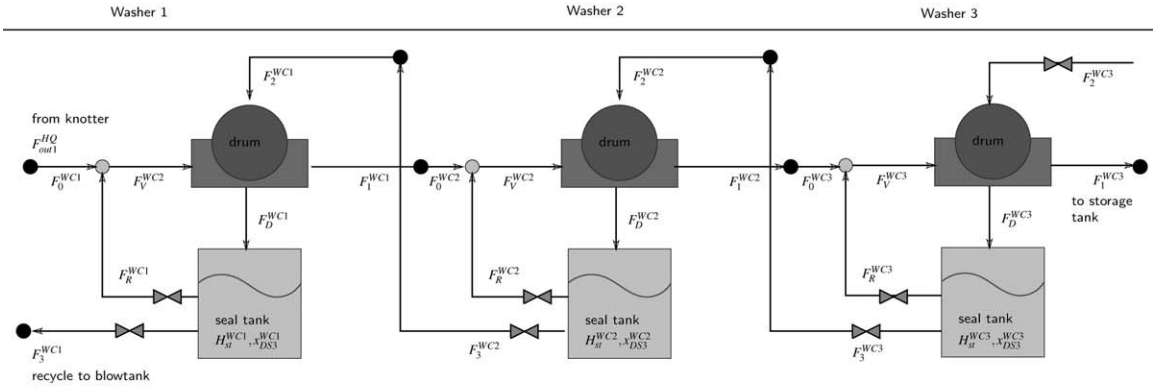

(b) Brownstock washing

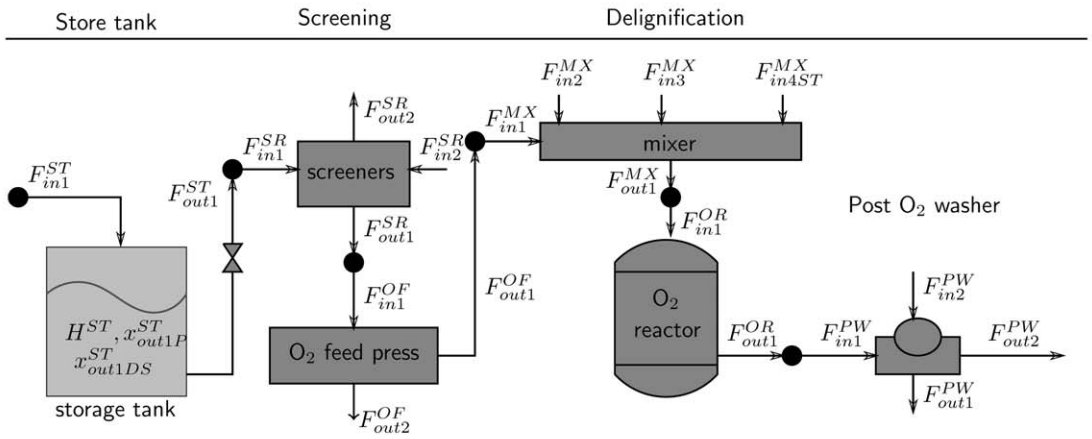

(c) Screening and delignification

Figure 4. Simplified schematic of plant model.

removed from the cellulose fibers in the cooking process, there is a loss in pulp mass that needs to be accounted for. This loss is captured using a defined quantity, "pulp shrinkage." The pulp shrinkage is related to another quantity ζ, the production factor. A cubic relationship between these two variables was developed in Balthazaar's[11] work. In our plant, the shrinkages were almost linear in the range of feasible operation, therefore, we present the linearized versions of the equations here. The pulp shrinkage regression equations are given below:

$$\zeta = \frac{F_1^{\mathrm{DG}}}{F_{\max}} \tag{32}$$

$$\alpha_{\mathrm{shrink},1} = \frac{-3.94 \cdot \zeta + 14.59}{100} \tag{33}$$

$$\alpha_{\mathrm{shrink},2} = \frac{-3.21 \cdot \zeta + 12.15}{100} \tag{34}$$

where F_1^{DG} is the flowrate into the digester (see Figure 4a), $F_{\max}$ is its upper bound, $\alpha_{\mathrm{shrink},1}$ is the pulp shrinkage for shrinkage occurring in the top portion of the digester, while

Table 1. Variable Superscript Tag Names and Corresponding Unit Names

Tag	Unit Name
DG	Digester
BT	Blowtank
HQ	Hi-Q knotter
WC1	Brownstock washer 1
WC2	Brownstock washer 2
WC3	Brownstock washer 3
ST	Storage tank
SR	Screeners
OF	O_2 feedpress
MX	Mixer
OR	O_2 delignification reactor
PW	Post-O_2 washer

$\alpha_{\text{shrink},2}$ is a similar quantity for the bottom part of the digester.

Simple mass balances govern the material flow in the different portions of the digester model. It is a common practice in industry to maintain a constant liquor-to-wood flowrate ratio.[20] The r_{lw} (liquor-to-wood ratio) parameter, which regulates the proportion of liquor $\left(F_{\text{in2}}^{\text{DG}}\right)$ used relative to the amount of wood chips fed $\left(F_{\text{in1P}}^{\text{DG}}\right)$, is assumed to be 3.6 by mass (Dubé,[10], Grace[20]).

$$F_{\text{in2}}^{\text{DG}} = r_{\text{lw}} \cdot F_{\text{in1P}}^{\text{DG}} \tag{35}$$

The knotting department consists of knotter machines that function to remove undigested chips and ill-sized wood pieces, or "knots," which hinder downstream processing. In the subsequent washing operation, the residual black liquor is separated from the pulp by feeding the pulp into a series of counter-current vacuum drum washers, in which black liquor is displaced. The counter-current washing circuit model is inspired by work by Wang.[21] Each drum washer unit is modeled as a quasi-steady-state state unit and seal tanks as dynamic units.

The displacement ratio (D_{R}) measures the washing efficiency of a washer in terms of percent dissolved solids removal (Noël et. al.,[22]). This ratio may be obtained by collecting samples on the washers. The displacement ratio is regressed against the wash-liquor ratio R_{W} using data from Grace et al.[20]:

$$D_{\text{R}} = \frac{x_{\text{DS1}} - x_{\text{DS2}}}{x_{\text{DSV}} - x_{\text{DS2}}} = a_1 R_{\text{W}}^3 + 0.1 a_2 R_{\text{W}}^2 + a_3 R_{\text{W}} \tag{36}$$

$$R_{\text{W}} = \frac{F_2}{F_1} \tag{37}$$

where $x_{\text{DS},i}$ are dissolved solid compositions in stream i of each washer, F_2 is the flowrate of the shower water stream, and F_1 is the flowrate of the outlet stream (see Figure 4b). The seal tank serves to create sufficient vacuum through the drop leg for proper washer operation. It is also used to store filtrate necessary for startup and functions as a buffer as well. The seal tank can be modeled simply as a continuously stirred tank.

The washed pulp is then conveyed to the screening department, to be prepared for delignification–the process by which the remaining lignin in the pulp is removed. The pulp stream is first mixed with chemical streams containing caustic soda and magnesium sulfate. This mixture is then fed in a counter-current direction with respect to an oxygen stream running within an O_2-delignification reactor. A reaction occurs in which the lignin separates from the cellulose fibers in the pulp stream. The delignified pulp stream is then ready for bleaching, where the pulp is whitened in a chlorination reactor.

