Linking Control Strategy Design and Model Predictive Control

James J. Downs
Eastman Chemical Company
Kingsport, Tennessee

Abstract

The purpose of this paper is to describe the importance of the underlying relationship between control strategy design and model predictive control. Successes and problems encountered when implementing model predictive control (MPC) on chemical processes have revealed that understanding this relationship provides insight into the nature of the process control problem. Model predictive control (MPC) has been used as an effective tool to gain the process control benefits that come from its ability to handle constraints, process interactions, and multiple time frames. The use of the MPC algorithm on a variety of chemical processes has led to insight on how to effectively use MPC along with traditional control strategy notions to improve process control. The development of control strategies using MPC has resulted in the typically reported benefits of increased throughput and reduced process variability. Several issues remain to be addressed. These include controller tuning, complex performance criteria, depth of integration of MPC with the regulatory control layer, redundant process information, and controller robustness to measurement loss or deterioration.

Keywords

Model predictive control, Control strategy design, Process control vendor, Process control education, Cascade control

Introduction

Process control strategy design has been the cornerstone of successful application of process control technology for many years. As new process control methods and algorithms have been conceived and developed, their successful application in the process industries has relied upon the underlying insight into the nature of the process. Certainly, if one chooses the right problem the success of a particular technology is enhanced. One technology can be shown to be superior to another simply by judicious choice of application. The success of model predictive control lies in its ability to cast process control strategy choices into a manageable framework. The capability to dynamically decouple process control loops, to handle process constraints, and to minimize deviation from set point are important but are more tactical in nature compared to the control strategy changes that take place.

As a model predictive controller is exercised throughout its allowable range, any number of control strategies may be manifested. The ability to understand the ramifications and consequences of each strategy or group of strategies is key to the successful implementation of model predictive control. In the past, we had a fixed control strategy. The process control designer was charged with the design of a strategy that would perform the best for as wide a range of circumstances as possible. Often if it were known beforehand that a strategy could not handle a particular set of conditions, those conditions were avoided during operation or protected against on a case-by-case basis. Given that the control strategy was fixed, it was studied by subjecting it to the variety of disturbances and operating scenarios that were plausible. As unexpected operating conditions and disturbances were encountered, the plant operators acted as “test pilots” having to manage the new operating regime as best they could.

Model predictive controllers provide the capability to change the fundamental control strategy while the process is operating—in its simplest sense acting as a control system override. The opportunity to exercise and explore complex MPC control designs is limited by the factorial number of possible strategies that can be in effect at a given time. How each of these strategies will respond to an array of disturbances and operating conditions must be answered or addressed to avoid having “test pilots” testing systems that look like black boxes.

The evolution of process control technology has expanded the role for the process control engineer. The notion of designing multiple, complex control strategies that can change during the normal course of operation is becoming more prevalent. Certainly high and low select overrides have been around for many years. However, the extent to which even mildly complex MPC applications result in unexpected control strategies is a new realm. The purpose of this paper is to discuss how this new focus is unfolding.

This paper starts with a perspective on the current process control work environment in the chemical industry. This perspective highlights what a process control engineer is likely to face in today’s world. With this perspective as a backdrop, a linkage is developed between the familiar territory of control strategy design and the newer, possibly unfamiliar tool, model predictive control. This linkage demonstrates the need and value of accumulated process knowledge and traditional process control notions when faced with reaping the widely acknowledged benefits of model predictive control. The paper concludes with examples that highlight the variety of problems benefiting from our application of model predictive control and that illustrate some of the implementation issues that we have encountered.

The Process Control Landscape in the Chemical Industry

New Plants

Normally for new plants the process control design is determined as a fixed strategy that provides regulatory control for the array of expected disturbances and operation regimes. The focus during this activity is maintaining the plant operation at a nominal operating condition from which operations can move to achieve product properties. New plants often contain new technology that involves uncertainty of operation and of performance. All that is needed of the control system is to maintain stability and to be understandable by personnel with a wide range of experience and education levels. Forays into the use of advanced control techniques on unfamiliar unit operations or processes employing new process technology have usually demonstrated that starting up with a simple, understandable control system is best. Once the operating characteristics are more known then the operation can be optimized employing more advanced control techniques.

The design of control strategies for new facilities warrants the need for control strategy design and analysis. The formation of rugged, well thought out regulatory control strategy designs that can withstand the variety of disturbances and operating abnormalities encountered during the first year of plant operation is a requirement for future process control enhancements. Undoubtedly, operating a process closer to optimum conditions and determining where that is requires some semblance of stable operation. There has been much written and presented to help integrate the process and the control strategy design. Recent examples include Barolo and Papini (2000), Groenendijk et al. (2000), and Tseng et al. (1999). These efforts are directed at the development of process designs and regulatory control strategies that achieve good regulatory control. This work is to be recognized for providing guidance where a few years ago there was none. In a similar vein, Luyben (1998c; 1998a; 1998b; 2000) has published a series of articles to guide control strategy development based on process situations. While the incorporation of advanced control technology into the process design phase is a notable goal, current practice is to get new processes up and running first and follow later with control system enhancements.

Existing Plants

Once a plant has been in operation long enough to find and fix problems that preclude stable operation then the initiation of process improvement activities is a natural consequence. It is during this time that enough is learned about the operation that advanced control techniques can be successfully applied. During this time the linkage among what is needed, what is feasible, and the appropriate technology to apply is most important. Among the many improvement opportunities and the flood of available technologies, it is necessary for the process control engineer to discriminate between process equipment problems and control strategy problems. Often process improvements that come from process control changes are of the control strategy variety rather than control algorithm changes. From our viewpoint, MPC is regarded as a "control strategy change agent" instead of an algorithm for improved high performance control. Indeed, it is capable of both. The effort needed to develop and maintain models accurate enough for high performance control, however, often outweighs the marginal benefits. The need to change control strategies for differing modes of operation has been more persistent.

The process knowledge available for existing plants provides insight into the true objective that needs to be achieved by the control system. While cursory overviews of plant operation may yield process control objectives that appear reasonable, often a deeper process insight is needed to arrive at the desired process objectives. This deeper process insight comes from understanding process chemistry, unit operation objectives, business objectives, and process flow structures.

Control Objectives

The definition of process control objectives often involves an evolutionary path. Often an initial statement of what the control system should do is oriented around what the current control system cannot do. "If only we could control the temperature, we would be happy with our operation" leads to "The temperature control we have is great, but what we really need is to control the composition of ...". This in turn may lead to other objectives that may change once the successive performance plateaus are reached. As process control systems hold key variables within narrower and narrower limits, the costs in terms of increased variation in other non-key variables becomes apparent and control objectives change. Furthermore, tighter control allows process engineers to see process improvement opportunities that are otherwise hidden.

When product requirements change, control objectives may need to be altered. These changes may involve a simple change such as altering controller weighting parameters or require an entire control system structure change. Labor, retraining, and opportunity costs to maintain and improve advanced control systems when process objectives change are compared with the economic benefits. This situation results in large, single product, unchanging plants to be obvious candidates for advanced control applications. This type of application is common and has been reported often in the literature. On the other hand, for plants that are smaller, multiple product, or undergo occasional change there is a need to be able to reap the benefits available from advanced control without prohibitive costs. An MPC structure that can represent various control strategies can be a very effective tool.

