Performance Monitoring and Advanced Process Control in the Pulp and Paper Industry

Leonardo C. Kammer, Bruce J. Allison, Alain A. Roche

Abstract: Modern, user-friendly tools are needed in process control, so that mill engineers are able to achieve high-quality product specifications via reduction of variability, as well as increased throughput and reduced energy consumption. This paper addresses the main features and implementation issues of two of those tools: a loop monitoring package and a multivariable advanced process controller. For each of these tools, we provide a mill-based case study with analyses of the benefits obtained. The use of the loop monitoring tool for predictive maintenance has eliminated most of the problems associated with control loop malfunctions and their lack of performance. The advanced process controller increased the throughput of a lime kiln while maintaining product quality, with an estimated annual savings of about US$120,000.

Application: The program maintenance tools described in this paper can help mills realize the benefits of well-functioning distributed control systems.

Computer control systems have been introduced to most pulp and paper mills to control various process variables. These systems have become so complex that monitoring and maintenance became a demanding task facing an increasingly limited staff. As a result, many controllers are not working optimally and some of them even increase process variability. According to a survey reported by Bialkowski [1], only 20% of control loops in the pulp and paper industry actually decreased process variability. The situation is likely to have improved since the survey was published. One of the main factors for our optimism is the development and mainstream use of modern user-friendly tools for process-control monitoring and design.

Control system auditing and maintenance involve three parts: performance monitoring, malfunction diagnosis, and problem correction. Performance monitoring is the most labor-intensive task among the three. Until recently, regular monitoring of control loop performance was rarely done in a mill. On the other hand, accurate diagnosis of problems requires comprehensive process knowledge; most mills rely on experienced engineers and operators for malfunction diagnosis. Once the cause of a problem is found, fixing the problem is relatively easy, at least manually.

A monitoring system is a tool that automatically monitors control loop performance by extracting performance-related information from process data. Most of the monitoring systems available on the market also assist mill personnel in malfunction diagnosis by pointing out the likely cause of poor performance. Given that these tools can only automate parts of the control-system maintenance task, the success of their application largely depends on efficient use by the mills.

Another factor that sometimes can contribute to process variability is the complex interactive nature of the process variables. In those cases, the best solution is to address the control design as a multivariable problem, instead of addressing independent single-input single-output (SISO) loops. The only advanced multivariable control methodology that has significantly affected industrial control engineering is model-based predictive control (MPC) [2]. MPC has been described as a *one-stop-shop* for the control of multivariable systems with constraints. Feedback, feed-forward, decoupling, and constraint handling are all integrated into one program. All the user needs to do is specify a model and the constraints, and choose tuning parameters with the help of simulation. This circumvents many of the design decisions and maintenance associated with applying classical SISO techniques to the control of complex multivariable processes.

Nearly all major suppliers of control equipment have their own products for performance monitoring and model-based predictive control. Even though these products are based on software implementation, their commercialization is frequently on a service-like basis, where a fee is charged on each individual application. To avoid excessive costs and be on the forefront of technological advancements, a pulp and paper company can alternatively establish a partnership with research institutes and/or universities. Given the research emphasis in

these institutions, a successful partnership requires a control engineer at the mill site committed to the installation and maintenance of the developed strategy or product.

This paper presents two of Paprican's control-related products: a loop monitoring and diagnosis system, and a model-based predictive controller. For each of these tools, we provide an overview of their characteristics and present a mill-based case study with analyses of the benefits obtained.

PERFORMANCE MONITORING

Control-loop performance monitoring became an active area for research and application in the field of process control after the publication of Harris's 1989 article on performance assessment [3]. Interpretation is simple and most of the process information required is readily available, making it very attractive for plant implementation. Paprican has adapted this technology for the pulp and paper industry [4, 5] and packaged the various functions for performance monitoring and diagnosis issues in a tool called LoopMD.

As an information tool, LoopMD extracts control loop performance-related information (performance indices) from large amounts of process data. Mathematically, it can be seen as a mapping from a large dimension data space to a significantly smaller dimension performance-indices space. This concept is illustrated in **Fig. 1**.

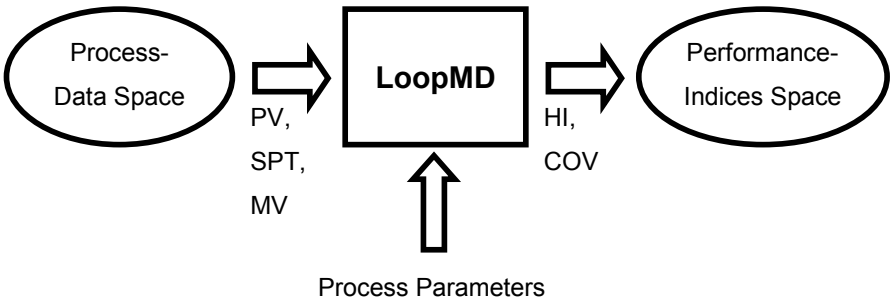

Figure 1. A loop monitoring system extracts relevant information from a large process data set.