The buffer tanks are modeled as simple mixing tanks. For example, the dynamic equations for the storage tank are as follows (see Figure 4c):

$$M_{\text{tank}} \cdot 0.01 \cdot \frac{d}{dt} H^{\text{ST}} = F_{\text{in1}}^{\text{ST}} - F_{\text{out1}}^{\text{ST}}, \qquad H^{\text{ST}}(0) = H_0^{\text{ST}} \tag{38}$$

$$M_{\text{tank}} \cdot 0.01 \cdot H^{\text{ST}} \cdot \frac{d}{dt} x_{\text{out1P}}^{\text{ST}} = F_{\text{in1P}}^{\text{ST}} - F_{\text{in1}}^{\text{ST}} \cdot x_{\text{out1P}}^{\text{ST}}, \qquad x_{\text{out1P}}^{\text{ST}}(0) = x_{\text{in1P}}^{\text{ST}}(0) \tag{39}$$

$$M_{\text{tank}} \cdot 0.01 \cdot H^{\text{ST}} \cdot \frac{d}{dt} x_{\text{out1DS}}^{\text{ST}} = F_{\text{in1DS}}^{\text{ST}} - F_{\text{in1}}^{\text{ST}} \cdot x_{\text{out1DS}}^{\text{ST}}, \qquad x_{\text{out1DS}}^{\text{ST}}(0) = x_{\text{in1DS}}^{\text{ST}}(0) \tag{40}$$

where M_{tank} is the tank capacity (mass), H^{ST} is tank level, specified as a percentage by mass of the total tank capacity (%). The inlet and outlet flowrates are F_{in1} and F_{out1}, and $F_{\text{in1,j}}$ for $j=\{P, \text{DS}\}$ are component flowrates of pulp and dissolved solids in stream F_{in1}, respectively. Similarly, $x_{\text{out1,j}}$ for $j=\{P, \text{DS}\}$ are the compositions of pulp and dissolved solids in stream F_{out1}, respectively.

When a process unit shuts down, an entire department is typically forced to shut down unless there is internal buffering capacity in the departments. A process unit is shut down and taken offline for a period of time, and is subsequently restored. Based on an estimate of the downtime (specified by the operator), the proposed control scheme is used to compute and implement a set of optimal control trajectories that accommodates the shutdown in the presence of nonlinear dynamics and recycles.

The manipulated variables and their bounds are given in Table 2. We assume full-state measurements are available. The sets of independent states and inputs required to specify the steady state at the end of the restoration period are:

$$\mathcal{I}_{\text{ss}}^{x} = \{i | x_i \in \mathcal{X}_{\text{ss}}\} \tag{41}$$

$$\mathcal{X}_{\text{ss}} = \{H^{\text{BT}}, H^{\text{WC1}}, H^{\text{WC2}}, H^{\text{WC3}}, H^{\text{ST}}\} \tag{42}$$

$$\mathcal{I}_{\text{ss}}^{u} = \{i | u_i \in \mathcal{U}_{\text{ss}}\} \tag{43}$$

$$\mathcal{U}_{\text{ss}} = \{F_{\text{in1}}^{\text{DG}}, F_{\text{R}}^{\text{WC1}}, F_{3}^{\text{WC1}}, F_{2}^{\text{WC1}}, F_{\text{R}}^{\text{WC2}}, F_{\text{R}}^{\text{WC3}}, F_{3}^{\text{WC3}}\} \tag{44}$$

The capacities of the buffer tank inventory units and their nominal levels are given in Table 3. Table 4 lists product

Table 2. Decision Variables and Their Bounds

Manipulated Variables	Min (ton/h)	Max (ton/h)
$F_{\text{in1}}^{\text{DG}}$	0	40
$F_{\text{S2}}^{\text{BT}}$	0	600
F_{3}^{WC1}	0	600
$F_{2}^{\text{WC1}} = F_{3}^{\text{WC2}}$	0	600
$F_{\text{R}}^{\text{WC1}}$	0	600
$F_{\text{D}}^{\text{WC1}}$	0	600
$F_{2}^{\text{WC2}} = F_{3}^{\text{WC3}}$	0	600
$F_{\text{R}}^{\text{WC2}}$	0	600
$F_{\text{D}}^{\text{WC2}}$	0	600
F_{2}^{WC3}	0	600
$F_{\text{R}}^{\text{WC3}}$	0	600
$F_{\text{D}}^{\text{WC3}}$	0	600
$F_{\text{out1}}^{\text{ST}}$	0	600
$F_{\text{out1}}^{\text{SR}}$	0	600
$F_{\text{out1}}^{\text{OF}}$	0	600
$F_{\text{out2}}^{\text{PW}}$	0	600

Table 3. Tank Names, Capacities, and Nominal Inventories

Tank Name	Capacity, V_{tank} (m^3)	Nom. Inven. (m^3/% of capacity)
Blowtank	2050.0	1230.0/60%
Sealtank #1	1067.0	248.6/23.3%
Sealtank #2	861.0	288.4/33.5%
Sealtank #3	683.0	227.4/33.3%
Storage tank	1620.0	972.0/60%

quality variables and their targets. This work uses the simultaneous method for dynamic optimization, where both controls and states are discretized. The model contains 217 inequality constraints, and 181 equations–of which 12 are differential equations. Orthogonal collocation on finite elements was used for discretization, resulting in an optimization problem (after presolve) on the order of 17,200 variables, 16,300 equality constraints, and 1000 inequality constraints.

The economic objective function is a linear function of the cumulative profit over the horizon of interest:

$$\Phi_{\text{econ}} = \Delta t \sum_{m=1}^{n_m} \left[C_m \left(\sum_{k=0}^{N-1} F_{m,k} \right) \right] \tag{45}$$

where Δt is the sample time, F_m is the flowrate (in ton/h) of some component m, C_m is the price that component ($/ton), and n_m is the total number of components. The prices C_m can be found in Table 5.

The solvers used in this work were IPOPT 3.8.2[24] (a large-scale sparse solver that implements a primal-dual interior point method) and CONOPT 3.14G[25] (a feasible path active-set solver). IPOPT was used in the open-loop studies, as well as to generate an initial guess for the first iteration of the shutdown controller. CONOPT was used for closed-loop studies due to its amenability to warm-starts. Its ability to follow a feasible path is also particularly advantageous for efficiently solving multiple multitiered optimization problems, as the solution for one tier is a feasible initial point for the next.

Case 1: Nominal open-loop optimization of shutdowns

Description. In this case study, we demonstrate the application of aforementioned methodologies for the dynamic optimization of three separate shutdown scenarios. In Scenario (a), the O_2 Delignifier unit is shut down, and Scenario (b) involves the shutdown of the Hi-Q knotting unit, while in Scenario (c), the digester unit goes offline.

The plant is assumed to be operating at steady state initially, and the shutdown occurs suddenly. Flowrates into the shut down unit (as well as adjacent quasi-steady-state units) dip to 0. Past the shutdown period, the optimizer responds by prescribing input actions to restore the plant to its steady state.