In this more fluid application environment, we are not driving for a lower IAE, ITAE, etc. as much as we are addressing opportunities to conveniently drive processes to the optimum steady state when constraints are encountered. In many of our applications it is the steady state targeting feature of model predictive control that is the important piece. Control objectives that we encounter are much more focused on where the process will line out under various conditions rather than on how a process will dynamically respond.

The Academic/Vendor/Industry Relationship

The implementation of process control technology, and in particular model predictive control, requires process control skills that may not be taught in the normal undergraduate curriculum. Certainly skills arising from formal training in this area may not be recent or deep enough to warrant a personal embarkation into an advanced control project. The role of the corporate process control group in the chemical industry is pivotal in channeling the technology to the appropriate applications and ensuring their success. Without such a group, the linkage is weak between those with the problems and those possessing the solutions. A central group can provide the standardization and stewardship needed for company wide application.

The process control vendor has historically provided the control toolkit needed to apply process control technology along with training and personnel to use their products. The current climate of specialization of service providers as modeled in the communications industries is becoming popular in the chemical industry as well. Vendors are moving from providing a product to providing a service. Academic institutions on the other hand provide trained personnel and technology ideas but no industrially hardened products. The process control toolkits on the market today have a variety of technologies that hopefully weave together to make their use easy. The relationship between vendors and academicians is becoming stronger. This is driven by a viewpoint that few companies have the wherewithal to incorporate new theoretical advancements into their day-to-day business. Process control vendors are becoming a more important avenue through which theoretical advancements make their way to industrial practice.

Companies are in transition to meet relentless market pressures on shorter and shorter time horizons and the lure of marketing suggestions that promise short payback times while requiring little long term corporate investment are strong. The choice of appropriate control technology requires an unbiased viewpoint. Often times the solution of a control problem can be accomplished via many technology avenues. If a vendor is selling hammers then the vendor sees most problems looking like nails. The implementation of advanced control is only warranted where simpler control techniques are inadequate. Corporate process control groups should have the knowledge to make choices among competing technologies based on life cycle costs and other intangible factors. Vendors while economically driven should nevertheless provide a similar unbiased approach to problem solution.

The identification of appropriate candidates for advanced control usually requires proprietary knowledge of process economics, process weaknesses, process chemistry, and even corporate politics. Relying upon operation personnel to identify candidates in the midst of regulatory, labor, and production demands is difficult for reasonably steady operations and nearly impossible for constantly changing production environments. Indeed, a process control specialist with knowledge of corporate objectives and a process viewpoint of the larger picture has a much better chance of identifying the best projects. Once projects are selected, the process knowledge needed to reach solutions is normally located within operations. Often that knowledge is shared with the in-house process control specialist because of past experiences, built up trust relationships, or personal relationships. The criticality of process knowledge cannot be overstated (Downs and Doss, 1991). How much process knowledge can be shared with non-company personnel, secrecy agreements notwithstanding, is always a subject of debate. Process discoveries during implementation, accumulated process operational savvy, and application tricks are all subject to loss after the project is complete. In addition, process control revelations arising from implementation become leveragable knowledge for the control vendor. Undeniably, early customers are in the role of guinea pig until adequate enhancements harden advanced control products.

One of the most important factors in the success of process control projects is the long-term maintenance of the finished product. Valves change, transmitter ranges change, processes change. There is an inevitable march toward a process that sooner or later does not match the process control system. For large volume plants with only a few products the process control system may remain valid for several years. However, as the variety of projects increases, the applications require more support. Local personnel can change simple items, however, software upgrades, process changes, and even retuning will probably require specialized support. This support can be provided by service contracts or by in-house specialist.

In this environment, Eastman has thus far benefited from having a corporate group to manage this activity. The Eastman process control group has maintained the strength to objectively evaluate the cost/benefit tradeoff in the spectrum ranging from an entirely in-house process control program to one that is entirely contracted to a service provider. Our current approach is to purchase those products and technology that provide value and are generally one-time in cost. Those products that entail on-going costs for each use or application have been used sparingly due to their continual drain on process control profitability. Each company trades off between the expense of maintaining in-house talent and purchasing that talent through vendors. However, recent informal survey data indicates that the need for in-house process control expertise remains strong (Downs, 2000).

We believe that there is a strong case for academia to continue to provide people knowledgeable in process control not only for service vendors but also for corporate in-house needs. The propagation of process control technology from the academic realm to the industrial shop floor requires both vendor and user comprehension. Movement toward a strictly "vendor/supplier sells control system hardware/solutions/knowledge to corporate consumer" may appeal to the compartmentally minded. However, once the corporate user becomes ignorant in the technology, the synergy between process design, control, and operation is lost. The lack of process control talent in any of academic research, process control vendor, or corporate consumer is a weakness and handicap for all concerned.

Instantiation of Supervisory Control Systems

Underlying the application of advanced technology is the computational platform and distributed control system in place. Decisions of complexity and distributive reliability are important factors in the definition of scope for advanced control applications. As advanced control algorithms become available on regulatory level distributed control systems, the process control engineer is faced not only with a technology decision but also with a choice of vendor instantiation of advanced control technology. While algorithm fundamentals are published and well known, it is often the subtle modifications of published technology that make the technology actually work in practice. Each vendor claiming that their implementation is superior to their competitors can create a confusing climate that clouds the more important issues of control system objectives.

While it may seem that the development of process control technology is a mature area, the application of the technology available is quite young. The field is much bigger than the $3M project to apply MPC on the next mega-sized olefins plant. It is much bigger than the application on the "off-the-shelf" polymers facility. Mining the industry for valuable applications that may not be of the high throughput/low margin genre is widespread for exploitation. However, to do this effectively the application costs must be low. Eastman has written its own MPC code to enable "free" replication of MPC technology in addition to the learning of the technology that comes with such an endeavor. We have applied the infinite horizon model predictive control algorithm as described by Muske and Rawlings (1993). The IHMPC algorithm is based on a state space description of the process. A Kalman filter is used as a plant observer to reconstruct plant states, a quadratic program formulation is used to determine steady state process targets, and an infinite horizon linear quadratic regulator problem is solved to determine process inputs. Additional implementation details are described in Downs and Vogel (2001). Certainly, as Qin and Badgwell (1997) point out, there are numerous implementations of model predictive control algorithms. With as much research and development effort that model predictive control has commanded it would be a shame if questions such as which implementation of MPC to use, implementation costs, etc. inhibited the harvesting of the fruit this technology offers.

Our experience suggests that there are numerous good applications that require MPC to be integrated with other process control technologies. This integration demands an understanding of our chemical processes, their regulatory control strategies, the array of process control technologies available, and how to apply various technologies effectively. The variety of process control needs, process control technology, and the underlying hardware available have led to an increased need for broad based, knowledgeable process control talent.