The process data consist of sampled process variables, setpoints, and manipulated variables, whereas the performance indices consist of monitoring and diagnosis results. These diagnosis results are only used when monitoring results indicate poor control performance. The monitoring results contain two quantities for each loop: the performance index or Harris index (HI) and the coefficient of variation (COV). The Harris index of a control loop is defined as the ratio of current variance of the process variable to the minimum achievable variance. The larger the Harris index, the poorer the performance. The coefficient of variation is also used as an indicator for control loop performance. COV is defined as the ratio of the standard deviation to the mean of a process variable. The Harris index shows the potential improvement that can be achieved in terms of reducing process variability by fixing control problems. The coefficient of variation is an indication of the damage caused by poor control. These two quantities are periodically computed and monitored. When both COV and HI are high, the loop in question requires investigation. The evolution of these two performance indices over time gives a good indication of the urgency of the problem, whether, for example, a valve repair is needed immediately or whether it can wait until the next maintenance shutdown.

The diagnosis results contain the following indicators that help users identify the source of the problem: *subgroup index*, *modified performance index*, *trend*, *Harris index variation*, *control mode*, and *likely cause*. Control-loop problems are classified into the following categories: poor controller tuning, actuator/sensor malfunction, and interaction with other control loops. Interactions between control loops are automatically detected, and loops interacting with each other are assigned the same *subgroup index*. *Modified performance indices* are re-evaluations of the performance indices once it has been established that the loops were interacting with each other [5]. A large *modified index* implies that there is a problem in that loop, independent of the common problem among the interacting loops. If a loop has a small *modified index*, then its poor performance is likely caused by interactions with other loops.

Other than loop interactions, the system also tries to identify from process data whether the problem is caused by poor tuning or malfunctioning devices. The indicator *likely cause* contains the result of this analysis. In some instances, poor performance might be caused by the fact that the loop is under manual control. Control modes are detected from analysis of process data, and the result is provided in the indicator *control mode.* Another possibility is that the process is going through an upset at the time when data is collected. This upset may lead to a poor performance index and a large COV. Upsets are detected and the trend of the process is reported. *Harris index variation,* the difference between the current and previous Harris index estimates, also provides helpful information for diagnosis. A sudden increase in the Harris index might be indicative of a sudden deterioration of control performance, usually associated with device malfunctioning. It must be stressed that the diagnosis indicators do not provide a definitive conclusion regarding the cause of a problem, but they can provide useful assistance to the user.

Limited knowledge about the underlying processes is required to set up the system: the process deadtime and the sampling period of process data. LoopMD also comes with a set of accessory tools. With these tools, users are able to plot time series, autocorrelation, and power spectrum of a selected process variable. Users can also view the performance history trend of a loop and the process and information (P&I) diagram of a selected part of the process.

The software structure is guided by the point of view that a control loop monitoring system is a tool for extracting information from process data, for instrumentation, operation, and management staff in a mill. As illustrated in Fig. 1, the program needs access to process data, and it needs to present control performance related information in a format that is accessible to all relevant staff. There are essentially two approaches for data collection: indirectly through the mill process information system, or with a direct connection using OLE for Process Control. When using the process information system, sampling rates and attributes for data compression are automatically and temporarily adjusted during data collection. **Figure 2** illustrates a typical system implementation.

CASE STUDY FOR PERFORMANCE MONITORING

Paprican has been developing and transferring this technology to its member companies, over the last 8 years. Some of the mill implementations have been previously reported [5]. Three implementations can be considered as prototypes with a major involvement from Paprican staff. Five others were implemented in-house with minimal contribution from our staff. In two cases, mill personnel reprogrammed the user interface and in one case the system was completely redesigned [6]. In addition, one other company is using the software off-line to quantify the potential benefit from improved control maintenance.

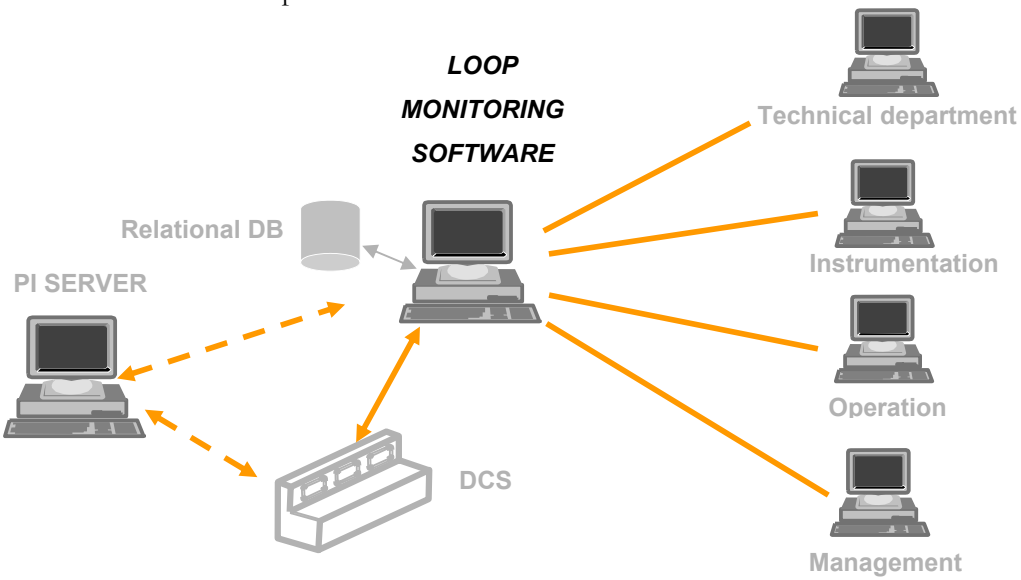

Figure 2. Typical system implementation.