Table 4. Product Quality Variables and Their Targets

Product Quality Variable	Description	Target Value (%)
C_M^{SR}	Outlet consistency of screeners	4.5
C_M^{OF}	Outlet consistency of O_2 feedpress	30.0
C_M^{PW}	Outlet consistency of Post-O_2 washer	6.5

Table 5. Average Prices of Materials[11,23]

Material	Prices, C_m ($/ton)
Pulp	725.00
Chips	−25.00
Chemicals	−100.00
Energy from black liquor	0.35

In each case below, the simulation horizon is 12 h and the shutdown duration is 5.5 h. For this initial case, we assume we know the downtime exactly and that there are no uncertainties. The details are summarized in Table 6.

Results. The results of the three scenarios are as follows:

• ***Scenario (a): Shutdown in the O_2 delignifier.*** The overall economics achieved was $92,963. The trajectories are represented by the dashed lines in Figure 5. The flowrate out the storage tank (upstream of the delignifier, see Figure 4c), F_{out1}^{ST} (Figure 5m) goes to 0 during the shutdown interval $t=0.25-5.75$ h. This has the effect of shutting down the O_2 delignifier and all the surrounding units (up to the screening unit), as there are no buffer capacities in units between it and the storage tank.

As the shutdown is progressing, we notice a rise in the storage tank level H^{ST} (Figure 5h), due to accumulation of material from upstream units. The digester outlet stream F_{in1}^{DG} (Figure 5i) continues to produce material at its nominal rate, while the outlet flowrate in the blowtank F_{S2}^{BT} (Figure 5j) drops to accommodate the constriction downstream. This causes the rise in the blowtank level H^{BT} (Figure 5a).

After the shutdown period is over $(t > 5.5 \text{ h})$, the system attempts to make up for lost production by discharging the contents of the storage tank, as seen in a decline in its level H^{ST}. Two product quality variables C_M^{SR} (Figure 5n) and C_M^{PW} (Figure 5o) (which are associated with the screening and post-O_2 washing units) can be seen to be deviating from their targets during the restoration period, in response to the material influx during the restoration period. All inputs and states return to their original steady-states post-restoration.

Because a shutdown in the O_2 delignifier is downstream of all the major production units and its adjacent buffering tank is capable of absorbing the majority of the product being produced upstream, the overall effect of this partial shutdown on the plant is less severe than that of partial shutdowns in units upstream. This is corroborated by the fact that the economics achieved is considerably higher than that of the next two cases.

• ***Scenario (b): Shutdown in the Hi-Q knotter.*** The overall economics achieved was $65,292, which is lower than that of a partial shutdown in the downstream delignifier unit above.

Table 6. Case 1 Summary, Scenarios (a), (b), and (c)

Case 1 Summary (Figures 5)	
Units shut down	Shutdown scenarios (a) O_2 delignifier; (b) Hi-Q knotter; (c) digester.
Simulation horizon (h)	12
Shutdown start time (h)	0.25
Estimated downtime, d_{est} (h)	5
Actual downtime, d_{act} (h)	5
Reoptimization point (h)	N/A
Disturbances	None