Motivating Example—A Low Selector

The linkage between control strategy design and model predictive control has provided insight into the problem formulation and design of advanced control systems. The examples in this paper are intended to highlight the variety of application needs, MPC control structure, and the control strategy viewpoint. A motivating example centered on a distillation column control problem demonstrates the relationship between control strategy design and MPC design. Each example illustrates the importance of control strategy concepts when developing good MPC problem statements.

The concept of controlling unit operations within process constraints has been around for many years. The use of high and low selectors to prevent an operation from violating constraints can be viewed as a control strategy change agent. Consider the column illustrated in Figure 1. One control objective is to manipulate steam flow rate to control the underflow composition. If the column feed rate becomes large enough then the steam rate may increase to the point of flooding the column. A low select can be used to choose the lower of two desired steam flow set points, that requested by the column delta pressure controller or that requested by the underflow composition controller.

The low selector changes the column control strategy from an underflow composition to steam strategy to a column delta pressure to steam strategy. In the former,

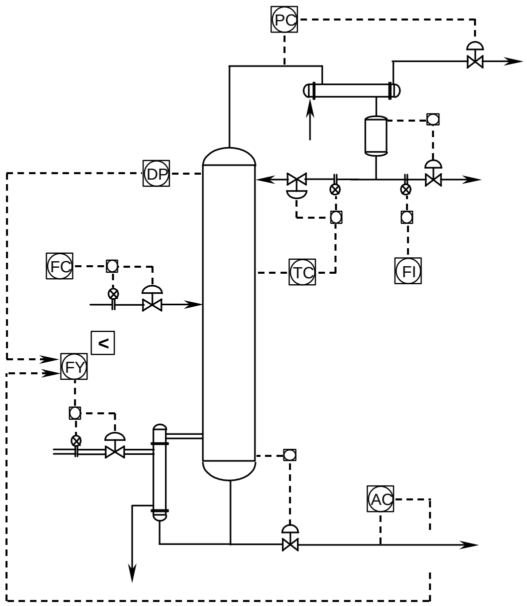

Figure 1: Distillation column with low selector on heat input.

the column loading will change in an attempt to maintain a relatively constant underflow composition. In the latter, the underflow composition will vary and yield to maintaining the column loading at a maximum value.

The steady state that this column will approach is apparent in this simple example, as are the two control loops that can be invoked. The two control loops, composition to steam and delta pressure to steam, will have different dynamics and may need to be tuned differently. Other control blocks can be added to make the transition from one strategy to another a smooth one. Current distributed control systems usually handle initialization of the non-selected controller. The influence of other control loops such as the temperature to reflux rate loop could also be incorporated. Other variables that depend on the steam rate and that need to be maintained within constraints could be added to the low selector. For example, the distillate flow rate may feed a downstream operation that has a maximum feed rate limit. If the column control demands a distillate rate that exceeds this limit, we may instead want to give up on the underflow composition to maintain temperature control. The list of possible constraint additions obviously could go on. As plant designs become more integrated this type of constraint escalation becomes more prevalent.

For the case of one manipulated variable, each constraint represents a different control strategy. However, if there is more than one manipulated variable then the number of possible control strategies is much larger. The understanding of what control strategy might be instantiated at any given time is an integral part of the design of the high and low selectors. Each possible pairing can be examined and verified for practical sense. If the number of possible strategies becomes too large to be reasonably evaluated, the high/low selectors are reconsidered and alternatives to achieve the control objective are developed. There is a self-regulating nature to the control design process—if the strategy becomes too complex to understand all that might happen, then simplifications are made.

Contrast the high/low selector design process to the model predictive control application mentality of today. It is so easy to add input or output constraints that a complete analysis of the resulting controller can become practically impossible. The unusual controller pairings that may result can be quite unexpected. Viewing model predictive control as a control strategy change agent can lead to insight into what the controller might end up controlling with what. This insight can provide guidance in what dynamic relationships are important in controller performance and robustness. Furthermore, do the operation regimes where the controller may end up make sense—even if they are stable? At Eastman we spend considerable time on determining why an undesirable outcome has occurred only to find out that the controller has done exactly what we programmed it to do. It has become evident to us that viewing MPC in light of control strategy design has made our MPC design job much easier and more intuitive.

Multivariable Control Applied to a Distillation Column

Problem Statement

Using high and low selectors for constraint control when more than one manipulated variable is involved quickly leads to application of MPC to more easily manage process constraints. When viewed as a control strategy change agent it is realized that the number of different control strategies that can be active at any one time is large. Each of these strategies can be evaluated based on numerous tools that have been developed over the years such as RGA, Niederlinski Index, SVD, etc. (Bristol, 1966; McAvoy, 1983). The issues of control loop interaction, degeneracy of degrees of freedom, and sensitivity to model error that control strategy analysis tools address can be applied to understand underlying problems in MPC applications.

Our approach in using MPC is one of understanding the control strategy that we want to invoke and how we want that strategy to change under different operating scenarios and then using MPC to accomplish this. As a result our MPC applications are studied more from the steady state viewpoint than a dynamic one. Certainly

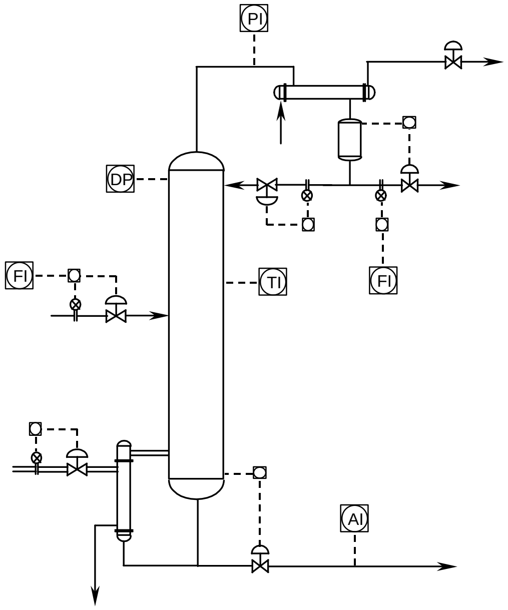

Figure 2: Distillation column.

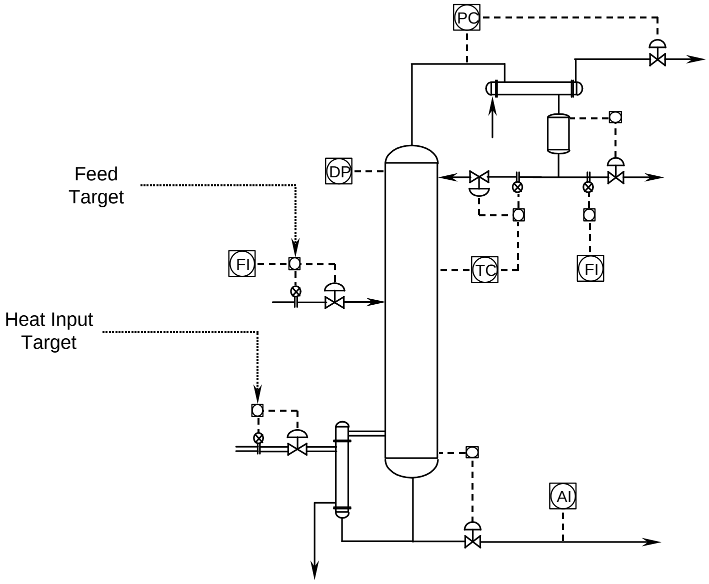

Figure 3: Distillation column hybrid control strategy.

we have cases where the dynamic advantages of MPC are exploited, but we have found that the steady state features are of most benefit.