This experience in a variety of mills has clearly shown that a monitoring and diagnosis system is only part of a maintenance program. Control maintenance is a multi-step process that involves identification of a potential problem, verification of the causes of the problem, correction of the problem, and to close the loop on this multi-step process, verification that the problem has been properly corrected. It is important to have a powerful and reliable tool to generate information on loop performance. However, this information will only have value if it is made available to the right people who will make corresponding decisions and take proper actions. Experience has shown that giving this information only to instrumentation personnel was not effective, and that operating staff were more concerned about getting the problem rapidly fixed. People are involved in the control maintenance process and it is essential to ensure that the human part of the loop is properly organized.

The procedure in place at a newsprint mill in Québec, Canada, is a good illustration of the key elements needed for successful application. We found that, with proper management, a network-oriented monitoring and diagnosis tool can significantly affect mill productivity.

LoopMD was initially implemented at the mill for their TMP process, boiler, and deinking plant, and then extended to the three paper machines, for a total of about 500 loops being monitored. It is installed on an NT server on the mill computer network. It retrieves data automatically from a distributed control system (DCS) once a day for the paper machines and every week for the rest of the mill, and then assesses loop performance. Results are recorded in the relational database of LoopMD on the server. Acceptable performance index and COV are defined for each loop and stored in the database, which also contains detailed information for each loop, such as tuning parameters, control span, or filtering of the process variable.

At present, the control engineer who carried out the implementation and developed a user interface is in charge of the system. After receiving an abnormal notification from the system, he will query the database through the interface. The primary information he uses is the performance index, the coefficient of variation, and the evolution of their last 10 values to confirm deterioration. Plots of time series, power spectra, and autocorrelations are available for more insight into the problem. Then, following the direction given by the diagnosis report, he will carry out bump tests for problematic loops and devices to verify the diagnosis. After correct diagnosis is made, he will follow a set procedure to ensure proper followup.

For each defective loop, a report is sent to the instrumentation superintendent and to the superintendent of the operating department concerned. The report identifies the defective loop and indicates the diagnostic for the problem, the performance index, and the coefficient of variation. After the instrumentation department fixes the problem, the report is sent back with a description of the corrective action. The control engineer audits the loops again to verify that the problems have been corrected and includes the results on the report that is redistributed. At the end of each month, he sends a report to the mill manager to summarize how many loops are monitored, how many of them have had problems, what are the problems, how many of the problems have been fixed, and for the unfixed ones, why they cannot be fixed. **Figure 3** summarizes the three main steps of the process. The software is used to identify the problem and to verify that it has been properly corrected, but the problem still has to be corrected.

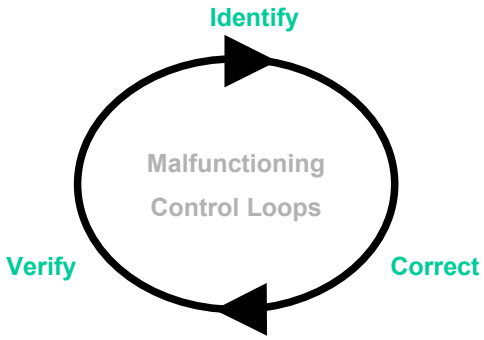

Figure 3 : Closing the loop on performance monitoring.

All related personnel have access to the database from their desktop computers through an Intranet system. With the help of the loop monitoring system, it takes the control engineer an average of 2 hours every

week to check the performance of about 500 loops. Loop monitoring on this scale was never done before the implementation of the system, because it would have been too time-consuming.

One year after implementation, tuning problems had practically disappeared. They show up only occasionally after a valve replacement. Slow deterioration of valves over many weeks is now the main source of poor loop performance: 70% due to backlash and 30% due to stiction. In addition, predictive maintenance associated with the continuous use of the system has reduced by about 70% the number of defective loops detected every month.

MODEL-BASED PREDICTIVE CONTROL

Model-based predictive control (MPC) is a term used to describe a very broad range of control methods that use a model of the process to calculate the control action, based on the minimization of a cost function [7]. The main difference between the MPC algorithm presented here and those that are commercially available is its ability to be implemented on distributed control systems (DCS) already available at the mills. Our particular MPC algorithm is based on generalized predictive control (GPC) [8]. GPC is general enough to handle processes with any number of controlled outputs, manipulated inputs, and measured disturbances. However, most of the main ideas in GPC can be conveyed by considering the SISO case.

The quadratic cost function to be minimized for SISO GPC is

$$J = E\left\{\alpha\sum_{i=n_1}^{n_2}\left[w(t+i)-y(t+i)\right]^2+\rho\sum_{i=1}^{n_u}\left[\Delta u(t+i-1)\right]^2\right\}, \tag{1}$$

where w is the setpoint, y is the controlled variable (output), u is the manipulated variable (input), $\Delta = 1-z^{-1}$ is the differencing operator [z^{-1} is the backward shift operator, i.e. $z^{-1}y(t)=y(t-1)$], and E is the expectation operator. In effect, Eq. (1) minimizes the predicted errors (setpoint minus output) over a future number of sampling instants (called the output horizon) defined by n_1 and n_2. The independent variables in the minimization are n_u future input moves, called the input horizon. The weighting factors α and ρ are used to adjust the relative importance of predicted errors and control moves, respectively. Only the first (current) input move is implemented and the optimization is resolved with new information at the next sampling instant.

