Plantwide Control Design Procedure

Michael L. Luyben and Bjorn D. Tyreus
DuPont Central Research & Development, Wilmington, DE 19880
William L. Luyben
Dept. of Chemical Engineering, Lehigh University, Bethlehem, PA 18015

Plantwide control involves the systems and strategies required to control an entire chemical plant consisting of many interconnected unit operations. A general heuristic design procedure is presented that generates an effective plantwide control structure for an entire complex process flowsheet and not simply individual units. The nine steps of the proposed procedure center around the fundamental principles of plantwide control: energy management; production rate; product quality; operational, environmental and safety constraints; liquid-level and gas-pressure inventories; makeup of reactants; component balances; and economic or process optimization. Application of the procedure is illustrated with three industrial examples: the vinyl acetate monomer process, the Eastman plantwide-control process, and the HDA process.

Introduction

Control analysis and control system design for chemical and petroleum processes have traditionally followed the "unit operations approach" (Stephanopoulos, 1983). First, all of the control loops were established individually for each unit or piece of equipment in the plant. Then the pieces were combined together into an entire plant. This meant that any conflicts among the control loops somehow had to be reconciled. The implicit assumption of this approach was that the sum of the *individual* parts could effectively comprise the *whole* of the plant's control system. Over the last few decades, process-control researchers and practitioners have developed effective control schemes for many of the traditional chemical unit operations. For processes where these unit operations are arranged in series, each downstream unit simply sees disturbances from its upstream neighbor.

Most industrial processes contain a complex flow sheet with several recycle streams, energy integration, and many different unit operations. Essentially, the plantwide control problem is how to develop the control loops needed to operate an *entire* process and achieve its design objectives. The problem is extremely complex and is very much open-ended. There are a combinatorial number of possible choices and alternative strategies. And there is no unique "correct" solution.

Certain features of complex and highly integrated processes are unique from the viewpoint of plantwide control. Recycle streams and energy integration introduce a feedback of material and energy among units upstream and downstream. They also interconnect separate unit operations and create a path for disturbance propagation. The presence of recycle streams profoundly alters the dynamic behavior of the plant by introducing an integrating effect that is not localized to an isolated part of the process.

Despite this process complexity, the unit operations approach to control-system design has worked reasonably well. In the past, plants with recycle streams contained many surge tanks to buffer disturbances, to minimize interaction, and to isolate units in the sequence of material flow. This allowed each unit to be controlled individually. Prior to the 1970s, low energy costs meant little economic incentive for energy integration. However, there is growing pressure to reduce both capital investment and working capital and to improve safety and environmental concerns. This has prompted design engineers to start eliminating many surge tanks, increasing recycle streams, and introducing heat integration. Often this is done without a complete understanding of the effect on plant operability.

So economic forces within the chemical industry are compelling improved capital productivity. Requirements for on-aim product quality control grow increasingly tighter. More energy integration occurs. Improved product yields provide additional plant capacity by lower reactant per-pass conversion and higher material recycle rates. These are all economi-

---

Correspondence concerning this article should be addressed to M. L. Luyben.

cally attractive in the steady-state flow sheet, but they present significant challenges to smooth dynamic plant operation.

Buckley (1964) proposed a control design procedure for the plantwide control problem that consisted of two levels. The first determined the material balance control structure to handle vessel inventories for low-frequency disturbances. The second established the product quality control structure to regulate high-frequency disturbances. This procedure has been widely and effectively utilized. It has served as the conceptual framework in many subsequent ideas for developing control systems for complete plants. However, this Buckley procedure provides little guidance concerning three important aspects of a plantwide control strategy. First, it does not explicitly discuss energy management. Second, it does not address the specific issues of recycle systems. Third, it does not deal with component balances in the context of inventory control. By placing the priority on material balance over product quality controls, the procedure can significantly limit the flexibility in choosing the latter.

Reaching a solution to the complex plantwide control problem is a creative challenge. It demands insight into and understanding of the chemistry, physics, and economics of real processes. Stephanopoulos (1983) clearly summarized the creativity involved in control-strategy synthesis for chemical plants. However, there is still little guidance in the literature on designing workable plantwide control systems. We present here such a systematic design procedure. Our framework in tackling a problem of this complexity is based upon heuristics that account for the unique features of plantwide control. The scope embraces continuous processes with reaction and separation sections. The procedure is based upon research by us and others, and our governing philosophy is *it is always best to utilize the simplest control system that will achieve the desired objectives.* Three industrial processes illustrate its application in this article. In addition, the procedure has been used successfully over several years in a number of commercial processes at DuPont.

Previous Work

Some recent work in the area of plantwide control includes that of Price and Georgakis (1993) and Price et al. (1994). They employed a tiered framework, following Buckley, that classifies the control loops according to their importance for plant operation to reach a "self-consistent" structure that minimizes the propagation of disturbances. Narraway and Perkins (1993) formulated an optimization problem that evaluates control structures based upon the effect of disturbances on process economics. Wolff and Skogestad (1994) presented a number of guidelines in finding control strategies for integrated plants.

Studies of plantwide control for the Eastman process control challenge problem presented by Downs and Vogel (1994) include: McAvoy and Ye (1994), Lyman and Georgakis (1995), Kanadibhotla and Riggs (1995), Ricker and Lee (1995), Banerjee and Arkun (1995), McAvoy et al. (1996), and Ricker (1996).

This article draws from our previous work within the area of plantwide control, with some previous publications by Tyreus and Luyben (1993), Luyben and Luyben (1995), Yi and Luyben (1995), and Luyben et al. (1996).

One development that will facilitate research on the plantwide control problem is the commercial release of computer software needed for nonlinear plantwide dynamic simulations. We strongly believe that the final evaluation of any plantwide control structure requires rigorous nonlinear dynamic simulations, not linear transfer-function analysis.

Some people claim that the plantwide control problem has already been solved by the application of several commercial forms of model predictive control (MPC). Our understanding is that MPC has found widespread use in the petroleum industry. The chemical industry, however, is still dominated by the use of distributed control systems implementing simple PID controllers. We are addressing the plantwide control problem within this context. We are not addressing the application of multivariable model-based controllers in this article. We recommend a careful reading of Ricker (1996) concerning the application of MPC to a chemical process such as the Eastman problem.

Plantwide Control Design Procedure

Our plantwide control design procedure satisfies the two fundamental chemical engineering principles of the overall conservation of energy and mass. Additionally, the procedure accounts for nonconserved entities within a plant such as chemical components (produced and consumed) and entropy (produced). In fact, five of the nine steps deal with plantwide control issues that would not be addressed by simply combining the control systems from all of the individual unit operations. Steps 1 and 2 establish the objectives of the control system and the available degrees of freedom. Step 3 ensures that any production of heat (entropy) within the process is properly dissipated. In Steps 4 and 5 we satisfy the business objectives concerning production rate, product quality, and safety. Step 6 involves total mass balance control, whereas in Step 7 we ensure that nonconserved chemical components are accounted for. That concludes the plantwide control aspects. In Step 8 we complete the control systems for individual unit operations. Finally, Step 9 uses the remaining degrees of freedom for optimization and improved dynamic controllability. This heuristic procedure will generate a workable plantwide control strategy, which is not necessarily the *best* solution. Because the design problem is open-ended, the procedure will not produce a unique solution.

1. Establish Control Objectives. Assess the steady-state design and dynamic control objectives for the process. This is probably the most important aspect of the problem because different criteria lead to different control structures. These objectives include reactor and separation yields, product quality specifications, product grades and demand determination, environmental restrictions, and the range of safe operating conditions.

2. Determine Control Degrees of Freedom. Count the number of control valves available. This is the number of degrees of freedom for control, that is, the number of variables that can be controlled. The valves must be legitimate (flow through a liquid-filled line can be regulated by only one control valve).

3. Establish Energy Management System. We use the term "energy management" to describe two functions: (1) We must provide a control system that removes exothermic heats of reaction from the process. If heat is not removed to utilities directly at the reactor, then it can be used elsewhere in the process by other unit operations. This heat, however, must ultimately be dissipated to utilities. (2) If heat integration does occur between process streams, then the second function of energy management is to provide a control system that prevents the propagation of thermal disturbances and ensures that the exothermic reactor heat is dissipated and not recycled. Process-to-process heat exchangers and heat-integrated unit operations must be analyzed to determine that there are sufficient degrees of freedom for control.