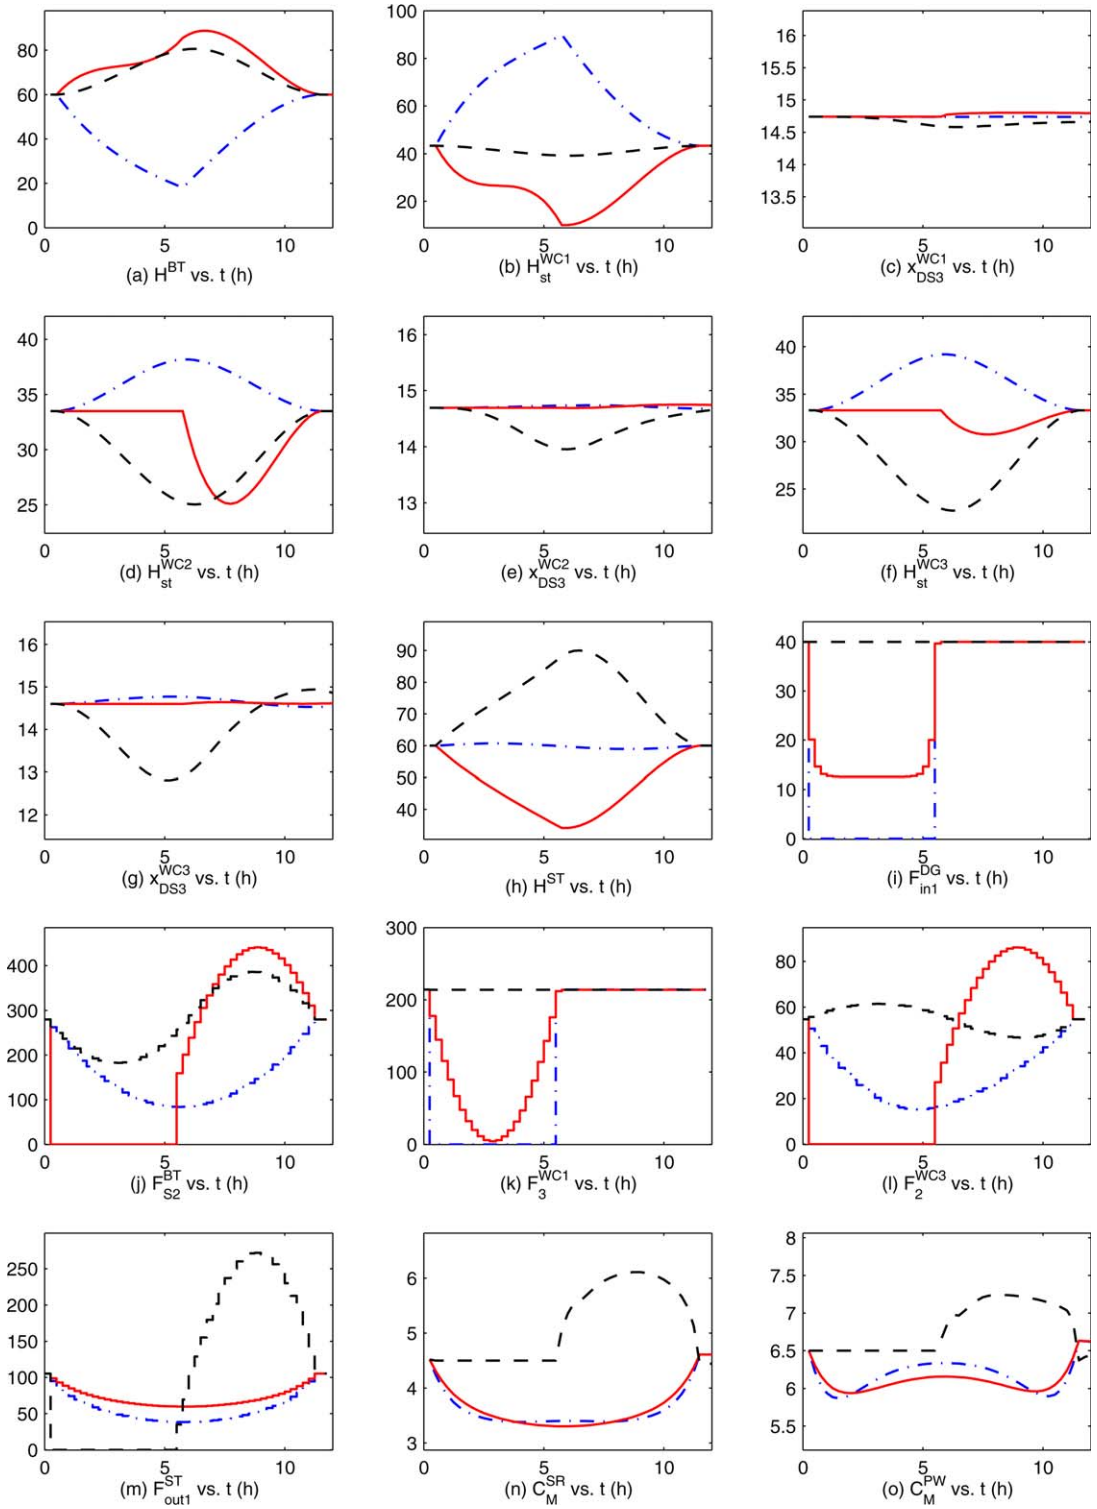

Figure 5. Case 1: Open-loop dynamic optimization.

Solution trajectories for three distinct shutdowns. Legend: Dashed lines = Scenario (a) delignifier shutdown. Solid lines = Scenario (b) Hi-Q knotter shutdown. Dotted dashed lines = Scenario (c) digester shutdown. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com].

The trajectories are represented by the solid lines in Figure 5. This shutdown is indicated by the turning off of flows to all units bounded by the blowtank and the storage tank, that is, from the Hi-Q knotter to Washer 3. The seal tanks of the washers however remain operational as buffer capacities. The blowtank outlet flowrate F_{S2}^{BT} (Figure 5j) is seen to go to 0. All internal recycles in the washing circuit are also at 0 during the shutdown. However, the main recycle stream from Washer 1 seal tank to the blowtank, F_3^{WC1} (Figure 5h) continues to be a degree-of-freedom, therefore the level of the Washer 1 seal tank, H_{st}^{WC1} (Figure 5b) can be seen to be changing during the shutdown.

During the shutdown period, production continues both upstream and downstream of the blowtank and storage tank,

Table 7. Case 2 Summary

Case 2 Summary	Scenario (1), $d_{est} > d_{act}$, Figure 6	Scenario (2), $d_{est} < d_{act}$, Figure 7
Units shut down	Hi-Q knotter	Hi-Q knotter
Simulation horizon (h)	12	12
Shutdown start time (h)	0.25	0.25
Estimated downtime, d_{est} (h)	3.5	3.5
Actual downtime, d_{act} (h)	2.5	4.5
Reoptimization points (h)	0.25, 1.25, 2.25 h	0.25, 1.25, 2.25, 2.75 h
Disturbances	None	None

respectively. This is seen from a rise in H^{BT} (Figure 5a) and a dip in H^{ST} (Figure 5h). The production rate of the digester also drops during the shutdown, as seen in F_{in}^{DG} (Figure 5i). This is to prevent overwhelming the blowtank, whose level H^{BT} has an upper-bound of 90%, with material exiting the digester. The level of the Washer 1 seal tank H_{st}^{WC1} (Figure 5b) also drops during the shutdown period because its contents are recycled to the blowtank, via recycle stream F_3^{WC1} (Figure 5k).

The restoration trajectories for F_{S2}^{BT} (Figure 5j) and F_2^{WC3} (Figure 5l) reveal the general strategy of the optimizer. During the restoration period, the optimizer is pursuing a course of action in which the flowrates are initially increased by a significant amount to increase production. However, in order to meet endpoint constraints, the optimizer quickly ramps down the said flowrates and settles at the steady-state values. As a result, the inventories in the tank are then built back up. The dissolved solids compositions in the washer seal tanks, namely $x_{DS3}^{WC1}, x_{DS3}^{WC2}, x_{DS3}^{WC3}$ (Figures 5c, e, g) do not appear to deviate significantly from their steady-state values, which indicates that the shutdown did not introduce signification variation to the compositions in the counter-current washing circuit.

For this partial shutdown case, the product qualities did not deviate significantly from the target. This is the result of two things: (1) the product quality optimization tier minimizes all deviations from target values; (2) the sizable capacity of the storage tank minimizes the propagation of variation in pulp consistency to downstream units.

• ***Scenario (c): Shutdown in the digester.*** The overall economics achieved was \$52,286, which is worse than the above two cases. This can be attributed to the fact that the digester is a key unit in the process, and the pulp it generates is the primary component for processing.

The overall strategy is similar in principle to what is described in the previous case. The trajectories are represented by the dotted-dashed lines in Figure 5. The digester is the first unit in this plant, therefore in this scenario we expect a greater disruption in the plant than the above two cases.

The chip feed to the digester, F_{in1}^{DG} (Figure 5i) is turned off during the shutdown period. We assume that the combined duration of the shutdown and restartup procedures are lumped into this downtime. The strategy is to keep all the units downstream of the blowtank operational while the digester is down. The implementation of this strategy can be seen by a decline in the level of H^{BT} (Figure 5a), and a gradual increase during its subsequent restoration.

During the shutdown, the washers continue to operate. However, the liquor recycle to the blowtank from the washing circuits is turned off as seen in the figure for F_3^{WC1} (Figure 5k). The inventories of the 1st and 3rd sealtanks build up as a result, as observed in H_{st}^{WC1} and H_{st}^{WC3} (Figures 5b, f). The pulp consistencies downstream can be seen to drop, particularly in the screeners (C_M^{SR}, Figure 5n) and post-O_2 washers (C_M^{PW}, Figure 5o), due to the loss of pulp output from the digester.

In the restoration phase, the inventories are being balanced in such a way as to enable the system to reach its original steady state at the end of the restoration period.

Case 2: Uncertainty in downtime estimate, d_{est} (MPC mode)

Description. This case study demonstrates the use of reoptimization as a feedback mechanism for updating downtime estimates. The downtime estimate is an uncertain quantity that often cannot be accurately determined from the outset. Over the course of time, as more information becomes available, the downtime estimate can be updated and a reoptimization is performed, and a midcourse correction strategy is implemented.

We will consider two scenarios. In Scenario (1), the actual downtime d_{act} is shorter than the initial estimate d_{est}. In Scenario (2), the actual downtime d_{act} is longer than the initial estimate d_{est}. In each scenario, we investigate the effect of updating the downtime estimate at different times, that is, we assume that at a chosen reoptimization point, the operator receives a corrected downtime estimate from the maintenance personnel. The estimate is entered into the control system, and the MPC algorithm reoptimizes the remainder of the trajectories based on this new estimate. In this study, we only update the downtime estimate once per scenario, but in principle, it can be done as many times as required. The attributes of the two scenarios are summarized in Table 7. Given that there are three reoptimization points in Scenario (1), and four reoptimization points in Scenario (2), seven independent subcases are carried out.

Results. The economics obtained in each scenario subcase (at each reoptimization point) are listed in Table 8.