Another important design consideration for supervisory control that employs model predictive controllers is the controller architecture. At one end of the spectrum is the flat architecture that has all the measurements and manipulated variables in one MPC controller. This structure takes on the appearance of a "black box" and it is sometimes difficult to diagnose underlying controller problems. At the other end of the spectrum is a vertical architecture that resembles a multi-layer cascade structure. This structure has the advantage of building control strategies using conventional process control notions and of segregating unrelated parts of the controller. Using the vertical structure, however, requires that the issues of controller initialization, constraint passing between layers, and controller speeds of response be managed.

To illustrate some of these issues consider the distillation column illustrated in Figure 2. We assume that process analysis has been completed to determine that the following process objectives are to be achieved:

1. Maintain tray temperature in the rectifying section at set point
2. Maintain underflow composition at set point
3. Maintain feed rate at set point
4. Maintain column pressure drop less than a given maximum
5. Maintain the column distillate rate less than a given maximum
6. Minimize energy usage

We have reflux rate, heat input rate, feed rate and cooling rate available to manipulate. Of course, even getting to this step required a decision to control reflux drum level with distillate rate and reboiler level with underflow rate. These controllers could also be added to the control strategy development problem but will be assumed as given here. Developing a control strategy to achieve these objectives is incomplete until we know what set points to give up on if the column becomes constrained—a ranking of importance is also required. For this example we will assume that the temperature in the rectifying section is the most important followed by column feed rate followed by underflow composition, which is the least important.

Case 1—A Hybrid Strategy

Consider first a hybrid strategy that employs an underlying SISO strategy that is illustrated in Figure 3. This may be a strategy that has been successfully used for many years and is effective in maintaining the rectifying temperature at its set point. Our job is simply to achieve the stated objectives by overlaying an "advanced control system" above the regulatory SISO strategy. This is common when the regulatory strategy is sound and provides stabilizing control during periods of process upsets, start-ups, etc.

A steady state process gain matrix for the resulting variables is

	Heat Input	Feed Rate
Underflow Composition	-1	0.5
Feed Rate	0	1
Distillate Rate	0.01	0.5
Pressure Drop	1	0

During normal operation with no constraint active an MPC controller will line out with the underflow composition and feed rate at set point and the distillate rate and pressure drop within limits. The control strategy active at this time has the steady state gain matrix,

	Heat Input	Feed Rate
Underflow Composition	-1	0.5
Feed Rate	0	1

Clearly we can see that the resulting MPC controller will look a lot like underflow composition to heat input and feed rate set constant with a feed forward term between feed rate and heat input. Interaction measures would say that this strategy should work fine—in fact, the relative gain for each loop is equal to one.

Consider how operations change if the pressure drop constraint becomes active. If we are to give up on underflow composition first and maintain feed rate then we end up with a gain matrix,

	Heat Input	Feed Rate
Feed Rate	0	1
Pressure Drop	1	0

which again indicates that MPC will work well. However, if it is desired to give up instead on feed rate and maintain underflow composition then the gain matrix is

	Heat Input	Feed Rate
Underflow Composition	-1	0.5
Pressure Drop	1	0

and we will have a more difficult control problem. In fact the model relating underflow composition to feed rate at constant heat input becomes more important because it is the only link that the feed rate has into the control strategy. During normal operation this model only influences the feedforward relationship between heat input and feed rate whereas in this constrained case it is the primary relationship for underflow composition control. This difference in control problem characteristics resulted from a change in the steady state weighting of the controlled variable importance. Certainly, this is an innocent change that has important ramifications on the resulting control problem.

Next consider the case where the distillate rate is constrained and again where we are to give up on underflow composition and maintain feed rate. This scenario yields the following gain matrix

	Heat Input	Feed Rate
Feed Rate	0	1
Distillate Rate	0.01	0.5

which is almost degenerate. Large steady state heat input changes are needed to have any effect on the distillate rate. Certainly for this simple example, process insight might key us into the fact that the distillate rate and the feed rate are so closely tied together that this requirement is unreasonable. However, this case was quite reasonable when it was the column pressure drop that was the constraint instead of the column distillate rate. This behavior points to the fact that steady state weighting preferences may easily lead to difficult dynamic control problems that have poor characteristics regarding interaction or robustness.

Finally consider the case where both constraints are active. The gain matrix again becomes docile and well behaved.

	Heat Input	Feed Rate
Distillate Rate	0.01	0.5
Pressure Drop	1	0

The heat input maximizes the column pressure drop to keep the underflow composition as close to target as possible and the feed rate is reduced to keep the column distillate rate within its limits. If we knew these constraints were always going to be active, this may even be a strategy we would design.

Each constraint scenario yields a different control strategy; some of these strategies are well behaved and some are clearly not strategies that we would want to deploy. The transition between different strategies is seamless and it may appear that if you can describe the control objective in terms of an MPC structure your problems are over—simply configure, tune, and start counting the savings. However, this is not the case and for MPC controllers that have these time bombs buried within them, it usually happens that the poor, unforeseen strategy gets invoked late at night or on holiday weekends.

Case 2—A Less Hybrid Strategy

If the temperature in the rectifying section of the column is in fact less important than the feed rate or underflow composition, then it is advantageous to move that control loop from the regulatory level to the supervisory level. This allows the importance of holding the temperature at set point to be given a lower weight in the steady state target calculation. As long as the temperature is controlled at the regulatory layer, the reflux will change in an attempt to get the temperature to target even though this may have a low priority. Consider a different hybrid strategy that employs an underlying SISO structure that is only used to handle level control loops. In this case the rectifying temperature and reflux rate are

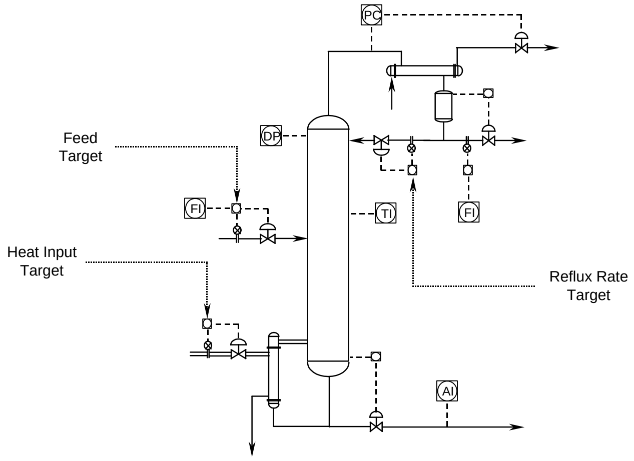

Figure 4: Distillation column hybrid control strategy with reflux rate in supervisory control layer.

in the supervisory layer. A candidate strategy is illustrated in Figure 4. There is advantage to leaving inventory loops at a PID level because it keeps the supervisory controller dealing only with self regulating loops and can help avoid reliability issues around advanced control systems. Our job is simply to achieve the stated objectives by overlaying an "advanced control system" above the regulatory control strategy. It is noted, however, that once the underlying inventory control strategy is chosen, many options for the overall control are eliminated.