The output predictions in Eq. (1) are made using a process model that is represented by the following controlled auto-regressive integrated moving-average (CARIMA) model:

$$A(z^{-1})y(t) = B(z^{-1})u(t-k)+T(z^{-1})v(t)/\Delta, \tag{2}$$

where v is an independent normally distributed (white) noise with variance σ_v^2; A, B, and T are polynomials in the backward shift operator:

$$\begin{aligned} A(z^{-1}) &= 1+a_1z^{-1}+\ldots+a_{na}z^{-na}\\ B(z^{-1}) &= b_0+b_1z^{-1}+\ldots+b_{nb}z^{-nb}\\ T(z^{-1}) &= 1+t_1z^{-1}+\ldots+t_{nt}z^{-nt} \end{aligned} \tag{3}$$

and k is the deadtime.

When there are no constraints, Eq. (1) may be minimized analytically. However, it is also possible to make the minimization subject to absolute constraints on the output,

$$y^{\min} \le y(t+i) \le y^{\max},$$

where $i=n_1$ to n_2, and absolute and rate of change constraints on the input,

$$\begin{aligned} u^{\min} &\le u(t+i-1) \le u^{\max}\\ \Delta u^{\min} &\le \Delta u(t+i-1) \le \Delta u^{\max}, \end{aligned}$$

where i=1 to n_u. Because the cost function is quadratic, quadratic programming (QP) techniques are well suited to minimizing Eq. (1) in the general case of constraints on both outputs and inputs. To avoid the computational burden of QP, output constraints are not considered in our implementation and input constraints are enforced by using the one-step optimal saturation correction algorithm proposed by Segall and colleagues [9]. Note that the one-step algorithm is exactly equivalent to a QP when n_u=1.

The tuning parameters in SISO GPC are n_1, n_2, n_u, a, and ρ. In addition, constraints that consist of minimum, maximum, and rate-of-change limits need to be chosen for the input. Consider the situation when the process may be represented by a first-order plus deadtime model, which is very often the case in the process industries. Now consider the choice of tuning parameters $n_1=n_2=n_y$, n_u=1, a=1, ρ=0, and $T=A$. When a step change in setpoint or an output step disturbance occurs, this controller will calculate a single control action (n_u=1) that will drive the output back to the setpoint at a point n_y sampling intervals into the future. Furthermore, the return to setpoint will occur in the minimum possible time when $n_y=k$ (recall that k is the deadtime). This is called *output deadbeat* control, and may lack robustness or require excessive control actions. Robustness is improved and control action magnitude reduced by setting

$$T(z^{-1}) = A(z^{-1})T^*(z^{-1}),$$

where $T^*(z^{-1}) = (1-\exp(-T_s/\lambda^d)z^{-1})$. Here T_s is the sampling interval and λ^d is the closed-loop time constant for an output step disturbance. Because the choice of T does not affect the setpoint response in GPC, it is necessary to filter the setpoint,

$$w^f(t) = \left[\exp(-T_s/\lambda^w)\right]w^f(t-1) + \left[1-\exp(-T_s/\lambda^w)\right]w(t),$$

and use $w^f(t)$ in place of $w(t)$ in the controller. The designer is free to choose λ^w different from λ^d, yielding a two-degree-of-freedom design. When $\lambda^w = \lambda^d$, this GPC is exactly equivalent to a Dahlin controller, which is a deadtime compensator widely used in the pulp and paper industry, particularly in paper machine control systems.

Of course, GPC is more general than a Dahlin controller. Different choices for the tuning parameters may be necessary when the process model is more complex than first-order plus deadtime. T may need to be modified when the goal is to reject input disturbances rather than disturbances at the output. Finally, GPC is extensible to multivariable control, and can handle non-square systems.

CASE STUDY FOR MODEL-BASED PREDICTIVE CONTROL

Paprican has been developing its MPC algorithm over the last 8 years, with many successful applications on digesters, high-consistency storage tanks, and lime kilns. This section presents an application of MPC to a lime kiln in Prince George, Canada [10]. The lime kiln is an important unit operation in the chemical recovery and causticizing system of a kraft mill. Ideally, the kiln should be able to produce sufficient, high-quality lime to meet the needs of the causticizing operation while minimizing operating costs and environmental emissions. A well-designed control system can help meet these objectives. Potential cost reductions from improved control result mainly from: (a) reduced requirement for purchased lime, (b) reduced refractory damage, and (c) lower fuel consumption. In a recent survey [11], savings of between US$ 300,000 and US$ 1.4 million per year were attributed to kiln control. According to Järvensuvu and colleagues [12], proper control of the kiln can also help reduce total reduced sulfur (TRS) emissions during transient operation. Recent trends in the cost of fuel are driving a renewed interest in kiln control.