Several observations may be gleaned from the results. The ideal results correspond to the subcases in each scenario where the reoptimization time is $t = 0.25$ h. This is because the shutdown starts at $t = 0.25$ h, therefore, if the correct downtime is given to the optimizer at this point, it can take the necessary corrective actions without any hindrances.

First, we observe that the longer downtime corresponds to poorer economic outcomes, which is consistent with what we would intuitively expect. Second, one observes that the economic outcomes generally become worse the later the reoptimization occurs. This is because the outcome of a

Table 8. Economics of Scenarios (a) and (b).

Scenario Subcases	d_{est}	d_{act}	Reopt Time (h)	Economics, Φ_{econ}^* (\$)
Scenario (1a), $d_{est} > d_{act}$ (ideal)	3.5	2.5	0.25	84,590
Scenario (1b), $d_{est} > d_{act}$	3.5	2.5	1.25	84,588
Scenario (1c), $d_{est} > d_{act}$	3.5	2.5	2.25	84,574
Scenario (2a), $d_{est} < d_{act}$ (ideal)	3.5	4.5	0.25	69,158
Scenario (2b), $d_{est} < d_{act}$	3.5	4.5	1.25	69,155
Scenario (2c), $d_{est} < d_{act}$	3.5	4.5	2.25	69,153
Scenario (2d), $d_{est} < d_{act}$	3.5	4.5	2.75	69,150

The ideal case is when the actual downtime is known at the start of the shutdown (at $t = 0.25$ h), and a reoptimization is performed at that point.

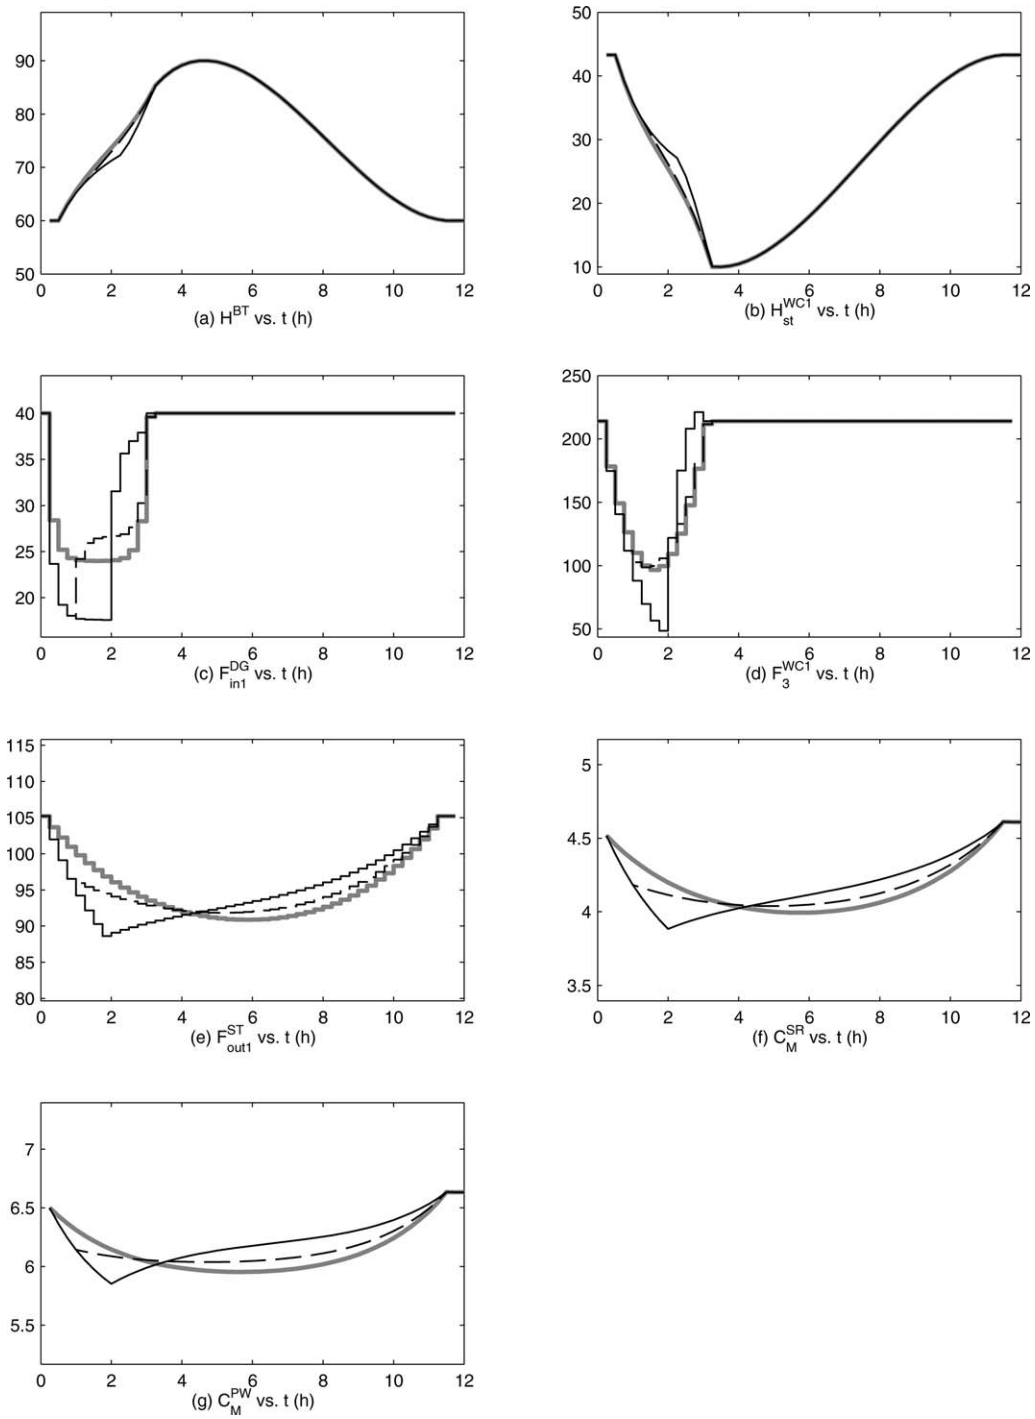

Figure 6. Closed-loop results. Scenario (1), Downtime estimate reoptimization.

Hi-Q knotter shutdown (duration d_{act} = 2.5h < d_{est} = 3.5h). Reoptimization at t = 0.25, 1.25, 2.25 h. Legend: Thick gray lines = (1a) reoptimized at 0.25 h. Dashed lines = (1b) reoptimized at 1.25 h. Thin solid lines = (1c) reoptimized at 2.25 h.

reoptimization depends on the degrees-of-freedom available at the point of reoptimization. It stands to reason that if the plant had been operated longer with the incorrect downtime, it may have moved to a point that is disadvantaged vis-à-vis the operating point of a plant that had been reoptimized earlier with a correct downtime. Further, owing to the fact that our controller has a finite-shrinking horizon, a later reoptimization step finds itself with few remaining degrees-of-freedom to adjust.

Figure 6 shows the trajectories obtained from the subcases of Scenario (1). Due to the large sizes of the buffer tanks, the differences in the tank level trajectories H^{BT} and H_{st}^{WC1} (Figures 6a, b) for the three subcases are small (though visible), despite the differences in their respective reoptimization times.