A steady state process gain matrix for the resulting variables is given by

	Heat Input	Feed Rate	Reflux Rate
Underflow Composition	-2	0.3	1
Feed Rate	0	1	0
Temperature	1	0.1	-0.8
Distillate Rate	2	0	-1
Pressure Drop	1	0	0.1

When no constraints are active, controlling the underflow composition, feed rate, and temperature using the heat input set point, feed rate set point, and the reflux rate set point leads to a well-behaved process gain matrix. MPC in this case simply provides mild decoupling of the implied SISO loops. The RGA for this case is

	Heat Input	Feed Rate	Reflux Rate
Underflow Composition	2.66	0	-1.66
Feed Rate	0	1	0
Temperature	-1.66	0	2.66

and using it to pair loops one can envision an SISO strategy as shown in Figure 5.

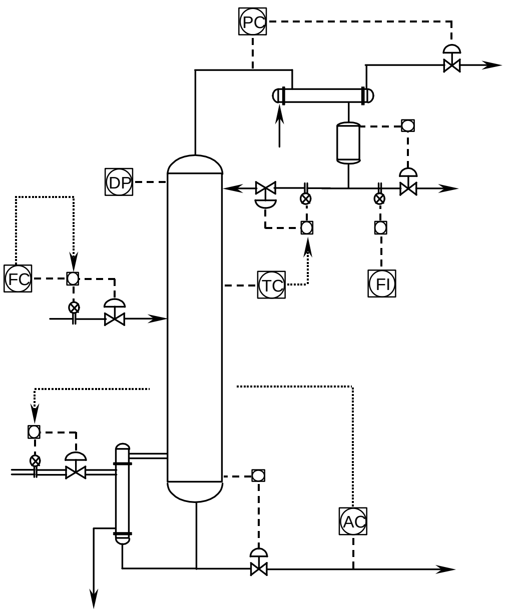

Figure 5: Distillation column control strategy during unconstrained operation.

When constraints become active, the implied control strategy changes and the resulting implied strategy is dependent on the importance placed on the variables having set points, that is, which controlled variables are allowed to deviate from their set point. Using a linear program (LP) for the determination of steady state operation results in answers that lie at a vertex and an implied control strategy that has controlled variables that do not line out at set point entering the problem in a sequential manner as constraints become active. Using a quadratic program (QP) for the determination of steady state operation results in answers that can look blended. For example, two controlled variables can be allowed to deviate from set point equally and can enter the problem in a parallel manner. There has been recent work to explore the steady state target problem formulation and calculation (Kassmann et al., 2000).

Consider the case where the column pressure drop constraint becomes active and the column rectifying temperature is to be allowed to deviate from set point. The resulting controller gain matrix is

	Heat Input	Feed Rate	Reflux Rate
Underflow Composition	-2	0.3	1
Feed Rate	0	1	0
Pressure Drop	1	0	0.1

An RGA calculation of this matrix is given by

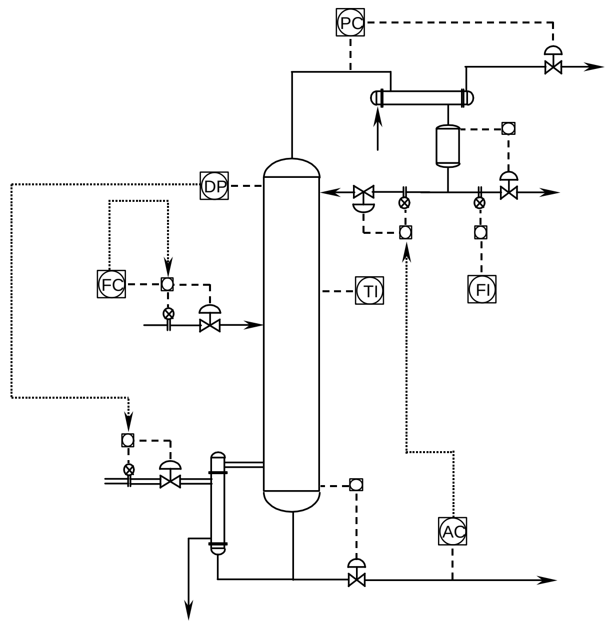

Figure 6: Distillation column control strategy when delta pressure becomes constrained

	Heat Input	Feed Rate	Reflux Rate
Underflow Composition	1/6	0	5/6
Feed Rate	0	1	0
Pressure Drop	5/6	0	1/6

and suggests a control strategy as shown in Figure 6. From the gain matrix we can see that this pairing is a pretty obvious one if the variables we are required to control are the ones shown. However, presented with this control design problem from an SISO point of view, we would probably be looking for alternatives—controlling the underflow composition with the reflux just doesn't look too promising. If we expect this case to occur then we would probably want to spend additional effort determining the relationship between reflux and underflow composition when the feed and pressure drop are constant. This is, of course, different than the relationship determined during open loop testing when the feed and heat input are constant. The understanding of what control strategies can look like under different constraint scenarios leads to insight into why an MPC controller might fail or perform poorly.

Finally, the question of including the reflux drum level, column reboiler level, and column pressure in the supervisory control layer must also be addressed. Including the level measurements and their control in MPC leads to handling a mix of self-regulating and integrating variables. It is not clear to us which predictive control technologies on the market are equipped to handle this case. Depending on the column reflux ratio, which may change during the course of operation, the level control strategy may be best left alone and on the regulatory layer or it may be paramount that it to be given over to the supervisory controller.

Incorporating the column pressure control into the supervisory layer may at first seem unwise. However, the ability to change the operating pressure of the column can lead to increased energy efficiency provided it can be done in a coordinated way with the other column controls. That, of course, is exactly what MPC does.

Process Applications

This paper contains examples of the variety of applications benefiting from our use of model predictive control. Our successful record of gaining benefit from this technology has relied upon several basic tenets. First, our ability to develop good regulatory control strategies has provided a solid foundation on which to build higher-level supervisory control systems. The benefit and results from this step sometimes indicate that this is all that needs to be done. Second, the identification of good advanced control candidates has required an understanding of the process economics to screen for high value applications. Third, the costs of solution development and implementation has been kept low and not hindered the "leap of faith" often required of operations. Fourth, a building block approach to reaching intermediate process control milestones has led to increasing complexity and value that could not be envisioned at the start of projects. Fifth, the integration of process improvement functions already in place (e.g. design of experiments, equipment design, process chemistry experiments) has led to control objectives that were unknown at project initiation. Finally, maintenance of our applications has led to new opportunities and relationships that have in turn grown this aspect of our work.