Heat removal in exothermic reactors is crucial because of the potential for thermal runaways. In endothermic reactions, failure to add enough heat simply results in the reaction slowing up. If the exothermic reactor is running adiabatically, the control system must prevent excessive temperature rise through the reactor (e.g., by setting the ratio of the flow rate of the limiting fresh reactant to the flow rate of a recycle stream acting as a thermal sink).

Heat transfer between process streams can create significant interaction. For reactor feed/effluent heat exchangers it can lead to positive feedback and even instability. Where there is partial condensation or partial vaporization in a process-to-process heat exchanger, disturbances can be amplified because of heat of vaporization and temperature effects.

Increased use of heat integration can lead to complex dynamic behavior and poor performance due to recycling of disturbances. If not already included in the design, trim heaters/coolers or heat exchanger bypass lines must be added to prevent this. Energy disturbances should be transferred to the plant utility system whenever possible to remove this source of variability from the process units.

4. Set Production Rate. Determine what valve will be used to set throughput. Often design constraints require that production be set at a certain point. An upstream process may establish the feed flow sent to the plant. A downstream process may require on-demand production, which fixes the product flow rate from the plant.

If no constraint applies, then we select the valve that provides smooth and stable production-rate transitions and rejects disturbances. We often want to select a variable that has the least effect on the separation section, but also has a rapid and direct effect on reaction rate in the reactor without hitting an operational constraint. This may be the feed flow to the separation section, the flow rate of a recycle stream, the flow rate of initiator or catalyst to the reactor, the reactor heat removal rate, the reactor temperature, and so forth.

Throughput changes can only be achieved by altering, either directly or indirectly, conditions in the reactor. To obtain higher production rates, we must increase overall reaction rates. This can be accomplished by raising temperature (higher specific reaction rate), increasing reactant concentrations, increasing reactor holdup (in liquid-phase reactors), or increasing reactor pressure (in gas-phase reactors).

The choice of production-rate control cannot be made arbitrarily because of the implications for component balances examined in Step 7. Tyreus and Luyben (1993) and Luyben et al. (1996) discuss some workable choices and limitations of what handles can be used.

5. Control Product Quality and Handle Safety, Operational, and Environmental Constraints. Select the "best" valves to control each of the product-quality, safety, and environmental variables. We want tight control of these important quantities for economic and operational reasons. Hence we should select manipulated variables such that the dynamic relationships between the controlled and manipulated variables feature small time constants and dead times and large steady-state gains. The former gives small closed-loop time constants, and the latter prevents problems with the rangeability of the manipulated variable (control-valve saturation).

The magnitudes of various flow rates also come into consideration. For example, temperature (or bottoms product purity) in a distillation column is typically controlled by manipulating steam flow to the reboiler (column boilup), and base level is controlled with bottoms product flow rate. However, in columns with a large boilup ratio and small bottoms flow rate, these loops should be reversed because boilup has a larger effect on base level than bottoms flow. However, inverse-response problems in some columns may occur when base level is controlled by heat input. High reflux ratios at the top of a column require similar analysis in selecting reflux or distillate to control overhead product purity.

6. Control Inventories (Pressures and Levels) and Fix a Flow in Every Recycle Loop. Determine the valve to control each inventory variable. These variables include all liquid levels (except for surge volume in certain liquid recycle streams) and gas pressures. An inventory variable should typically be controlled with the manipulated variable that has the largest effect on it within that unit. Inventory may also be controlled with fresh reactant makeup streams. Liquid fresh feed streams may be added to a location where level reflects the amount of that component in the process. Gas fresh feed streams may be added to a location where pressure reflects the amount of that material in the process.

Proportional-only control should be used in nonreactive level loops for cascaded units in series. Even in reactor-level control, proportional control should be considered to help filter flow-rate disturbances to the downstream separation system. There is nothing necessarily sacred about holding reactor level constant.

In most processes a flow controller should be present in all liquid recycle loops. This is a simple and effective way to prevent potentially large changes in recycle flows that can occur if all flows in the recycle loop are controlled by levels. Two benefits result from this flow-control strategy. First, the plant's separation section is not subjected to large load disturbances. Second, consideration must be given to alternative fresh reactant makeup control strategies rather than flow control. In a dynamic sense, level controlling all flows in a recycle loop is a case of recycling disturbances and should be avoided.

Exceptions to this recycle flow-control strategy are possible if composition analyzers are available, either directly or indirectly (Luyben et al., 1996). For example, the fresh reactant feed flow can be used to control reactor feed or effluent composition. By controlling this composition, we prevent the buildup or depletion of this component. We then have a handle to avoid large changes in the recycle flow containing this reactant.

Gas recycle loops are normally set at maximum circulation rate, as limited by compressor capacity, to achieve maximum yields.

7. Check Component Balances. Ensure that the overall component balances for each chemical species can be satisfied either through reaction or exit streams by accounting for the component's composition or inventory at some point in the process. Light, intermediate, and heavy inert components must have an exit path from the system. Reactants must be consumed in the reaction section or leave as impurities in the product streams. Fresh reactant makeup feed streams can be manipulated to control reactor feed composition or a recycle stream composition (or to hold pressure or level as noted in the previous step). Purge streams can also be used to control the amount of high- or low-boiling impurities in a recycle stream.

Component balances can often be quite subtle. They depend upon the specific kinetics and reaction paths in the system. They often affect what variable can be used to set production rate or rate in the reactor. Downs (1992) provides an insightful discussion of some of these issues.

8. Control Individual Unit Operations. Establish the control loops necessary to operate each of the individual unit operations. Many effective control schemes have been established over the years for chemical units (Shinskey, 1988). For example, a tubular reactor usually requires control of inlet temperature. High-temperature endothermic reactions typically have a control system to adjust fuel flow rate to a furnace supplying energy to the reactor. Crystallizers require manipulation of refrigeration load to control temperature. Oxygen concentration in the stack gas from a furnace is controlled to prevent excess fuel usage. Liquid solvent feed flow to an absorber is controlled as some ratio to the gas feed.

9. Optimize Economics or Improve Dynamic Controllability. After satisfying all of the basic regulatory requirements, we usually have additional degrees of freedom involving control valves that have not been used and setpoints in some controllers that can be adjusted. These can be used either to optimize steady-state economic process performance (e.g., minimize energy, maximize selectivity) or to improve dynamic response.

For example, an exothermic chemical reactor can be cooled with both jacket cooling water and brine to a reflux condenser. A valve-position control strategy would allow fast, effective reactor temperature control while minimizing brine use.

Justification of Sequence

Although the order of the steps in the design procedure may initially seem arbitrary, the sequence comes from a consideration first of choices that have already been assigned due to equipment or business constraints and then the importance in a hierarchy of priorities. Steps 1 and 2 are straightforward in determining the objectives and available degrees of freedom.

Step 3 is next because the reactor is typically the heart of an industrial process and the methods for heat removal are intrinsically part of the reactor design. So it is usually not optional what degrees of freedom can be used for exothermic reactor control. When the heat generated in an exothermic reactor is used within the process by energy integration, we must ensure that the energy is dissipated and not recycled. Hence we examine process-to-process heat exchangers and heat-integrated unit operations to determine that we have sufficient degrees of freedom (bypass lines or trim heaters/coolers).

The choice of where production rate is set (Step 4) is often a pivotal decision, but it frequently is determined externally by a business objective. This removes another degree of freedom that cannot be used. If we are free to choose the handle for production rate, then Steps 5 through 7 are the priority order. However, at Step 7 we may determine that the choice will not work in light of other plantwide control considerations, in which case we would return to Step 4 and select something different. Determining the *best* choice at Step 4 can only be done by nonlinear dynamic simulation of disturbances with a complete control strategy.

Step 5 is done next because the control of product quality is closely tied to Step 1 and is a higher priority than the control of inventories. Hence it should be done early when we still have the widest choice of manipulators available. Its importance is based on the issue of variability, which we want to be as small as possible for on-aim product quality control. Variability in inventory control tends to be not as critical, which is the reason it is done in Step 6.