From a control perspective, the lime kiln may be viewed as a multivariable distributed parameter process. The output variables to be controlled are the solids temperature at the product (hot) end, and the temperature and excess oxygen content of the combustion gas at the feed (cold) end. The manipulated input variables are the fuel flow rate and the induction draught (ID) fan speed or damper position. The lime mud feed rate is a measured disturbance. Thus, the kiln is a non-square system with two inputs, three outputs, and one measured disturbance (**Fig. 4**). For any given feed characteristics or production rate, there is an optimal set of operating conditions that produces the best quality lime while minimizing energy consumption. The purpose of the kiln

control system is to ensure consistent, high-quality lime production by maintaining these operating conditions, despite disturbances. This is made difficult by significant interactions between variables, long time delays, nonlinearities, operational constraints, and noisy measurements. However, as indicated above, the expected benefits through cost reductions are significant.

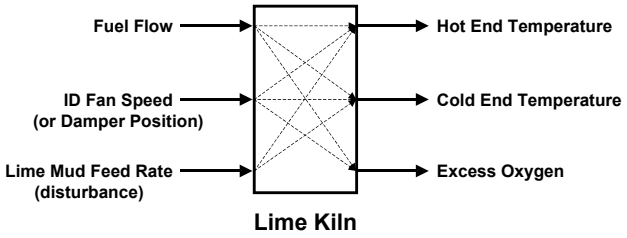

Figure 4: Block diagram of a lime kiln showing inputs, disturbances and outputs. Dashed lines indicate interactions between variables.

The most straightforward application of MPC to the lime-kiln problem would be to control the hot-end and cold-end temperatures to a target, while maintaining the excess oxygen within constraints. The disadvantage of this MPC approach is that it requires the ability to handle output constraints. This is usually achieved by solving a quadratic program (QP) on-line, which comes at a large computational cost. Thus, MPC often requires dedicated hardware for implementation.

To overcome this problem, we use a cascaded control structure that exploits the inherent advantages of MPC (deadtime compensation, decoupling, input constraint handling) while avoiding the computational complexities of a QP. The result is a constrained, model-based, multivariable lime kiln control strategy that can be implemented on standard DCS hardware available in most mills. **Figure 5** shows the control strategy. This is exactly equivalent to the strategy described by McIlwain [13], but with the individual PID controllers, decouplers, etc., replaced by two cascaded MPCs. The inner loop is a 2×2 MPC that controls the cold-end temperature and excess oxygen with the fuel flow rate and the ID fan speed. The outer loop is a 2×1 MPC that manipulates the cold-end temperature and excess oxygen set points to control the hot-end temperature while optimizing (mid-ranging) the excess oxygen over the long term. In effect, both the cold-end temperature and excess oxygen setpoints are immediately manipulated when there is a setpoint change or disturbance in the hot-end temperature. Eventually, the oxygen setpoint is brought back to its target value (O_2^{opt}) when the cold-end temperature setpoint changes begin to take effect. Note that the lime mud feed rate measurement is not used as a feedforward signal because the effect of feed rate changes is seen almost immediately in the cold-end temperature.

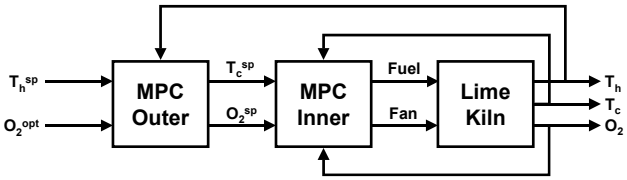

Figure 5: Block diagram of cascaded MPC control structure.

The inputs of the inner loop are the fuel flow rate in m^3/h denoted by F, and the ID fan speed in *rpm* denoted by S. The outputs are the cold-end temperature in °C denoted by T_c, and the excess oxygen in percent, denoted by O_2. A preliminary inner-loop model was fit from bump test data using standard open-loop identification methods. The basic sampling interval was 10 s, but the data was filtered and re-sampled at a 1-min interval before identification. The preliminary model was then refined through several closed-loop experiments. After converting to continuous time, the final model in deviation variables was

$$\begin{bmatrix} T_c(s) \\ O_2(s) \end{bmatrix} = \begin{bmatrix} \frac{0.16}{140s+1} & \frac{0.21}{10s+1} \\ \frac{-0.013}{6.1s+1} & \frac{0.026}{3.5s+1} \end{bmatrix} \begin{bmatrix} F(s) \\ S(s) \end{bmatrix}. \quad (4)$$

Note that although the individual transfer functions are simple first-order systems, there is a wide range of dynamic behavior, ranging from time constants of several minutes to several hours. The Eq. (4) model was used to implement the inner-loop controller.

Tuning parameters for the inner loop were chosen according to the method discussed in the previous section, and are shown in **Table I**. Thus, the quadratic cost function for the inner loop is

$$J_i = E\left\{\left[T_c^{sp}(t) - T_c(t+50)\right]^2 + 1000\left[O_2^{sp}(t) - O_2(t+5)\right]^2\right\}. \quad (5)$$

TABLE I
TUNING PARAMETERS FOR INNER-LOOP GPC.