However, for input trajectories—notably F_{in1}^{DG}, F_{3}^{WC1} and F_{out1}^{ST} (Figures 6c, d, e)—the differences between the scenarios are much more significant. The shape of the input trajectories corresponding to Scenario (1c) in particular depart visibly from the previous two scenarios. This is because the reoptimization point in Scenario (1c) occurs much later, therefore, trajectories based on an incorrect (longer) downtime have already been implemented up to that point.

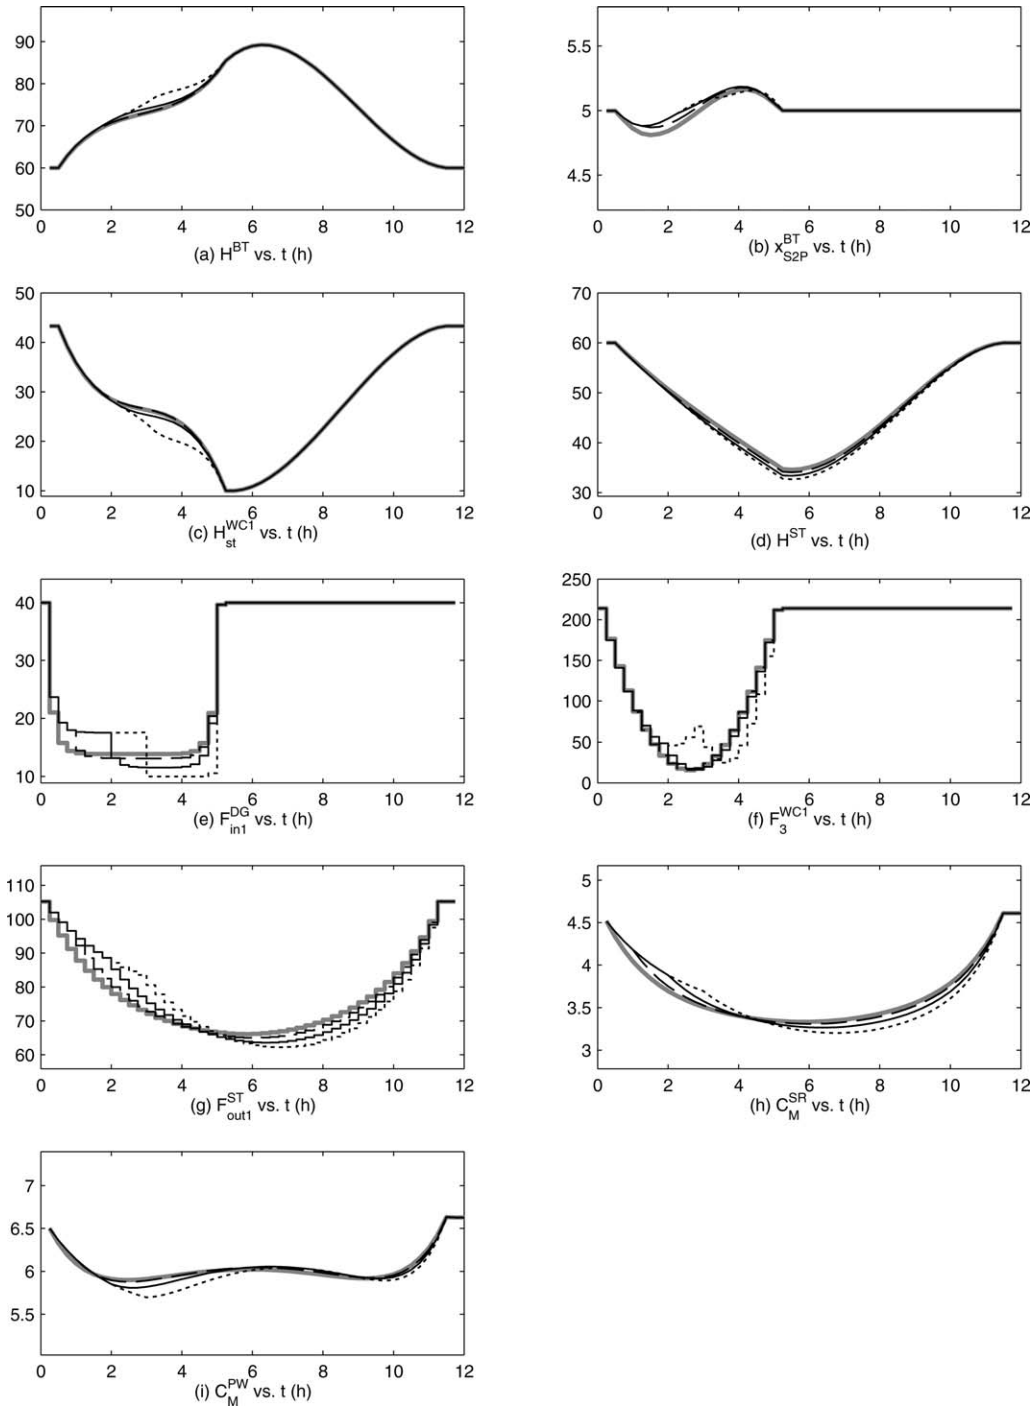

Figure 7. Closed-loop results. Scenario (2), Downtime estimate reoptimization.

Hi-Q knotter shutdown (duration $d_{act}=4.5h>d_{est}=3.5h$). Reoptimization at $t=0.25, 1.25, 2.25, 2.75$ h. Legend: Thick gray lines = (2a) reoptimized at 0.25 h. dashed lines = (2b) reoptimized at 1.25 h. Thin solid lines = (2c) reoptimized at 2.25 h. Dotted lines = (2d) reoptimized at 2.75 h.

In Scenarios (1b) and (1c), since the controller originally assumed that the downtime was longer than it actually is, F_{in1}^{DG}, F_3^{WC1} and F_{out1}^{ST} (Figures 6c, d, e) were adjusted in such a way as to simultaneously decrease the amount of material entering the blowtank (to prevent it from hitting its upper level bound) and slowing down the rate of the amount of material being discharged from the storage tank. The product quality variables, namely C_M^{SR} and C_M^{PW} (Figures 6f, g)—the consistencies of the pulp exiting the screeners and the post-O_2 washers—reflect the fact that due to the initially assumed longer downtime, a larger excursion from the targets was permitted. As soon as reoptimization occurs, and the correct downtime was set, the product qualities are driven back to their nominal values more rapidly.

In Scenario (2), depicted in Figure 7, the initial (incorrect) estimated downtime is the same as in Scenario (1), but now the actual downtime is longer than the estimated one. Similar input adjustments are made prior to reoptimization. Differences between the trajectories for F_{in1}^{DG}, F_3^{WC1}, and F_{out1}^{ST} (Figures 7e, f, g) are observed for different reoptimization times.

Table 9. Case 3 Summary

Case 3 Summary (Figure 8)	
Units shut down	Hi-Q knotter
Simulation horizon (h)	12
Shutdown start time (h)	0.25
Estimated downtime, d_{est} (h)	5
Actual downtime, d_{act} (h)	5
Disturbances	Actual F_2^{WC3} stream flowrate in plant is 30% higher than in model, between $t = 6.25$–8.25 h.