Reactor Product Crystallization Train

Illustrated in Figure 7 is a common situation where the control strategy needs to change during operation. Consider the problem of controlling the four crystallizer levels and a total throughput rate using the five manipulated variables shown. A common SISO strategy, Figure 8, would be to set the throughput rate using the reactor feed and then have a level to outflow pairing for the four crystallizers. This strategy has two problems. First the process variation will be propagated downstream and the fourth and perhaps the most important crystallizer will be the one getting the most variability in its feed flow rate. Second, if the process bottleneck is somewhere other than the reactor, then the throughput rate needs to be lowered enough to insure that the valve or crystallizer that is the bottleneck does not exceed its capability when throughput rates cycle through a maximum.

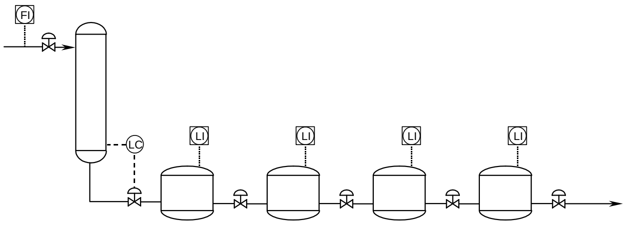

Figure 7: Reactor followed by crystallization train.

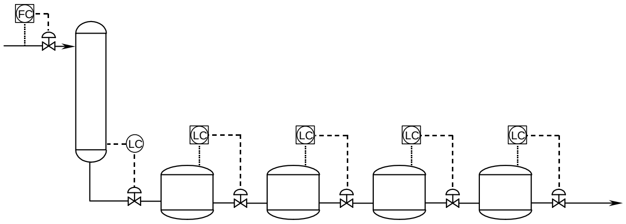

Figure 8: Original control strategy for reactor followed by crystallization train.

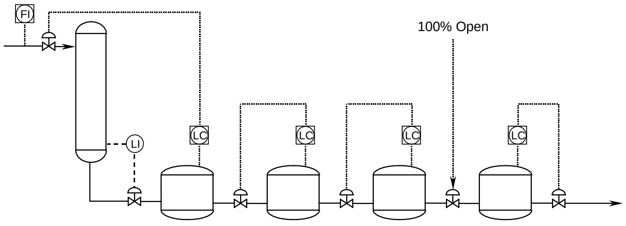

Figure 9: Level control strategy for reactor followed by crystallization train when downstream valve becomes constrained.

A model predictive controller has the ability to change the level control strategy as needed via the constraint handling. If the feed to the fourth crystallizer is the process constraint then the level control strategy becomes that shown in Figure 9. Of course this constraint, which may be more complicated than a simple valve limit, can move to different locations and a model predictive controller can accommodate this. Another advantage of a model predictive controller in this application is that it can be tuned to distribute the variability to the units that are least upset by flow variations. The level control variability in this example can be directed more toward the first and second crystallizers and away from the later ones.

There are implementation considerations needed to maintain operation when more than one manipulated variable becomes constrained or gets put into manual. Procedures are needed to control the process to some extent when no steady state solution to the model equations exists. This example is one of many where management of process inventories is important. These inventory control problems can span a single process like this example or cover large networks of in-process tankage. The ability to handle integrating variables and to distribute level and flow variability is important in this category of problems. Benefits of reduced flow variability often translate into increased production rates. The ability to handle unit operation feed constraints that move from unit to unit based on processing conditions is also an important benefit arising form this type of problem.

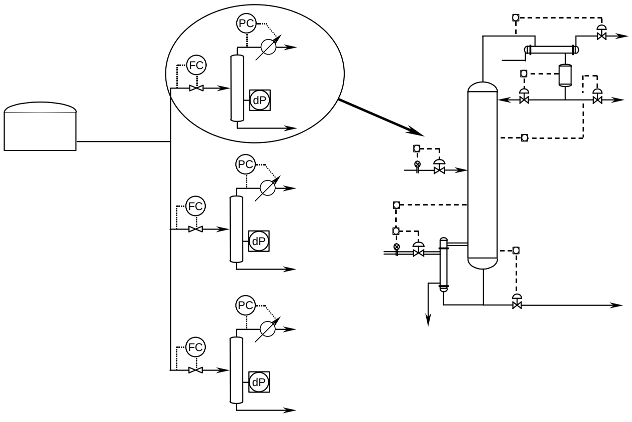

Figure 10: Parallel distillation column loading.

Parallel Distillation Column Loader

Illustrated in Figure 10 is a common situation where several parallel unit operations need to be used in an efficient manner. In this example there are three isomer separation distillation columns that process a reactor effluent. The control objective is to maintain the total feed to the system at a specified target and to load the columns in an efficient manner. Each column has an effective SISO control strategy that controls end compositions by manipulating distillate rate and heat input. Manipulating cooling duty controls the operating pressure of each column. The feed rate capacity as measured by column differential pressure and the separation efficiency are a function of the column operating pressure.

The control objective can be met by manipulating the feed rate to each column and the operating pressure of each column. Certainly other choices can be made. In particular, the depth to which the individual column regulatory control strategy is included in the model predictive controller is an important decision. The column regulatory controls will have to respond when the column operating pressure is changed. If this is expected to be a slow change the supervisory system may simply layer

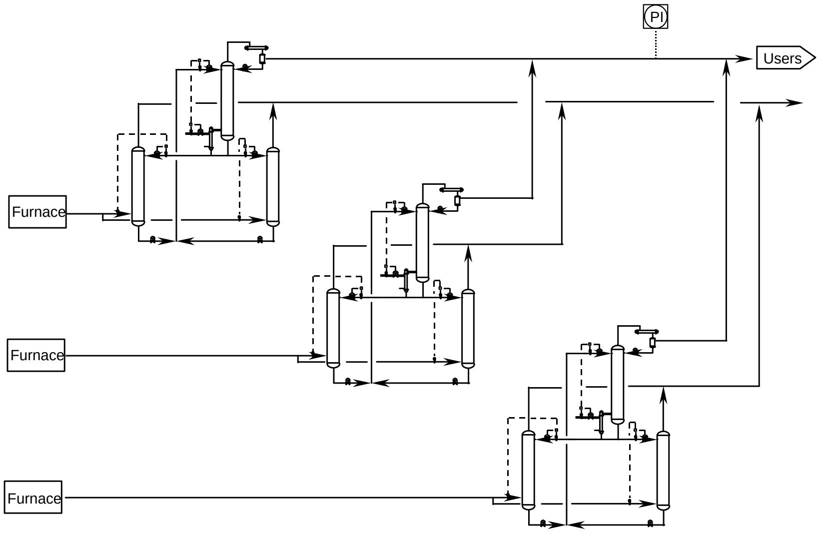

Figure 11: Process gas absorber/regenerator systems.

on top of the existing regulatory control. However, if the total feed rate and hence the pressure are expected to be changed more quickly, then the composition loops may need to be incorporated into the supervisory layer. If the economics strongly suggest that operation at minimum pressure is required, then the notion of running each column at its maximum pressure differential and manipulating the column pressures to control feed rates is not too far away. Unusual strategies like this one require process understanding to uncover pitfalls and unusual unit operation behavior that may make such a suggestion laughable.

This type of process loading to parallel unit operations is common. Usually the parallel operations have differing efficiencies that can be determined to minimize the processing costs. Often the operating efficiencies of the units are a function of how loaded the unit is. The efficiency often goes through a maximum indicating more efficient operation at higher loads up to a point after which efficiency drops off, usually very quickly.