Output	n_y	α	λ^d *(min)*	λ^v *(min)*	*Input*	ρ	$\nabla u^{min}, \nabla u^{max}$
Tc	*50*	*1*	*15*	*15*	*F*	*0*	*-15,15*
O_2	*5*	*1000*	*10*	*10*	*S*	*0*	*-7.5,7.5*

Here we chose n_u=1 for each input and n_1=n_2=n_y for the outputs. When ρ=0, the control move penalty terms drop completely, and α has no effect except at saturation. The α values shown above put more emphasis on control of excess oxygen than cold-end temperature in the event that either the fuel flow or fan speed saturate. Tuning of the unconstrained closed-loop response was achieved by adjusting n_y and λ. The absolute input constraints for the inner loop were determined by the operating personnel, and were fixed at u^{min}=0 and u^{max}=2100 $m^3/hour$ for the fuel flow and u^{min}=0 and u^{max}=1100 *rpm* for the fan speed. In addition, rate-of-change constraints (see Table I) were used to limit the size of any single control move. This is particularly useful during upset conditions.

As for the outer loop, the inputs are the cold-end temperature setpoint denoted by T_c^{sp} and the excess oxygen setpoint denoted by O_2^{sp}. The output is the hot-end temperature in °C, denoted by T_h. A similar procedure of identifying a preliminary model based on open-loop bump tests (i.e. inner loop closed, outer loop open) followed by closed-loop experiments to improve the model was followed. This resulted in a model with 35 min of deadtime in the response to O_2^{sp} changes, which is a problem for implementation because a large number of past control actions would be required for deadtime compensation. One way to deal with this problem was to increase the control interval for the outer loop. However, that was considered undesirable because of the potential negative effect on hot-end temperature control. Thus, a Padé approximation of the deadtime was used to arrive at the following reduced-order model:

$$T_h(s) = \begin{bmatrix} \frac{3.0(-90s+1)}{(60s+1)^2} & \frac{-50(-17.5s+1)}{(40s+1)(17.5s+1)} \end{bmatrix} \begin{bmatrix} T_c^{sp}(s) \\ O_2^{sp}(s) \end{bmatrix}. \quad (6)$$

This model exhibits a response to a step in cold-end temperature setpoint that initially goes the wrong way. The explanation for this is that the fuel flow and fan speed need to increase together for the cold-end temperature to increase and excess oxygen to remain constant. Initially, the increase in fan speed cools off the hot end. Later, the increase in fuel causes the hot end to heat up. Thus, the wrong-way response is simply due to the differing dynamic effects of the fuel flow and fan speed on the hot-end temperature. The model (6) was used to implement the outer-loop controller.

The tuning parameters for the outer loop are shown in **Table II**, and the quadratic cost function for the outer loop is

$$J_0 = E\left\{\begin{matrix} \left[T_h^{sp}(t) - T_h(t+180)\right]^2 + 400\left[\Delta T_c^{sp}(t)\right]^2 \\ +1.75\left[O_2^{opt}(t) - O_2^{sp}(t)\right]^2 \end{matrix}\right\}. \quad (7)$$

Again, we chose n_u =1 for each input and n_1=n_2=n_y for the output. Because there is only one output, α defaults to one. Furthermore, there are more inputs than outputs, so one of the inputs needs to be constrained to a steady-state value by modifying the cost function. The excess oxygen (last term in Eq. 7) was chosen for reasons mentioned earlier, and is constrained to never deviate too far from the optimization setpoint O_2^{opt} during normal operation. However, the weighting factor for excess oxygen was chosen so that the controller would contain integral action in the event of cold-end temperature saturation. That is, an offset will exist between O_2^{sp} and O_2^{opt} only if required to maintain the hot-end temperature at setpoint when the cold-end temperature setpoint saturates. The weighting factor for cold-end temperature determines how fast the oxygen is mid-ranged. The absolute input constraints for the outer loop were fixed at u^{min}=1 and u^{max}=5% for excess oxygen and u^{max}=350°C for cold-end temperature. The lower limit for cold-end temperature was changed depending on the production rate. Operating experience was that a certain minimum production rate dependent temperature was required to achieve proper drying of the lime mud before reaching the chain section of the kiln. As in the case of the inner loop, rate-of-change constraints (see Table II) were used to limit the size of any single control move. Recall that the manipulated variables for the outer loop are the cold-end temperature and excess oxygen setpoints, so the above limits apply to the setpoints only, and should not be interpreted as output constraints.

TABLE II
TUNING PARAMETERS FOR OUTER-LOOP GPC.

Output	n_y	α	λ^d *(min)*	λ^w *(min)*	*Input*	ρ	∇u^{min}, ∇u^{max}
T_b	*180*	*1*	*3*	*3*	T_c^{sp}	*400*	*-5,5*
				60	O_2^{sp}	*1.75*	*-0.1,0.1*

Figure 6 shows 60 hours, or 2.5 days, of operation with the cascaded GPC strategy and the kiln undergoing a number of large production rate changes. Apart from a disturbance coinciding with the large sudden decrease in production at the 10-hour point, which was later found to be caused by a mechanical problem with the venturi scrubber throat damper, the hot-end temperature is extremely well controlled, despite frequent rate changes. The hot-end temperature 2-sigma is 32°C over the entire 60-hour period and only 21°C over the second half when the production rate is not so variable. Note that the cold-end temperature is nearly always riding the lower constraint (which is determined from the production rate). This is because the excess oxygen is above the optimization target value of 1%. To decrease the excess oxygen, the cold-end temperature must decrease, but then the mud might not dry sufficiently and might plug the chain section. If there were no lower constraint on cold-end temperature, the excess oxygen would eventually reach the optimization target, and thus oxygen would become the constrained variable. This is an example of how a process must sometimes be run against constraints to operate economically. Of course, part of the success of the outer loop in providing tight control of the hot-end temperature and minimizing energy consumption is due to the inner loop, which provides tight control of the cold-end temperature and excess oxygen via coordinated manipulation of the fuel flow and fan speed.