Only after the total process mass balance has been satisfied can we check on the individual component balances in Step 7. That then settles the plantwide issues. We now apply our knowledge of unit operation control in Step 8 to improve performance and remain consistent with the plantwide requirements. Finally, Step 9 addresses higher level concerns above the base regulatory control strategy.

Having presented our design procedure and our justification for its sequence, we now illustrate its application with three industrial process examples.

Vinyl Acetate Monomer Process

Background

Our first application comes from Luyben and Tyreus (1997), who present design details of an industrial process for the vapor-phase manufacture of vinyl acetate monomer. DuPont currently operates a vinyl acetate process at its plant in La-Porte, Texas. To protect any proprietary DuPont information, all of the physical property and kinetic data, process flow-sheet information, and modeling formulation come from sources in the open literature. The process flow sheet (Figure 1) is based upon the one given in Report 15B by SRI International (1994).

The process contains eleven basic unit operations: packed tubular reactor with steam drum, vaporizer, four heat exchangers, vapor–liquid separator, recycle compressor, absorber with cooler, CO_2 removal system, and azeotropic distillation column with decanter. Two irreversible exothermic reactions are considered to generate vinyl acetate, water, and carbon dioxide from reactants ethylene, oxygen, and acetic acid:

$$C_2H_4 + CH_3COOH + 1/2\,O_2 \rightarrow CH_2 = CHOCOCH_3 + H_2O$$

$$C_2H_4 + 3\,O_2 \rightarrow 2\,CO_2 + 2\,H_2O.$$

The reactor contains tubes packed with catalyst, and heat is removed by generating steam on the shell side of the tubes. The rate expressions for the two reactions are nonlinear functions of the partial pressures of oxygen, ethylene, acetic acid, and water. Products vinyl acetate and water leave from the decanter and go downstream to a refining system not

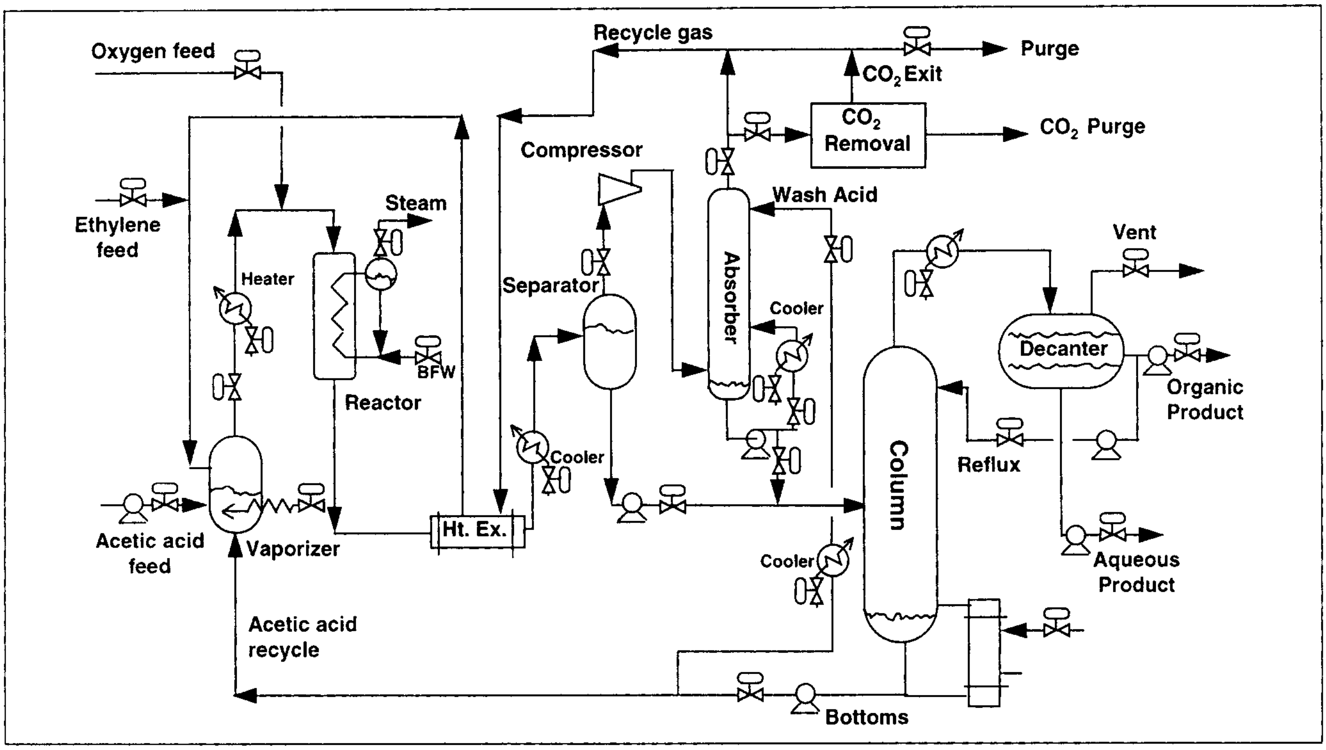

Figure 1. Vinyl acetate monomer process flowsheet.

considered in this study. Carbon dioxide is removed from the recycle gas stream in a system that is simplified to be a component separator with a certain efficiency that is a function of rate and concentration. A small purge stream also removes inert ethane that enters in the fresh ethylene stream.

There are three fresh reactant makeup feed streams (two gas and one liquid). The gas stream from the overhead of the absorber recycles unconverted ethylene and oxygen (plus carbon dioxide and ethane) back to the reactor via the vaporizer. The recycle gas is heated in a process-to-process heat exchanger with the reactor effluent gas stream. The liquid stream from the base of the azeotropic distillation column recycles unconverted acetic acid. Part of the bottoms stream is used in the absorber to recover water and vinyl acetate from the gas stream.

Plantwide control strategy

Step 1. For this process we must be able to set the production rate of vinyl acetate while minimizing yield losses to carbon dioxide. During the lifetime of the catalyst charge, catalyst activity decreases and the control system must operate under these different conditions. To maintain safe operating conditions, the oxygen concentration in the gas loop must remain outside the explosivity region for ethylene. The azeotropic distillation column must produce an overhead product with essentially no acetic acid and a bottoms product with no vinyl acetate.

Step 2. There are 26 control degrees of freedom in this process. They include: three feed valves for oxygen, ethylene, and acetic acid; vaporizer and heater steam valves; reactor steam drum liquid makeup and exit vapor valves; vaporizer overhead valve; two coolers and absorber cooling water valves; separator base and overhead valves; absorber overhead, base, wash acid, and liquid recirculation valves; gas valve to CO_2 removal system; gas purge valve; distillation column steam and cooling water valves; column base, reflux, and vent valves; and decanter organic and aqueous product valves.

Step 3. Energy management is critically important because of the highly exothermic reactions and potential for runaway or catalyst damage at high temperatures. By design, heat is removed from the reactor via transfer from the tubes to the shell, generating steam. Hence reactor temperature is controlled by steam temperature, which is set by controlling the pressure in the steam drum via the steam exit valve. This is a good example of how a degree of freedom needs to be used based upon the process design. If temperature measurements are available along the length of the tubes, then a reactor peak temperature or a profile can be controlled. Otherwise, the controlled variable is the reactor exit temperature.

The reactor effluent stream is cooled in a process-to-process heat exchanger with the gas recycle stream. A bypass line and control valve are necessary here only if we want to control one of the exchanger exit temperatures. If this exchanger is designed for only vapor flow, then the hot-side exit temperature must be controlled to a value above the dewpoint temperature by manipulating the bypass valve around the exchanger on the cold side (to avoid a control valve on the hot-stream side). The bypass line would have to be added if it were not included in the original design or the heat exchanger must be redesigned to handle two-phase flow.

Step 4. Ethylene and oxygen makeup feeds come from headers, and the acetic acid feed is drawn from a supply tank.

The vinyl acetate and water products go to downstream units. As a result, there are no design constraints that require production rate to be set either on supply or demand. Therefore, we look at reactor conditions to determine how to change production rate. Because the reactor feed contains both excess ethylene and acetic acid, manipulating the partial pressure of either component would not be effective. The partial pressure of oxygen is constrained by the safety limit, and once this is reached no further adjustments could be made. Pressure is limited by the process equipment design maximum. Hence the most direct handle for setting production rate is by changing the reactor exit or peak temperature.