If we compare the trajectories for Scenarios (2a) and (2d), we note that in the latter case—in which reoptimization occurs later—the recycle F_3^{WC1} (Figure 7f) to the blowtank is throttled down at the reoptimization point to avoid overwhelming the blowtank. The state trajectories for H^{BT}, x_{S2P}^{BT}, H_{st}^{WC1}, and H^{ST} (Figures 7a–d) can all be seen to evolve slightly differently in the different subcases, as do the product quality variables C_M^{SR} and C_M^{PW} (Figures 7h, i). We remark that for this particular plant/economic configuration, the penalty (in dollar terms) of not knowing the downtime duration accurately from the onset appears to be minimal.

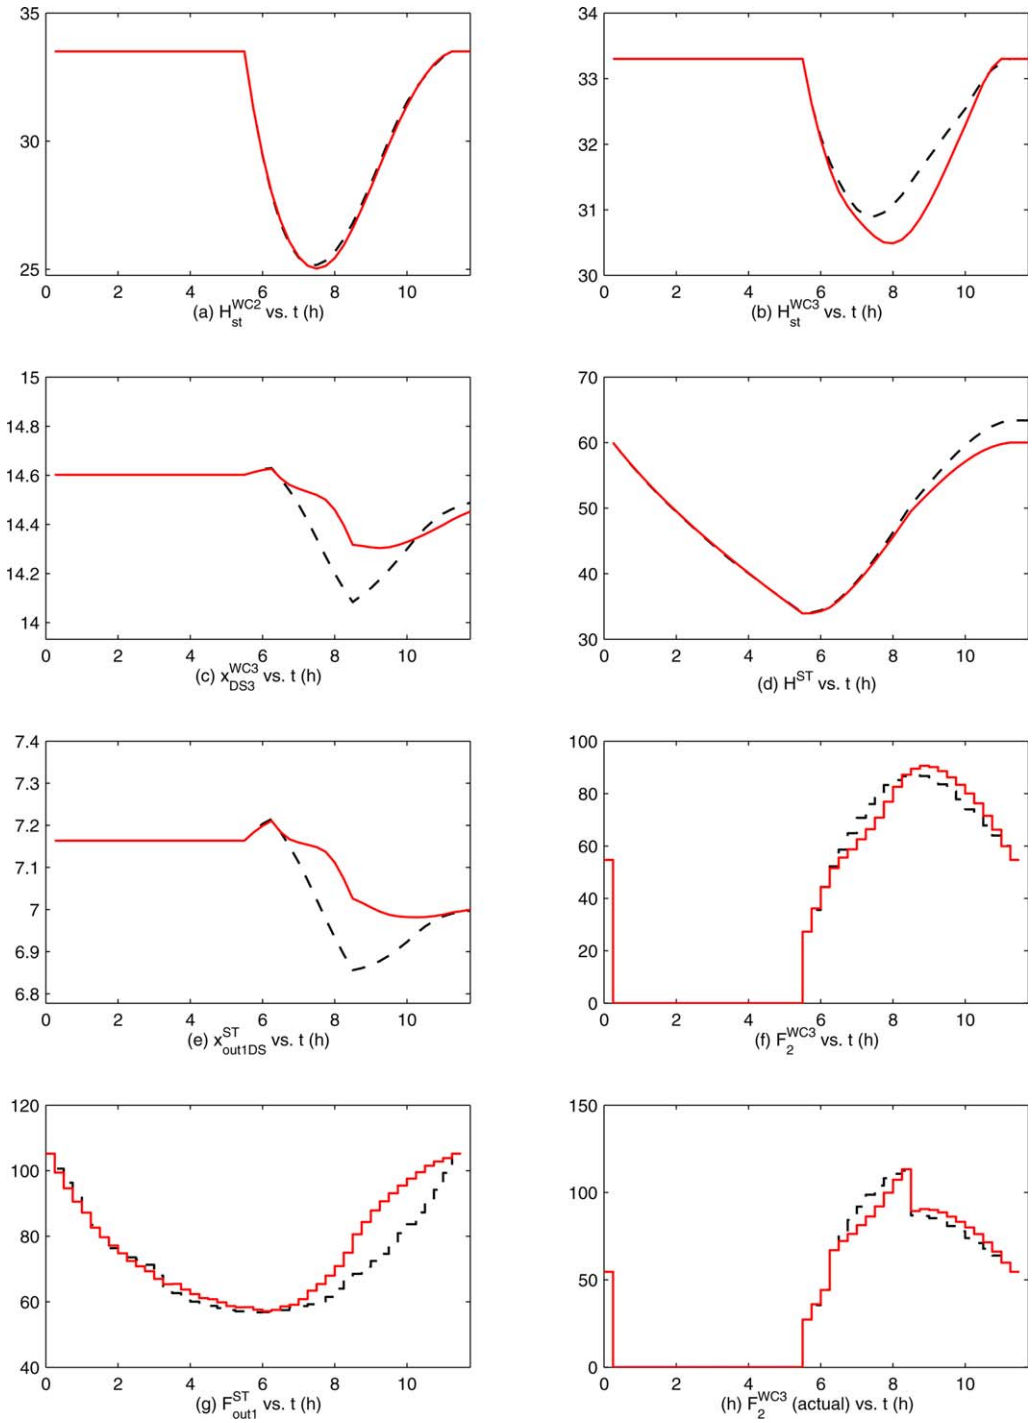

Figure 8. Closed-loop results with a disturbance entering the Washer 3 shower stream F_2^{WC3} between $t = 6.25$–8.25 h.

Legend: Dashed lines = plant trajectories from open-loop solution (no feedback). Solid lines = plant trajectories from closed-loop solution (with feedback). [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com].

However, other plant/economic configurations may yield different results.

In general, parametric feedback of the downtime estimate allows the operators to take corrective action, which in turn leads to smaller losses in economics owing to the way the plant is operated during a shutdown.

Case 3: Closed-loop study with unmeasured disturbances

Description. In this case study, we compare the implementation of an open-loop dynamic optimization solution (with no feedback) with that of a closed-loop solution (with feedback). In the former case, we assume that an open-loop dynamic optimization tool was used to generate the desired input trajectories, that the trajectories are used as a recipe for plant operation, and that the implementation is devoid of any closed-loop correction mechanism. The purpose of this study is to compare the performance of recipe-driven operation with an actual feedback control system on a plant.

A shutdown controller with full-state measurements is used to compute and implement the closed-loop solution. An unmeasured step disturbance in the shower stream of Washer 3, F_2^{WC3} enters the plant between $t = 6.25$–8.25 h. During this period, the actual flowrate of the shower stream that enters Washer 3 is 30% higher than modeled. No noise or plant-model mismatch (except for the disturbance) is present. Table 9 below summarizes this case.

Results. Figure 8 shows a comparison between the open-loop operation (dashed-lines) and closed-loop operation (solid lines).