Table III compares several different control strategies that have been implemented on this kiln since the early 1980s. In each case, the values in the table are averages taken over a 48-hour period representing typical operating conditions for the kiln at *peak* production and constant $CaCO_3$ residual. The kiln ran under manual control until 1993 when the strategy discussed by McIlwain [13] (which we subsequently refer to simply as PID) was implemented. This change brought significant benefits, including a dramatic increase in lime production and decrease in operating temperatures and energy consumption. These improvements were used to reduce the need for chemical makeup and to increase pulp production. However, to allow comparisons, we have expressed the benefits in terms of potential fuel cost savings at constant production, which in the case of PID amounted to more than US$ 600,000 per year (see Table III). The cascaded GPC strategy was commissioned in March 2000, and is estimated to have achieved a slight (3 metric tons/day, or 1.5%) increase in lime production and further decreases in operating temperatures and energy consumption, resulting in an additional US$ 140,000 per year in potential fuel cost savings. Note that the lime production increase resulted from a substantial reduction in dusting, from 19% of feed (or 91 metric tons/day) to 14% of feed (or 64 metric tons/day), and was achieved despite reducing the mud feed rate.

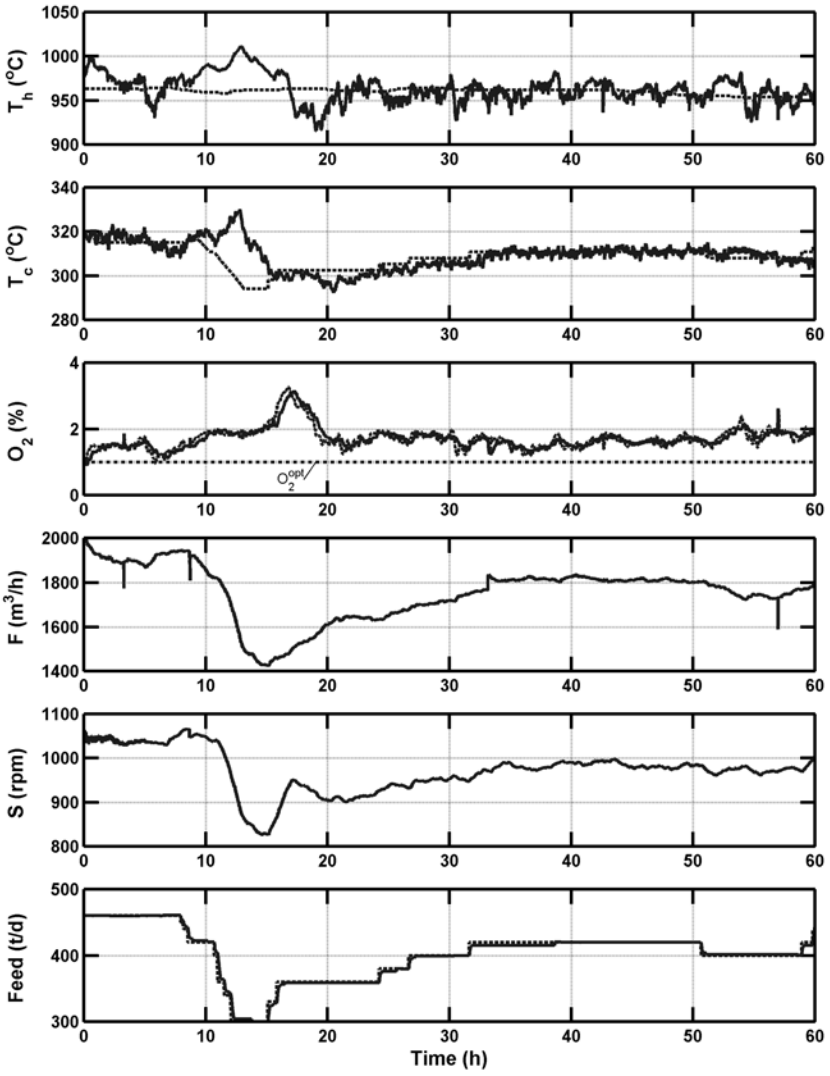

Figure 6: Cascaded GPC control of the lime kiln under variable operating conditions. Solid line is process variable, dashed line is set point, dash-dotted line in the third plot is O_2^{opt} . The temperature T_c and oxygen O_2 are operating near or at their lower constraints, minimizing energy consumption at high throughput.

TABLE III
TYPICAL PEAK OPERATING CONDITIONS.