Flue Gas CO_2 Absorber Control

A similar but different situation is illustrated in Figure 11 where three CO_2 absorber systems recover CO_2 from three different furnaces. There is a varying demand for recovered CO_2. The control objective is to recover the demand amount of CO_2 at the minimum costs. Each system has a different recovery efficiency and also has varying amounts of CO_2 that are available for recovery. Each recovery system has its own process constraints that must be honored.

A model predictive controller can be employed to manage the system. There are at least two major model predictive control strategies that are suggested. One is a flat, horizontal architecture and the other is a vertical architecture. The horizontal architecture has all the manipulated variables for each recovery system in the same MPC. This has the advantage of making all the information available in one MPC. As constraints become active in one system this information is part of the MPC calculations for the other systems. The downside is that changes in one system directly influence the other systems when local handling of disturbances might be a better alternative. In addition, there is always the possibility of a system being down or off-line requiring it to be removed from the controller.

Process control strategy notions suggest measurements that naturally reject some common process disturbances. For example, controlling percent CO_2 recovery, Figure 12, for a system may reject most of the feed rate disturbances that a local system may experience without propagating them to the other systems. Similarly, energy efficiency may also be normalized by feed rate. The characteristics and patterns of the CO_2 users can also be incorporated into the design of measurements that reject common disturbances. Process knowledge that absorber/regenerator systems of this type are suitable for ratio-oriented strategies can unload the supervisory control system to perform primarily the optimization work that it can do best.

Energy Recovery Pressure Controller

Illustrated in Figure 13 is a process environmental control/energy recovery situation where the control strategy needs to change during operation. In this example a process effluent gas stream contains components that need

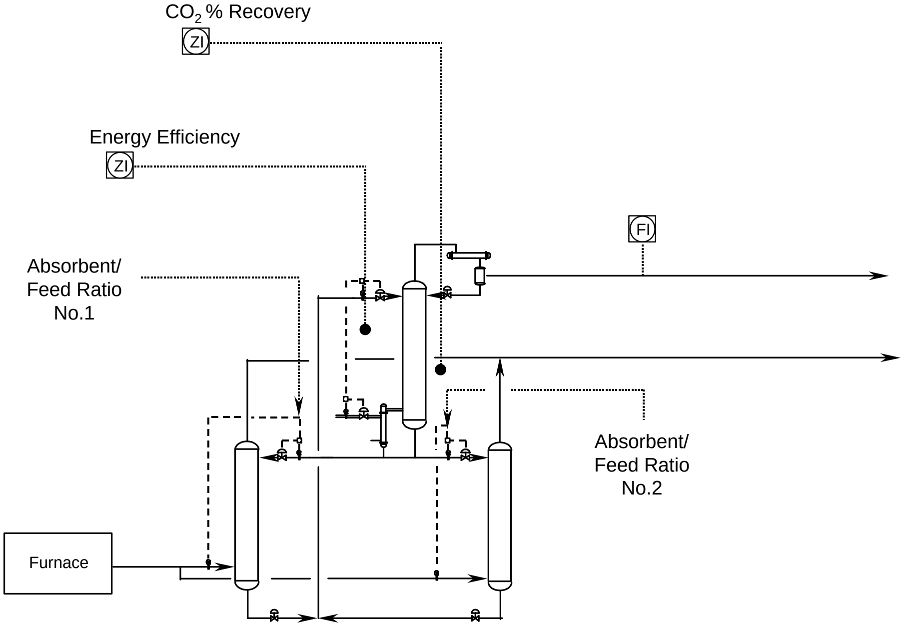

Figure 12: Individual process gas absorber/regenerator system.

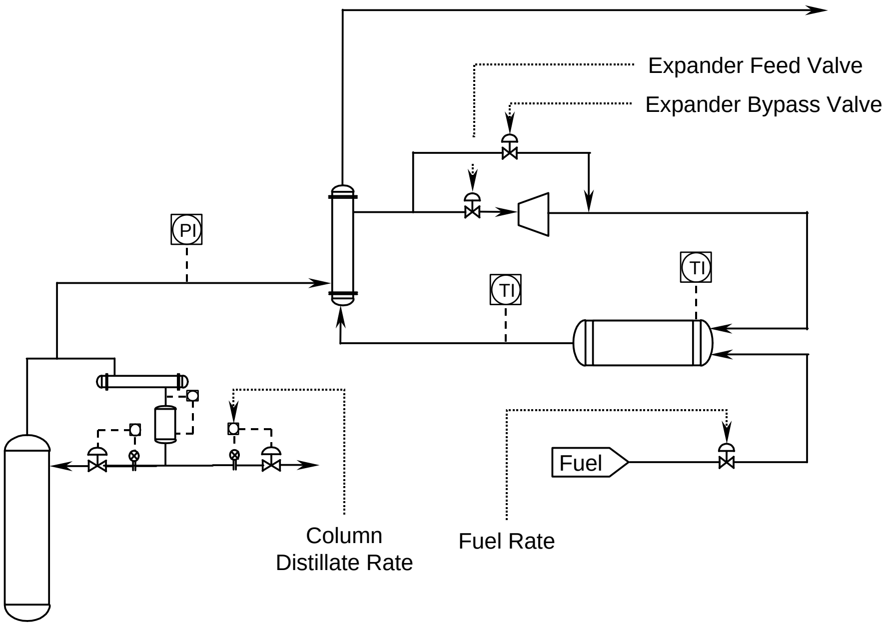

Figure 13: Process effluent catalytic oxidation process.

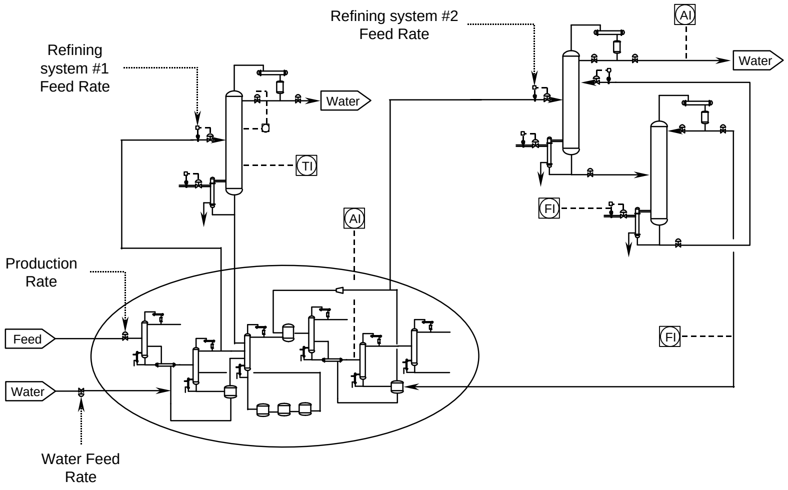

Figure 14: Plant wide water balance control.