Alternatively we could use the fresh oxygen feed flow to set production rate since it is the limiting component. However, there are two issues with this choice. Since oxygen is not completely consumed, we must worry about its accumulation in the system (component balance), which is constrained here by the safety limit. If oxygen were completely consumed, we must still worry about the reactor inlet oxygen composition because of the safety constraint. In either case with oxygen feed rate, we would have to control oxygen composition with some other variable to change production rate safely.

Step 5. The azeotropic distillation column does not produce the final salable vinyl acetate product. Its primary role is to recover and recycle unreacted acetic acid and to remove from the process all of the vinyl acetate and water produced. So we want little acetic acid in the overhead because this represents a yield loss. Also, the bottoms stream should contain no vinyl acetate since it polymerizes and fouls the heat-exchange equipment at the elevated temperatures of the column base and the vaporizer. Hence, we have two control objectives: base vinyl acetate and top acetic acid compositions. And we have two manipulators: steam and organic reflux flows. The temperature profile has a sharp break representing the change in vinyl acetate composition near the bottom of the column. Column steam (boilup) is the appropriate choice for temperature control because of its fast response compared with reflux. Then the overhead acetic acid composition must be controlled with reflux.

The overriding safety constraint in this process involves oxygen concentration in the gas loop, which must remain below 8 mol % to remain outside the explosivity envelope for ethylene mixtures at process conditions. The most direct manipulated variable to control oxygen composition at the reactor feed is the fresh oxygen feed flow.

If in Step 4 we chose to set production rate by flow controlling the fresh oxygen feed, then we would need an alternative manipulator to control oxygen composition. The only choice would be to use reactor temperature. We now can consider which choice—safety or production rate—is better. Temperature is *not* the most direct handle to control oxygen composition since its effectiveness hinges on incomplete oxygen conversion in the reactor and oxygen recycle. Because of the safety implications, we would choose to use fresh oxygen flow to control reactor inlet oxygen composition, which means production rate is set by reactor temperature.

Step 6. Two pressures must be controlled: in the column and in the gas loop. The most direct handle to control column pressure is by manipulating the vent stream from the decanter. We have three choices to control gas-loop pressure: purge flow, flow to the CO_2 removal system, and the fresh ethylene feed flow, since fresh oxygen flow has been previously selected. Both the purge flow and the flow to the CO_2 removal system are small relative to the gas recycle flow rate. Any changes in either one would not have a large effect on gas-loop pressure. Since ethylene composes a substantial part of the gas recycle stream, pressure is a good indication of the ethylene inventory. So we choose the fresh ethylene feed flow to control gas recycle loop pressure.

Note that there are three vessels (vaporizer, separator, absorber) within the gas loop where apparently pressure can also be controlled. However, these pressures actually cannot be selected arbitrarily once the compressor capacity and gas-loop pressure drop are established. In fact, to minimize pressure drop in the gas loop we would open completely or remove the overhead vapor control valves on these units, saving both compressor and valve costs.

Seven liquid levels are in the process: vaporizer, reactor steam drum, separator, absorber, column base, and two decanter layers. Control of the decanter levels is straightforward. The organic product flow controls the organic phase inventory while the aqueous product flow controls the aqueous phase inventory. Reactor steam drum level is maintained with boiler feed-water makeup flow.

The most direct way to control the remaining levels would be with the exit valves from the vessels. However, if we do this, we see that all of the flows around the liquid recycle loop would be set based upon levels, which would lead to undesirable propagation of disturbances. Instead we should control a flow somewhere in this loop. Acetic acid is the main component in the liquid recycle loop. Recycle and fresh acetic acid feed determine the component's composition in the reactor feed. A reasonable choice at this point is to control the total acetic acid feed-stream flow into the vaporizer. This means that we can use the fresh acetic acid feed stream to control column base level, since this is an indication of the acetic-acid inventory in the process. Vaporizer level is then controlled with the vaporizer steam flow, and separator and absorber levels can be controlled with the liquid exit valves from the units.

Step 7. Ethane is an inert component that enters with the ethylene feed. It can be removed from the process only via the gas purge stream, so purge flow is used to control ethane composition. Carbon dioxide is an unwanted byproduct that leaves in the CO_2 removal system. As long as the amount of carbon dioxide removed is proportional in some way to the CO_2 removal system feed, we can use this valve to control carbon dioxide composition. Oxygen inventory is accounted for by composition control with fresh oxygen feed. Inventory of ethylene can be controlled to maintain gas-loop pressure since ethylene composes the bulk of the gas recycle.

Acetic-acid inventory is regulated by using the fresh acetic-acid feed to control the base level in the distillation column. The temperature control loop in the distillation column achieves vinyl acetate composition control. Water, however, is an intermediate component with a boiling point between vinyl acetate and acetic acid. The inventory of water in the process will not be automatically accounted for by controlling those other two components. Instead, we must use reflux flow to control the water composition in the bottoms stream. Otherwise, there is no regulation of water to ensure it is removed from the process. By using column reflux to control base water, we are forced to give up on using reflux to control acetic acid recovery (Step 5). To ensure that the acetic acid composition in the decanter is acceptable, the column must have a sufficient number of stages. This means we may have to revisit its design.

Step 8. Several control valves now remain unassigned. Steam flow to the trim heater controls reactor inlet temperature. Cooling water flow to the trim cooler is used to control the exit process temperature and provide the required condensation in the reactor effluent stream. Liquid recirculation in the absorber is flow controlled to achieve product recovery, while the cooling water flow to the absorber cooler controls the recirculating liquid temperature. Acetic-acid flow to the top of the absorber is flow controlled to meet recovery specifications on the overhead gas stream. Cooling water flow to the cooler on this acetic-acid feed to the absorber is regulated to control the stream temperature. Cooling water flow in the column condenser controls decanter temperature.

Step 9. We have now established the basic regulatory plantwide control strategy (Figure 2). Based upon the heuristic established by Fisher et al. (1988) that recycle gas flows should be maximized to improve reactor yield, we open or remove the separator, vaporizer, and absorber overhead valves and run the compressor full out. To minimize the decanter temperature for improved organic recovery, the column condenser cooling water is set at maximum flow. Optimization of several controller setpoints can be done based upon economics. We must balance the trade-offs in maximizing vinyl acetate production and recovery with minimizing carbon dioxide production and energy consumption. This involves considering the reactor temperature setpoint, reactor feed temperature setpoint, composition of carbon dioxide and ethane in the gas recycle loop, oxygen composition setpoint (up to the maximum constraint), total flow rate of acetic acid to the vaporizer, and water composition in the recycle acetic acid. Additionally, economic evaluations must account for the effects of catalyst deactivation with time over the duration of the charge. To achieve the same vinyl acetate production rate, we must increase reactor temperature. However, this produces more carbon dioxide representing yield loss.

We have constructed a rigorous nonlinear first-principles dynamic model of this process with TMODS, DuPont's in-house interactive dynamic simulator. We have used the model to demonstrate that we have developed a workable control strategy.

Eastman Process

Background

Our next example is the Eastman process published by Downs and Vogel (1993) as a benchmark industrial process control problem. Here we illustrate the effect of different objectives on the resulting control structure (in this case, what determines production rate). First, production rate is set by the flow rate of the product stream leaving the process. Second, production rate is set by the flow rate of a fresh feed stream entering the process. These two criteria lead to completely different strategies, both of which are effective but accomplish different objectives under normal conditions. Overrides are used in both schemes, so under abnormal conditions both strategies adopt the appropriate structure to handle the unusual conditions.