Variable	*Manual*	*PID*	*GPC*
Feed of CaCO3 (a.d. metric tons/day)	*410*	*480*	*460*
Recycle (% of feed)	*37*	*19*	*14*
90% CaO Produced (metric tons/day)	*130*	*196*	*199*
Hot End Temperature (°C)	*1037*	*980*	*965*
Cold End Temperature (°C)	*357*	*335*	*305*
Fuel Consumption (m^3/hour)	*1925*	*2000*	*1850*
Energy (GJ/t 90% CaO)	*13.8*	*9.5*	*8.6*
Relative Savings (US$ x1000/year)	*0*	*630*	*770*

Table IV contains an evaluation of the benefits of cascaded GPC over the first year of operation, where PID control was used as the benchmark. Over the year with GPC, the kiln was able to support 2.1% more digester production while eliminating the equivalent of roughly 11 days of digester production that were formerly supported with purchased quicklime and caustic. Overall, this represents a calculated lime production increase of 5.1%. Total fuel consumption was roughly the same for the two periods. Thus, as a result of improved control, the kiln was able to carry 5.1% more of the causticizing load, without increasing fuel consumption. Had kiln production not been increased, fuel consumption would have been reduced by roughly 5%. The net annual saving in chemical makeup was estimated at roughly US$ 123,000. Other benefits of reducing makeup chemicals are that quicklime addition causes deterioration in causticizing plant operation and caustic makeup makes it more difficult to maintain white liquor % sulfidity. We did not attempt to attach a dollar value to these benefits. It should be noted that this kiln is producing at more than 1.5 times its original design production rate. Although the PID controls installed in 1993 were quite successful and represented a major debottlenecking of the kiln, GPC still allowed us to improve this operation by an additional 5%. As all the changes associated with going from PID to GPC were software related, and implemented on hardware that was already available in the mill, no capital expenditures were required.

TABLE IV
ONE YEAR ECONOMIC EVALUATION OF GPC CONTROL.

Control Strategy	Fuel Consumption (m^3/a.d. metric tons pulp)	Quicklime Addition (kg/a.d. Metric tons pulp)	Caustic Addition (kg/a.d. Metric tons pulp)	Net Savings (US$ x1000/year)	90% CaO Production Increase (%)
PID	51.9	7.57	2.41	Base Case	Base Case
GPC	52.2	3.24	1.00	123	5.1

Another benefit of GPC is that it has proven to be easier to apply and maintain. Because the PID control strategy was composed of a number of separate elements (feedback controllers, decouplers, anti-windup schemes, etc.), it had become somewhat cumbersome to maintain.

CONCLUSION

Despite all the benefits that can be provided by performance monitoring and model-based predictive control, their application in the pulp and paper industry is still limited. We foresee a time when MPC and performance monitoring tools will be widely used in the industry. Then, these tools will be an integral part of any distributed control system, similar to the way common PID controllers are currently available. Although successful applications of these technologies still depend on a certain degree of human expertise, much effort is being put into the development of user-friendly interfaces and simplified mechanisms that remove a significant part of the complexity of modern process-control tools. **TJ**

LITERATURE CITED

1. Bialkowski, W., *Dream versus reality: A view from both sides of the gap*. *Pulp Paper Can.*, 94(11): T19(1993).
2. Maciejowski, J.M., *Predictive Control with Constraints*, Prentice-Hall, Harlow, U.K., 2001.
3. Harris, T.J., *Can. J. Chem. Eng.* 67: 856(1989).
4. M. Perrier, M., and Roche, A.A., Towards a mill-wide evaluation of control loop performance, In *Control Systems '92*, p. 205–211, PAPTAC, Montreal, QC, Canada.
5. Owen, J.G., Read, D., Blekkenhorst, H., Roche, A.A., A mill prototype for automatic monitoring of control loop performance, In *Control Systems '96*, p. 171–178, PAPTAC, Montreal, QC, Canada.
6. Ogawa, S., A data analysis and graphical presentation system for control loop performance assessment, In *Proceedings of 1998 TAPPI PCE&I Conference*, TAPPI PRESS, Atlanta, Georgia, USA.
7. Camacho, E.F., and Bordons, C., *Model Predictive Control*, Springer-Verlag, London, U.K., 1998.
8. Clarke, D.W., Mohtadi, C., Tuffs, P.S., *Automatica*, 23(2): 137(1987).
9. Segall, N.L., MacGregor, J.F., Wright, J.D., *Automatica*, 27(1): 135(1991).
10. Allison, B.J., and Ball, J., Cascaded model predictive control of a rotary lime kiln. *Pulp Paper Can.* 105(1): T20(2004).
11. Roche, A.A., Jack, J.S., Savoie, M., *Pulp Paper Can.* 98(12): T440(1997).
12. Järvensuvu, M., Kivivasara, J., Saari, K., *Pulp Paper Can.*, 100(11): T347(1999).
13. McIlwain, J.A., Kiln control, *Pulp Paper Can.* 93(11): T307(1992).

Kammer and Allison are with Paprican in Vancouver, and Roche is with Paprican in Pointe-Claire, Canada. Email Kammer at lckammer@netscape.net.

This paper was originally published in O Papel 65(8): 55(2004), the official publication of the Brazilian Technical Association of Pulp and Paper (ABTCP). It is being reprinted by *TAPPI JOURNAL* through a cooperative agreement between TAPPI and ABTCP.

This paper is also published on TAPPI's web site <www.tappi.org> and summarized in the January Solutions! for People, Processes and Paper magazine (Vol. 88 No. 1).