	Refining System #1 Feed	Refining System #2 Feed	Process Production Rate	Process Water Feed (measured disturbance)
Composition	-1	-2	1	1
Temperature in Refining System #1	-2	0	-0.2	-0.1
Steam Flow in Refining System #2	0	1	0	0
Composition in Refining System #2	0	0.2	0	0
Recycle from Refining System #2	0	0.5	0	0

Table 1: Steady state process gain matrix for plant wide water balance control.

to be catalytically oxidized. The effluent gas comes from the top of a distillation column whose condenser can condense and remove some of the effluent as liquid products. The uncondensed gas is routed to a gas expander used to recover energy from this high pressure stream before it is sent to the catalytic oxidizer. The gas expander cools the gas. There is natural gas fuel that inexpensively preheats the gas to temperatures needed for catalytic oxidation to occur. There are maximum and minimum inlet and outlet temperatures that must be honored to insure proper component destruction.

The process gain matrix is given by

	Column Distillate Rate	Fuel Rate	Expander Feed Valve	Expander Bypass Valve
Pressure	-2	0	-1	-1
Inlet Temp.	0.5	2	-1	0
Exit Temp.	0.5	2	-1	0

This example incorporates the fast process dynamics associated with pressure control of gas systems. The incorporation of pressure control overrides for safety systems and the combustion control system used for the fuel can complicate the design of the advanced control strategy. The use of the advanced control system is to layer on top of the existing safety and burner control systems and not compromise their operation. In fact, these safety systems may operate on process control hardware that is distinct and loosely linked to the platform used for supervisory control. These issues may dictate the incorporation of information indicating how such systems are interacting with the supervisory control, if at all. Understanding the process control hardware, the process operational requirements, and the safety and environmental consequences are as much a part of the control system design as the control technology.

Plant Wide Water Balance Control

Illustrated in Figure 14 is a plant wide control situation where an overall control strategy needs to change during operation. In this example a plant contains a component that travels throughout the process. The component in this case is water and it is in a plant feed, is produced by reaction and is removed via two distillation systems. Refining system #1 is a simple distillation column and refining system #2 is a combination of columns. The cost of removal is different for each system and each system has process constraints that must be honored. An important control issue in this case is what measurement or combination of measurements indicates the status of the water balance. The measurement used in this case is an on-line analyzer on a key stream within the process.

The objective of a model predictive controller is to maintain the water in balance by manipulating the feed rates to the two refining systems and, if needed, the overall plant production rate. The steady state process gain matrix is shown in Table 1.

The complications that arise in this example are: (1) the process time frames are widely diverse, (2) the process economics suggest a preferred use of the manipulated variables, and (3) the best approach to identify the status of the plant water composition may be unclear. These issues can be addressed using a layered, vertical hierarchy or a flat, "all-in-one" strategy. How one chooses to design the structure depends on not only understanding the process control ramifications of each design but also the appropriateness and timing of making changes in one part of the plant in response to upsets in another part of the process.

The ability to add additional process constraints as they are discovered integrates a longer term support role for applications of this type. While the initial installation may incorporate only the constraints listed in the process gain matrix, good process improvement work will probably eliminate them and identify new unforeseen constraints. This in turn will require the identification of new models and support for this application. For many of our applications process improvement work thrives when process constraints are clearly identified and operated against. This creates an environment of process development whose costs and benefits can clearly be identified and realized.

Conclusion

The application of process control and, in particular, model predictive control remains an active and profitable area for the chemical industry. Considerable progress has been made to provide a theoretical foundation for model predictive control and to move it into the mainstream of application. As this technology becomes more widespread the implementation issues encountered every day will begin to be addressed and this technology will mature into a powerful tool with routine application. In the meantime there remains much to develop to reach this destination. As demonstrated in the process examples, there are numerous application issues that complicate the use of MPC. These issues include controller tuning, characterizing complex performance criteria, using redundant process information, and controller robustness. The solution to these issues requires identification of the problem, short term and possibly ad-hoc fixes, and research to address the problem in more fundamental ways. The need for process control talent for academic research, process control vendors, and corporate consumers alike remains strong. Each group brings a unique and indispensable viewpoint to the effective application of process control.

Incorporating the process control strategy viewpoint into advanced control design has provided Eastman with a very high success rate when applying advanced control technology. The understanding and incorporation of process knowledge continues to be invaluable in the successful application of new control techniques. The advent and routine use of model predictive control has not supplanted the need nor the value of process understanding in the successful application of process control technology. On the contrary, model predictive control has amplified the need for process and control strategy analysis and understanding.

References

Barolo, M. and C. A. Papini, "Improving Dual Composition Control in Continuous Distillation by a Novel Column Design," *AIChE J.*, **46**(1), 146–159 (2000).

Bristol, E. H., "On a new measure of interaction for multivariable process control," *IEEE Trans. Auto. Cont.*, pages 133–134 (1966).

Downs, J. J. and J. E. Doss, Present status and future needs—a view from North American industry, In Arkun, Y. and W. H. Ray, editors, *Chemical Process Control CPC IV*, pages 53–77, Austin, TX. CACHE (1991).

Downs, J. J. and E. F. Vogel, Industrial Experience with State-Space Model Predictive Control, In Rawlings, J. B., B. A. Ogunnaike, and J. W. Eaton, editors, *Chemical Process Control CPC VI*, Austin, TX. CACHE (2001).

Downs, J. J., Private communication with PIP representatives in process control (2000).

Groenendijk, A. J., A. C. Dimian, and P. D. Iedema, "Systems Approach for Evaluating Dynamics and Plantwide Control of Complex Plants," *AIChE J.*, **46**(1), 133–145 (2000).

Kassmann, D. E., T. A. Badgwell, and R. B. Hawkins, "Robust steady-state target calculation for model predictive control," *AIChE J.*, **46**(5), 1007–1024 (2000).

Luyben, W. L., "Plantwide Design and Control of Processes with Inerts. 1. Heavy Inerts.," *Ind. Eng. Chem. Res.*, **37**, 528–534 (1998a).

Luyben, W. L., "Plantwide Design and Control of Processes with Inerts. 1. Intermediate Inerts.," *Ind. Eng. Chem. Res.*, **37**, 535–546 (1998b).

Luyben, W. L., "Plantwide Design and Control of Processes with Inerts. 1. Light Inerts.," *Ind. Eng. Chem. Res.*, **37**, 516–527 (1998c).

Luyben, W. L., "Design and Control of Gas-Phase Reactor/Recycle Processes with Reversible Exothermic Reactions," *Ind. Eng. Chem. Res.*, **39**, 1529–1538 (2000).

McAvoy, T. J., *Interaction Analysis—Principles and Applications*. Instrument Society of America, Research Triangle Park, NC (1983).

Muske, K. R. and J. B. Rawlings, "Model Predictive Control with Linear Models," *AIChE J.*, **39**(2), 262–287 (1993).

Qin, S. J. and T. A. Badgwell, An overview of industrial model predictive control technology, In Kantor, J. C., C. E. García, and B. Carnahan, editors, *Proceedings of Chemical Process Control—V*, pages 232–256. CACHE, AIChE (1997).

Tseng, J. L., W. R. Cluett, and W. L. Bialkowski, "New Analysis and Design Tool for Achieving Low Variability Process Designs," *AIChE J.*, **45**(10), 2188–2202 (1999).