The process (Figure 3) contains five unit operations: reactor, condenser, vapor–liquid separator, recycle compressor,

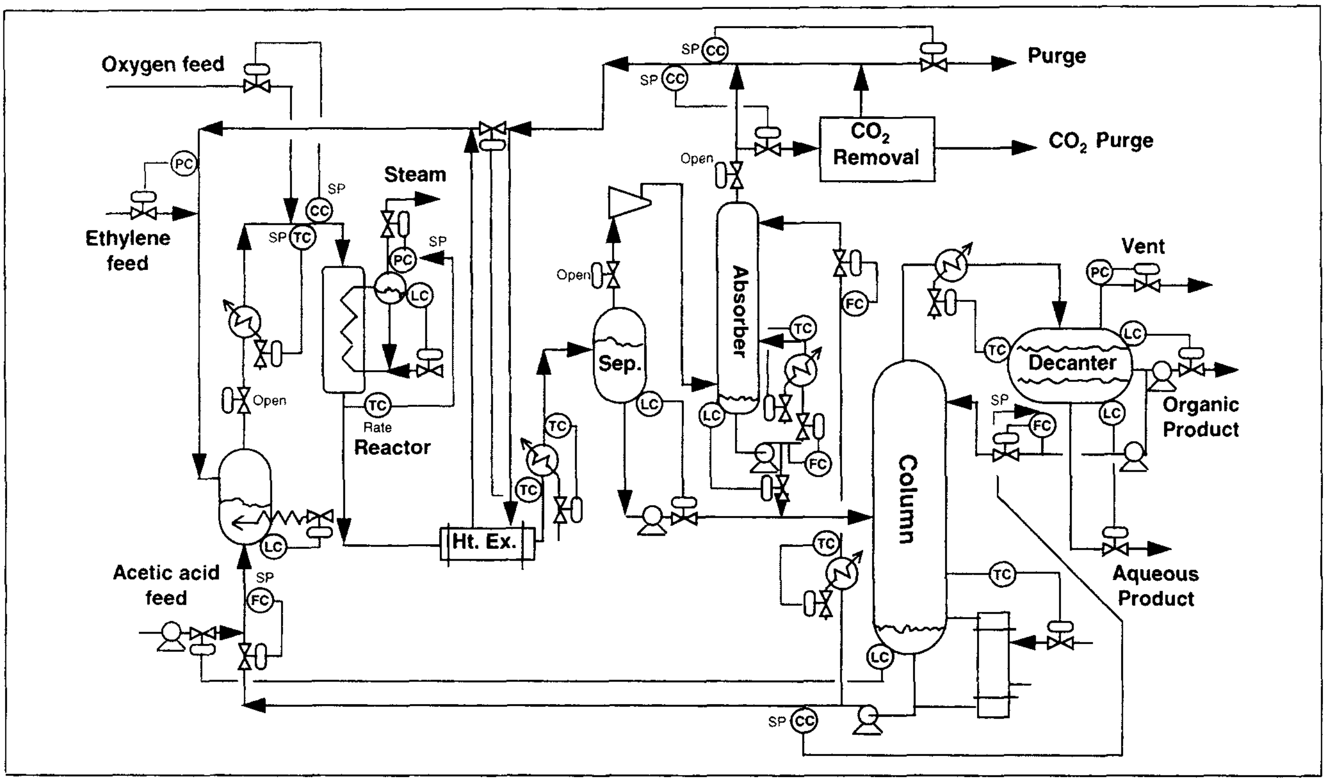

Figure 2. Plantwide control strategy for vinyl acetate monomer process.

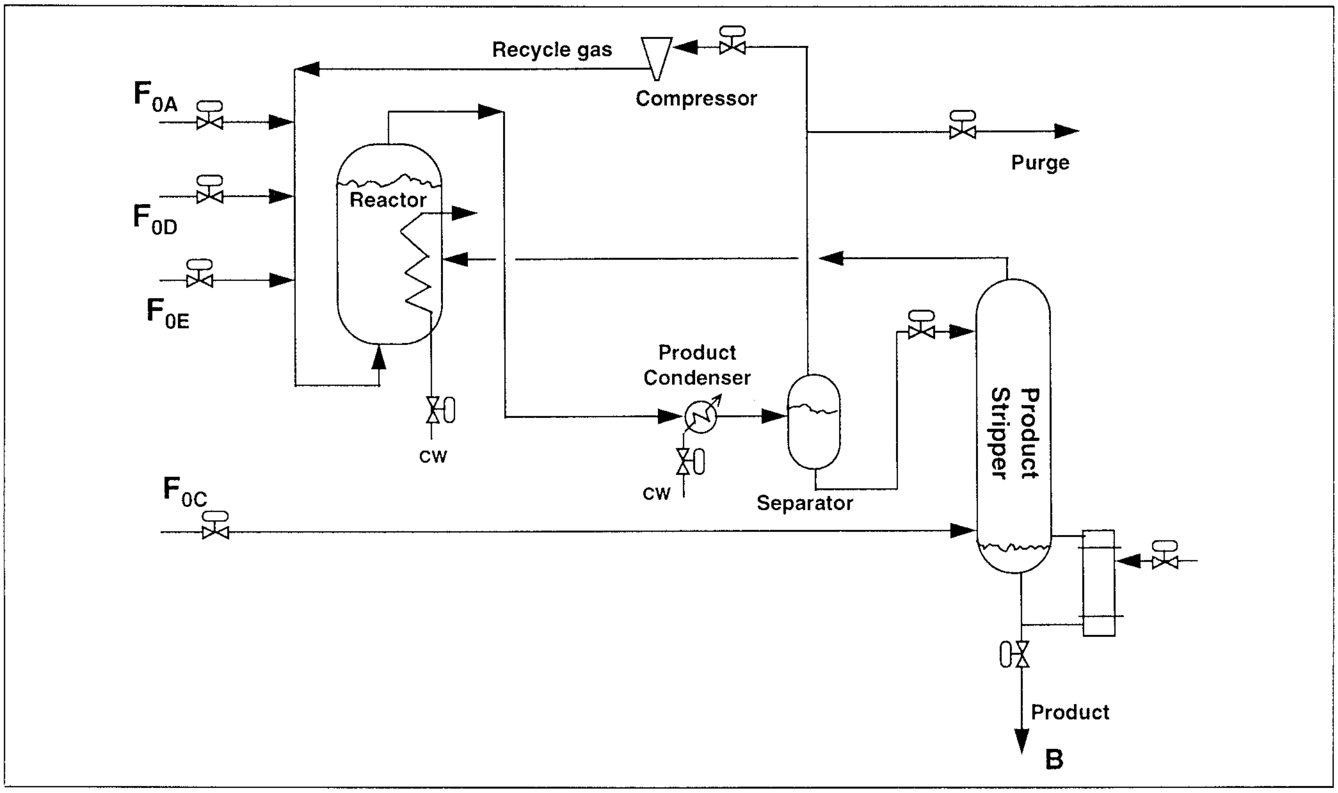

Figure 3. Eastman process flow sheet.

and product stripper. A small purge stream removes an inert component. The product stream is the bottoms from the stripper. The reactor, which is open-loop unstable, contains both liquid and vapor phases, but no liquid stream leaves the reactor. Four irreversible exothermic reactions are considered to generate products G and H and byproduct F. The reactions are approximately first order with respect to reactant concentrations. Temperature, pressure, and liquid level in the reactor all interact and their behavior is nonlinear. There are four fresh makeup gas feed streams. Fortunately, composition analyzers are available.

Criterion No. 1

Step 1. In this case we assume that the process must meet an on-demand production rate requirement because the flow rate of the product stream is set by the downstream process.

Step 2. This process has 12 degrees of freedom. One of these is agitation rate, which we simply hold constant. A second is the valve that controls the flow rate of the gas recycle. As described in the vinyl acetate process, we open this to maximize gas recycle flow and improve reactor yield. This leaves 10 degrees of freedom: four fresh feeds F_{0A}, F_{0D}, F_{0E}, and F_{0C}; purge valve; separator base valve; stripper base and steam valves; reactor cooling-water valve; and condenser cooling-water valve.

Step 3. The open-loop instability of the reactor acts somewhat like a constraint, since closed-loop control of reactor temperature is required. By design, the exothermic reactor heat is removed by cooling water in the reactor and product condenser. We choose to control reactor temperature with reactor cooling-water flow because of its direct effect. There are no process-to-process heat exchangers and no heat integration in this process. Disturbances can then be rejected to the plant utility system by cooling water or steam.

Step 4. Because of the objective to achieve on-demand production rate, the product stream leaving the stripper base is flow controlled via the bottoms control valve. This is a good example of how a degree of freedom must be used to satisfy a design or business constraint.

Step 5. The manipulator of choice to control product quality is stripper steam flow because of its fast response. Stripper temperature can be used to infer product composition. Another manipulated variable that directly affects stripper bottoms purity is the flow rate of feed F_{0C}. However, this fresh feed makeup stream affects the component balances of A and C in the system, while steam does not, and we would have recognized this at Step 7. Therefore, we choose reboiler steam to control product purity.