The controller has no direct means of measuring the *actual* shower water flowrate F_2^{WC3} that enters the plant; it merely assumes that the F_2^{WC3} input trajectory it prescribes is actually implemented as is in the plant. Figure 8f shows the shower water F_2^{WC3} trajectory that the controller "thinks" it has implemented on the plant. For illustration purposes, we also show the trajectory for "F_2^{WC3} (actual)" (Figure 8h), which depicts the actual shower water flowrate that enters the 3rd washer. The disturbance can be observed to occur between $t = 6.25$–8.25 h, where actual F_2^{WC3} (Figure 8h) is 30% higher than the presumed F_2^{WC3} (Figure 8f). As mentioned, the controller is not cognizant of the actual input trajectory.

We observe that the open-loop and closed-loop results are identical up to the point when the disturbance enters (at $t = 6.25$ h), except for a slight discrepancy between the two results in F_{out1}^{ST} (Figure 8g); this can be ascribed to the fact that a small positive value (1×10^{-4}) for the multitier relaxation factors ϵ_e and ϵ_s were used for numerical purposes. The unmeasured disturbance results in an influx of material into the storage tank, which causes its level H^{ST} (Figure 8d) to rise. Through the state measurements, the shutdown controller is able to detect this, and hence computes a change of course in the storage tank outlet flowrate, F_{out1}^{ST} (Figure 8g), where the excess material is discharged in order to compensate for the influx. Note also that the H^{ST} returns to nominal levels at the end of the restoration, as required.

By contrast, in the open-loop situation no feedback mechanism is present, therefore, the disturbance in the shower water flowrate is not detected. The F_{out1}^{ST} (Figure 8g) trajectory continues to operate as if no disturbance had occurred. As a result, the tank level H^{ST} (Figure 8d) does not return to its intended restoration target value at the end of the restoration, which is an undesirable result as it deprives the tank from anticipating further shutdowns. Other transient discrepancies between the open-loop and the closed-loop are seen in H_{st}^{WC3}, x_{DS3}^{WC3}, and x_{out1DS}^{ST} (Figures 8b, c, e).

The overall economics obtained in closed-loop $\left(\Phi_{econ}^{*} = \$66,915\right)$ is higher than that obtained by implementing the open-loop solution on the plant $\left(\Phi_{econ}^{*} = \$63,368\right)$. This demonstrates the advantage of using a framework for implementing dynamic optimization solutions in closed-loop through reoptimization.

Conclusion

The primary focus of this paper has been on the use of dynamic optimization for coordinating the buffer capacities in a plant in order to mitigate the effects of a unit shutdown. The model employed was a large-scale DAE system which was solved using a simultaneous solution strategy. Uncertainty in the downtime estimate parameter was handled through reoptimization, where updated downtime estimates are supplied directly to the controller. Case studies were carried out to examine the effects of reoptimization times with respect to erroneous downtime estimates (either longer or shorter than the actual). A multitiered economics MPC-based control framework for implementing the above optimal control policies was tendered for the detection and correction of disturbances and/or plant-model mismatch. The multitiered strategy is put forward as a means for systematically handling nonunique solutions in an optimization problem. Open-loop case studies demonstrate the application of multitiered dynamic optimization to the partial shutdown problem, while closed-loop case studies illustrate its application in a feedback loop.

Acknowledgments

Funding for this work through the McMaster Advanced Control Consortium is gratefully acknowledged.

Literature Cited

1. Pettersson B. Production control of a complex integrated pulp and paper mill. *TAPPI J*. 1969;52(11):2155–2159.
2. Pettersson B. Optimal production schemes coordinating subprocesses in a complex integrated pulp and paper mill. *Pulp Pap Canada*. 1970;71(5):59–63.
3. Pontryagin LS, Boltyanskii V, Gamkrelidge R, Mishchenko E. The Mathematical Theory of Optimal Processes. NY, Interscience, 1962.
4. Lee ES, Reklaitis GV. Intermediate storage and the operation of periodic processes under equipment failure. *Comput Chem Eng*. 1989; 13(11-12):1235–1243.
5. Lee ES, Reklaitis GV. Intermediate storage and operation of batch processes under batch failure. *Comput Chem Eng*. 1989;13(4-5):491–498.
6. Allison B. Averaging Level control for Surge Tanks in Series. In: Control Systems Conference. Technical Association of the Pulp and Paper Industry Press, Montreal, Canada, 2004:159–167.
7. Campo PJ, Morari M. Model predictive optimal averaging level control. *AIChE J*. 1989;35(4):579–591.
8. Huang YJ, Reklaitis GV, Venkatasubramanian V. Dynamic optimization based fault accommodation. *Comput Chem Eng*. 2000;24(2): 439–444.
9. Huang YJ, Venkatasubramanian V, Reklaitis GV. A model-based fault accommodation system. *Ind Eng Chem Res*. 2002;41(16):3806–3821.
10. Dubé JFH. Pulp Mill Scheduling: Optimal Use of Storage Volumes to Maximize Production. Master's thesis, McMaster University, Canada, 2000.
11. Balthazaar AKS. Dynamic Optimization of Multi-Unit systems under Failure Conditions. Master's thesis, McMaster University, Canada, 2005.

12. Cuthrell JE, Biegler LT. On the optimization of differential-algebraic process systems. *AIChE J*. 1987;33(8):1257–1270.
13. Chong Z, Swartz CLE. A modeling language for dynamic optimization. In: 56th CSChE Canadian Chemical Engineering Conference, Sherbrooke QC, 2006.
14. Fourer R, Gay DM, Kernighan BW. A modeling language for mathematical programming. *Manage Sci*. 1990;36(5):519–554.
15. Swartz CLE, Balthazaar AKS. Dynamic optimization of an integrated multi-unit system under failure conditions. AIChE Annual Meeting, Cincinnati, OH, 2005.
16. Swartz CLE. Algorithm for hierarchical supervisory control. *Comput Chem Eng*. 1995;19(11):1173–1180.
17. Qin SJ, Badgwell TA. A survey of industrial model predictive control technology. *Control Eng Practice*. 2003;11(7):733–764.
18. Chong Z. Dynamic optimization formulations for plant operation under partial shutdown conditions. PhD thesis, McMaster University, Canada, 2012.
19. Wisnewski PA, Doyle FJ III, Kayihan F. Fundamental continuous-pulp-digester model for simulation and control. *AIChE J*. 1997; 43(12):3175–3192.
20. Grace TM, Leopold B, Malcolm EW. Pulp and Paper Manufacture—Alkaline Pulping, Vol. 5, 3rd ed. Montreal, Quebec, Canada: The Joint Texbook Committee of the Paper Industry, 1989.
21. Wang R. Dynamic Simulation of Brownstock Washers and Bleach Plants. Master's thesis, University of British Columbia, Canada, 1993.
22. Noël A, Savoie M, Budman H, Lafontaine L. Advanced Brownstock Washer Control: Successful industrial implementation at James Mclaren. In: 2nd IEEE Conference on Control Applications, Vancouver B. C.,1993.
23. Gebart R, Westerlund L, Landalv I. Black liquor gasification–the fast lane to the biorefinery. In: Energy Conference (Risø National Laboratory), Risoe, Denmark, 2005.
24. Wächter A, Biegler LT. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Math Programming*. 2006;106(1):25–57.
25. ARKI Consulting and Development. Documentation for GAMS/CONOPT, available at http://www.gams.com/docs/solver/conopt.pdf. 2002. Last accessed 7 June 2013.

Manuscript received Nov. 15, 2012, and revision received Feb. 5, 2013.