High reactor temperature (175°C) is one safety constraint. Reactor-cooling water flow has previously been selected to control reactor temperature. The other safety constraint is high pressure, which must not exceed the shutdown limit of 3,000 kPa. The gas fresh feed streams, cooling-water streams, reboiler steam flow, and purge directly affect pressure. Any of these could potentially be used to control reactor pressure. Of course, reaction rate also affects pressure, so a variable that changes reaction rate could potentially be used to control pressure indirectly. Reactor-cooling water flow and steam have already been selected. Condenser cooling rate is smaller than reactor cooling rate, so it may not be very effective in controlling pressure. This leaves one of the gas flows. Fresh feed F_{0C} is the largest of these by far, so we choose it to control pressure.

Step 6. Three liquid levels need to be controlled: reactor, separator, and stripper base. We must use the Buckley strategy of *level control in the reverse direction to flow* since the stripper base product flow is fixed by production rate. Therefore liquid flow from the separator must be used to control the stripper base level. To control level in the separator, we select the cooling water flow to the condenser.

Now we must decide how to control the liquid level in the reactor. This liquid consists of mostly the heavy products, components G and H. The more fresh reactant components D and E are fed into the process, the more products will be produced. So we select the two fresh feed flow rates F_{0D} and F_{0E} to control reactor liquid level. We ratio one to the other, depending upon the desired split between components G and H in the final product. Simple flow ratios should be accurate enough to maintain the desired product distribution without any feedback of product compositions. So on-line analyzers on the product stream should not be required.

Step 7. A light inert component B enters in one of the feed streams. It can be removed from the process only by the purge stream, so purge flow rate is used to control its composition in the purge gas stream. Stripper temperature control keeps the volatile gas reactants within the gas recycle loop. Components D and E are accounted for by reactor level. The component balance for C is maintained by pressure, assuming we can control the composition of the other major component A in the gas loop. There must be some feedback mechanism to guarantee that precisely the correct number of moles of this component is fed into the system to react with the number of moles of component C. The only manipulator available to satisfy the component balance for A is the fresh feed stream flow F_{0A}. So we choose this to control the composition of A in the purge gas stream y_A. The composition of either the purge gas or the reactor feed could be used, but both are not necessary.

Step 8. Control of each individual unit operation has been established and all of the control valves have been assigned.

Step 9. Of the original 10 control valves, we have used one for production rate, one for product quality, one for pressure control, one for reactor temperature control, three for liquid levels, and two for compositions. An additional one was used to set the G/H ratio. However, we have an additional degree of freedom to be specified, which is the setpoint of the reactor temperature controller. The best way to manage the reactor temperature setpoint is not immediately obvious. It might be used in conjunction with the production-rate controller, that is, higher temperatures may be needed to increase throughputs. It might be adjusted to maximize yields and suppress undesired byproducts.

However, after making simulation runs with several of the disturbances suggested in the original article, it became apparent that the temperature in the separator was changing quite substantially and adversely affecting the stripper. Low separator temperature drops too many light components into the stripper, and the reboiler steam has trouble maintaining product quality. Therefore a separator temperature controller was added, whose output signal is the setpoint of the reactor temperature controller. The final basic regulatory control structure (Figure 4) is simple, effective, and easily understood by operating personnel.

This basic regulatory control structure was able to hold the process at the desired operating point for most of the disturbances. However, when manipulated variables hit constraints, it was unable to prevent a unit shutdown. The disturbance that shuts off the feed flow rate F_{0A} is probably the most drastic. The resulting imbalance in the stoichiometric amounts of components A and C drives the concentration y_A down quite rapidly. The reaction rate slows up, reactor temperature drops, and the process shuts down on high pressure. Since one degree of freedom has been removed by this disturbance, the control structure must be modified with overrides to handle the component balances.

The F_{0C} stream contains more C than A, so the excess C must be removed from the system. The only place available is the purge stream. Therefore, a low F_{0A} flow override controller is used to open the purge valve. The other action that must be taken is to prevent the concentration of component A in the system from dropping down too low and reducing reaction rates. This is achieved by using a low y_A concentration override controller to pinch the fresh feed flow rate F_{0D} to slow up the rate of consumption of A. Of course, F_{0E} is also reduced through the ratio.

Now the liquid level loops must also be modified since we no longer can specify production rate, and reactor level control cannot use F_{0D}. This is easily accomplished by using low-level override controllers on each of the three levels. Low stripper level pinches base product flow rate. Low separator level pinches separator liquid flow rate. Low reactor level pinches the condenser-cooling water flow rate. In an override situation the level control structure has been reversed from the basic structure, and now levels are held *in the direction of flow*.

All controllers (basic regulatory and override) were made proportional-only because there are no tight specifications on any variables. This process acts as an integrator in terms of pressure, liquid levels, and chemical components with little self-regulation, and this makes it difficult to tune controllers with integral action.

Criterion No. 2

Step 1. Instead of on-demand production required by a downstream unit, here we assume that the process must take whatever amount of fresh feed stream F_{0C} the upstream unit supplies because of insufficient holdup in the upstream process.

Step 2. Again, there are 10 degrees of freedom.

Step 3. This is constructed as before.

Step 4. Production rate is set by specifying F_{0C}.

Step 5. Reboiler steam controls product purity. Condenser cooling rate is used to control pressure.

Step 6. We use the Buckley strategy of *level control in the direction to flow* to regulate two liquid levels. The stripper base product flow rate is used to control stripper level. The flow rate of liquid from the separator is used to control the level in the separator. The liquid level in the reactor is controlled by the two fresh feed flow rates F_{0D} and F_{0E}. We ratio one to the other, depending upon the desired split between components G and H in the final product.

Step 7. Purge flow controls the composition of component B in the purge gas stream. The fresh feed stream F_{0A} is used to control the composition of A in the purge gas stream.

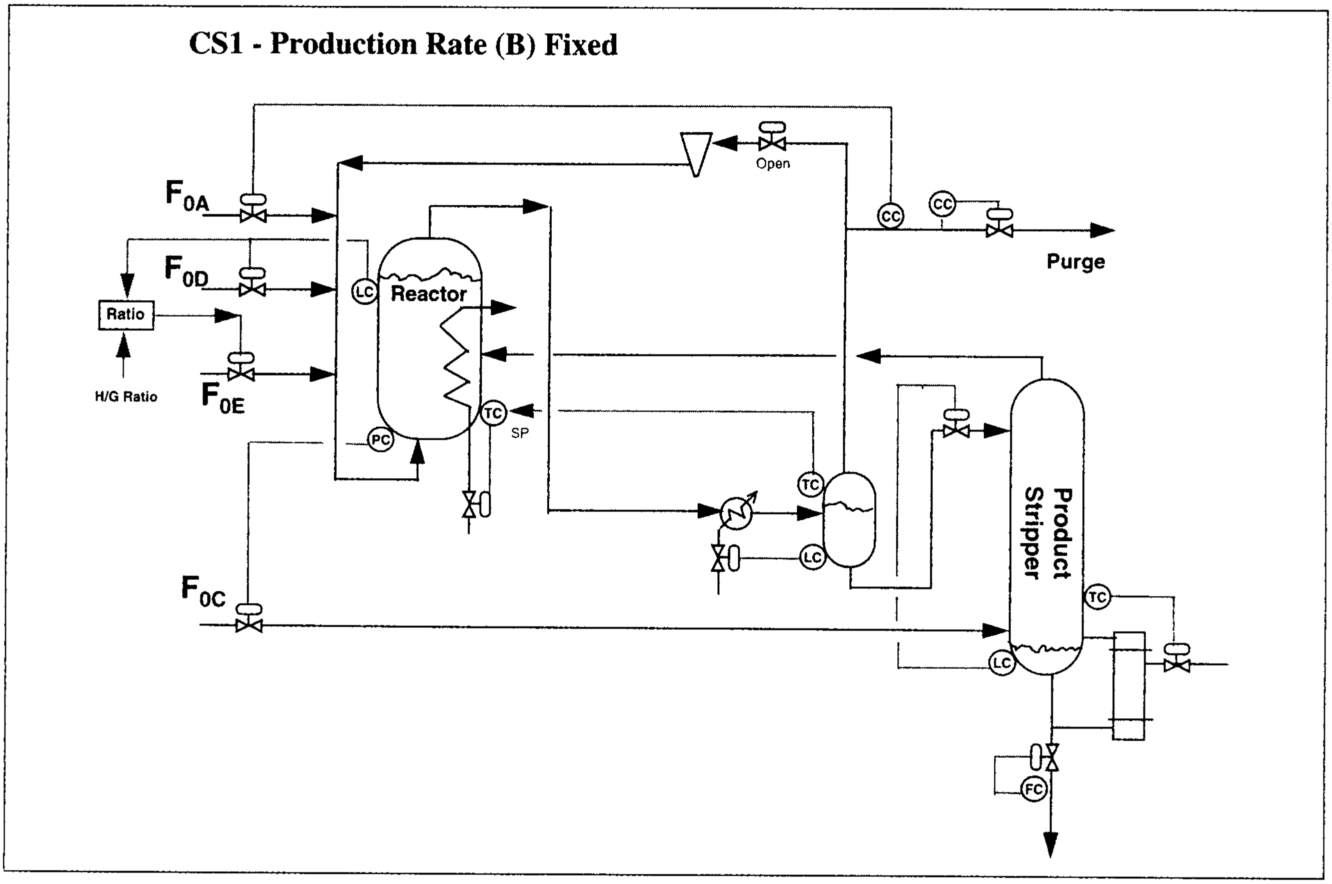

Figure 4. Plantwide control strategy for Eastman process—criterion No. 1.

Step 8. All control valves have been assigned.

Step 9. Separator temperature is controlled by resetting the setpoint of the reactor temperature controller. The final basic regulatory control structure is shown in Figure 5.

As before, a low F_{0A} flow override controller is used here to open the purge valve. A low y_A concentration override controller pinches the fresh feed flow rate F_{0D} to slow up the rate of consumption of A.

Both control systems were tested using the Fortran simulation provided by Downs and Vogel and were shown to provide effective control of the Eastman process. Although the two strategies handle disturbances differently, they both work without showing a clear advantage of one compared with the other.

HDA Process

Background

The design of the process for the hydrodealkylation of toluene (HDA) is extensively covered by Douglas (1988), from which we directly use the process flow sheet (Figure 6). All equipment design, physical property, and kinetic data are given by Douglas (1988) or are available from sources in the open literature. We have determined certain design and control variables (e.g., column feed locations, temperature control trays, overhead receiver and column base liquid holdups) that Douglas does not specify.

The process contains nine basic unit operations: reactor, furnace, vapor–liquid separator, recycle compressor, two heat exchangers, and three distillation columns. Two vapor-phase reactions are considered to generate benzene, methane, and diphenyl from reactants toluene and hydrogen. The reactor is adiabatic and must be run with an excess of hydrogen to prevent coking. The reactor effluent is quenched with liquid from the separator to prevent fouling in the process-to-process heat exchanger. The rate expressions are functions of the partial pressures of toluene, hydrogen, benzene, and diphenyl. Byproduct diphenyl is produced in an equilibrium reaction. Product benzene leaves as the distillate from the product column. Byproduct methane is removed in the purge stream and in the vapor overhead stream from the stabilizer column. Byproduct diphenyl exits in the bottoms stream from the recycle column.

There are two fresh reactant makeup feed streams (one gas and one liquid). The gas stream from the overhead of the vapor–liquid separator recycles unconverted hydrogen plus methane back to the reactor. The reactor feed is heated first in a process-to-process heat exchanger with the reactor effluent stream. The reactor feed is then brought up to temperature in the furnace, which is heated by combustion of fuel. The liquid distillate stream from the recycle column returns unconverted toluene to the reactor.

Plantwide control strategy

Step 1. For this process we must produce essentially pure benzene while minimizing yield losses of hydrogen and diphenyl. The reactor feed ratio of hydrogen to aromatics

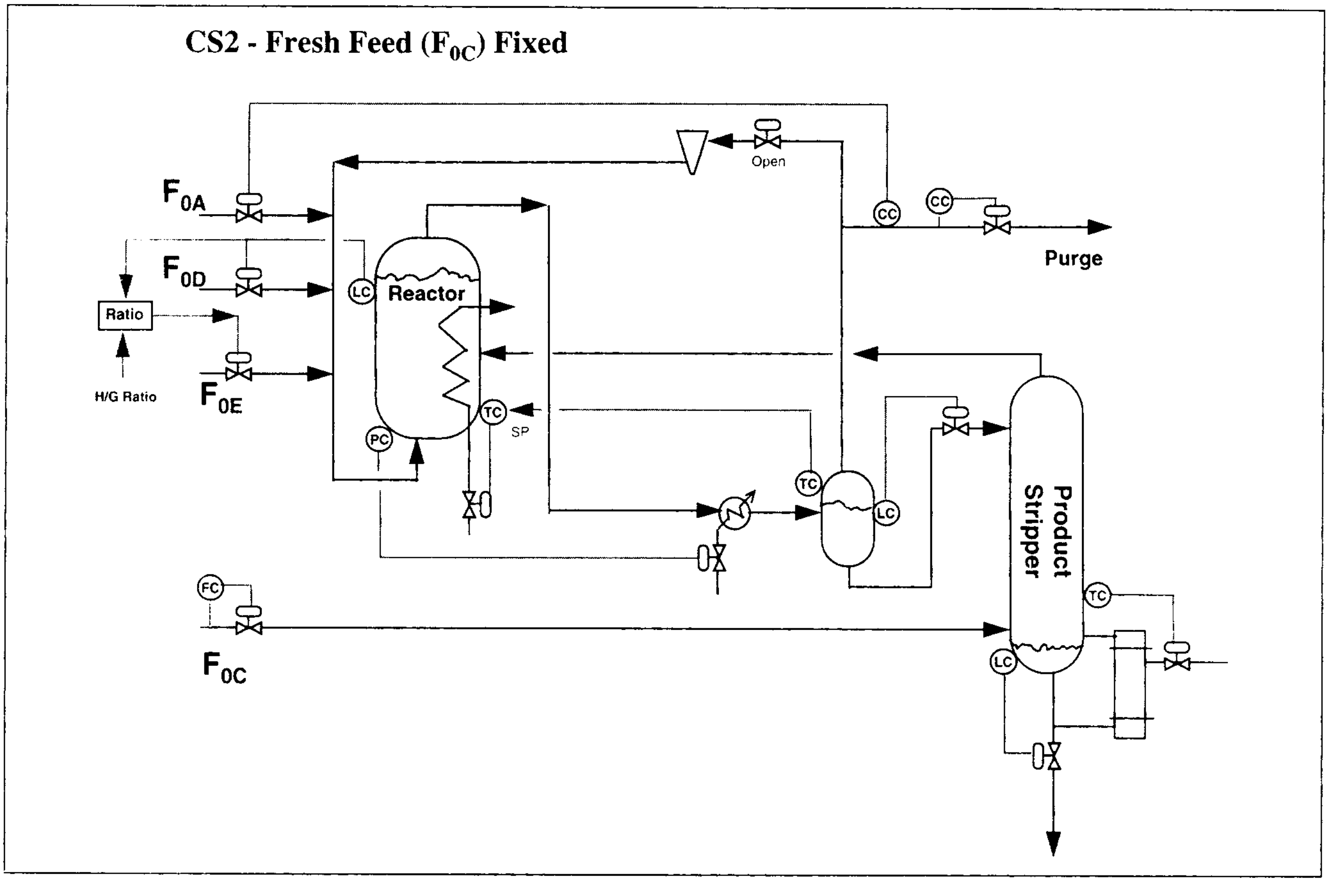

Figure 5. Plantwide control strategy for Eastman process—criterion No. 2.

must be greater than 5:1. The reactor effluent gas must be quenched to 620°C.

Step 2. There are 23 control degrees of freedom. They include: two fresh feed valves for hydrogen and toluene; purge valve; separator base and overhead valves; cooler cooling water valve; liquid quench valve; furnace fuel valve; stabilizer column steam, bottoms, reflux, cooling water, and vapor product valves; product column steam, bottoms, reflux, distillate, and cooling water valves; and recycle column steam, bottoms, reflux, distillate, and cooling water valves.

Step 3. The reactor operates adiabatically, so for a given reactor design the exit temperature depends upon the heat capacities of the reactor gases, reactor inlet temperature, and reactor conversion. Heat from the adiabatic reactor is carried in the effluent stream and is not removed from the process until it is dissipated to utility in the separator cooler. If we make the process-to-process heat exchanger too large, this means we may recycle the reactor heat and not dissipate it unless we have a bypass line around the exchanger. This is particularly true if the furnace is used only for trim control because it may not be able to provide sufficient control for a significant disturbance. If the process-to-process heat exchanger is not too large, this means that the furnace is operating under normal conditions. Then we are allowing more heat to go to utility at the separator cooler than we are generating in the reactor, and a bypass is not needed. For the HDA reactor with the quench loop, the furnace must be in operation at all times.

To ensure exothermic heat removal from the process, we are constrained by the process design to assign two control loops. We must control reactor inlet temperature with the furnace and control reactor exit temperature with the quench flow. Only by adjusting fuel to the furnace do we allow the reactor heat to be dissipated to the cooler. And only when the quench loop works do we guarantee that the furnace is operational. Because of this design we do not need a bypass line around the process-to-process heat exchanger.

Step 4. Hydrogen feed comes from a header, and toluene feed is drawn from a supply tank. The benzene, methane, and diphenyl products go to headers or tanks. Hence we are not constrained to set production either by supply or demand. There are conceptually four variables that alter reaction rate. Two are reactor pressure and total reactor feed flow at constant composition. Neither of these is viable because we want to run at maximum system pressure and compressor capacity for yield purposes. The other two are reactor inlet temperature and reactant composition (since hydrogen is in excess, toluene is the limiting component). This gives us two viable options: change reactor inlet temperature or inlet toluene composition. Another effect of a change in production rate is that there has to be a net change of both the hydrogen and toluene fresh feed rates into the process.

If we select temperature, we would like the reactor flow and composition to be nearly constant and we are constrained by the upper reactor temperature limit of 705°C. If we select toluene composition, we can control it either di-

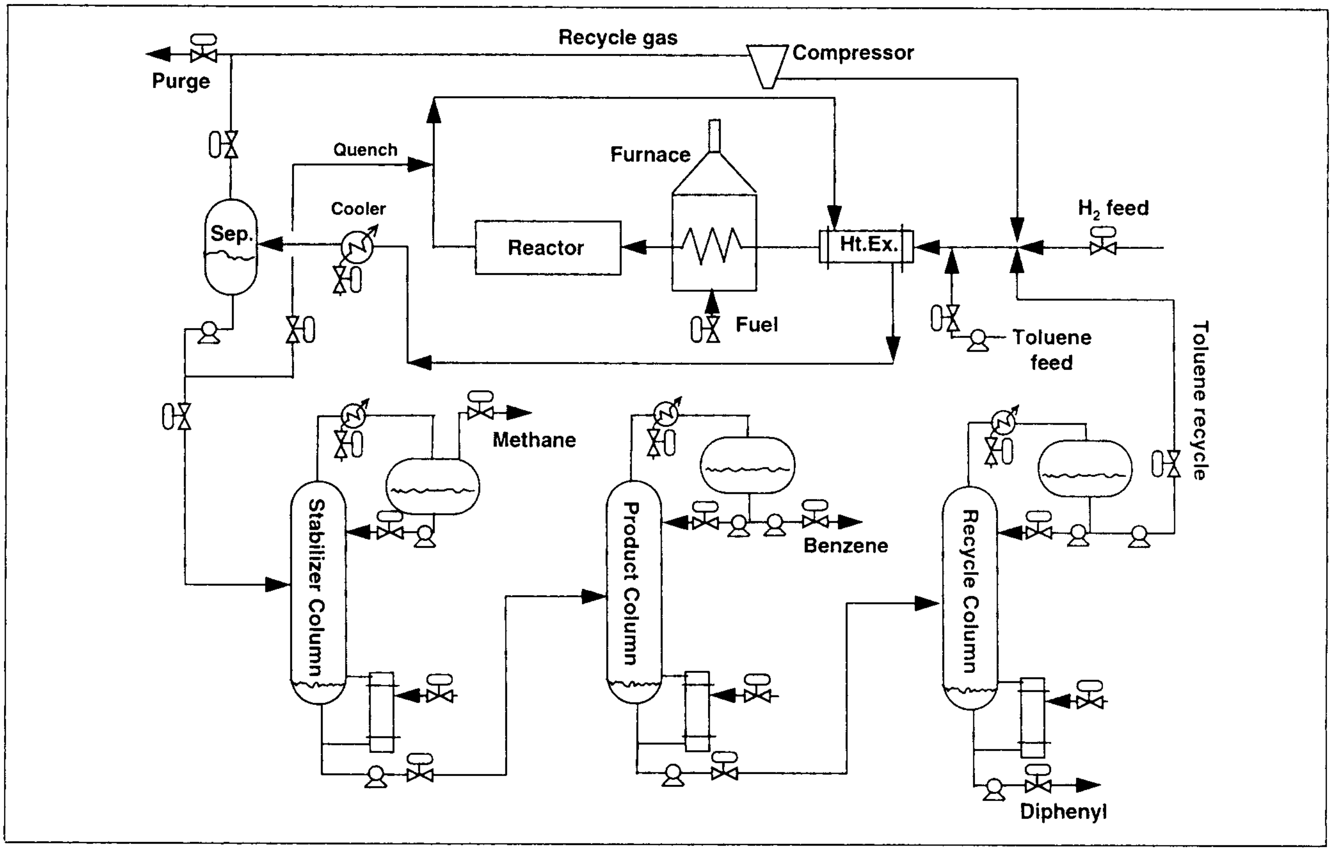

Figure 6. HDA process flow sheet.

rectly or indirectly. If directly, a reactor-feed composition analyzer is needed and is used to adjust either the fresh toluene feed rate or the total reactor toluene feed rate. If indirectly, the separation section is used as an analyzer for toluene. This allows us to control the total flow of toluene to the reactor (recycle plus fresh). Fresh toluene feed flow is used to control toluene inventory reflected in the recycle column overhead receiver level as an indication of the need for reactant makeup. Controlling the total toluene flow sets the reactor composition indirectly and is advantageous because it is less complicated and does not require an on-line analyzer.

Step 5. The distillate stream from the product column is salable benzene. Benzene quality can be affected primarily by two components, methane and toluene. Any methane that leaves in the bottoms of the stabilizer column contaminates the benzene product. The easy separation in the stabilizer column allows us to prevent this by using a temperature to set column steam rate (boilup). Toluene in the overhead of the product column also affects benzene quality. In this column the separation between benzene and toluene is also fairly easy. As a result, we can control product column boilup using a tray temperature. To achieve on-aim product quality control, we most likely would use an on-line overhead composition analyzer to adjust the setpoint of this temperature controller.

Important operational constraints have already been addressed with the quench and reactor inlet temperature control loops set in Step 3. If reactor inlet temperature is used to set production, we would include an override on this controller to keep the reactor exit temperature below the maximum allowable value. If toluene flow (fresh or recycle) is used to set rate, then we would need an override on this to maintain the minimum 5:1 ratio of hydrogen to aromatics in the reactor feed.

Step 6. Four pressures must be controlled: in the three distillation columns and in the gas loop. In the stabilizer column, vapor product flow is the most direct manipulator to control pressure. In the product and recycle columns, pressure control can be achieved by manipulating cooling water flow to regulate overhead condensation rate. As in the vinyl acetate process, we open the separator overhead valve and run the compressor at maximum gas recycle rate to improve yield. We then have two choices to control gas loop pressure: purge or fresh hydrogen feed flow. Since pressure indicates hydrogen inventory in the gas recycle loop, we choose the hydrogen feed to control gas loop pressure.

Seven liquid levels are in the process: separator and two (base and overhead receiver) in each column. The most direct way to control separator level is with the liquid flow to the stabilizer column. Then the stabilizer column overhead receiver level is controlled with cooling water flow and base level is controlled with bottoms flow. In the product column, distillate flow controls overhead receiver level and bottoms flow controls base level.

After these choices, we must decide about level control in the recycle column. Contrary to the other columns, here the boilup ratio is large since the bottoms diphenyl flow is quite small compared with the toluene recycle rate. For this case,

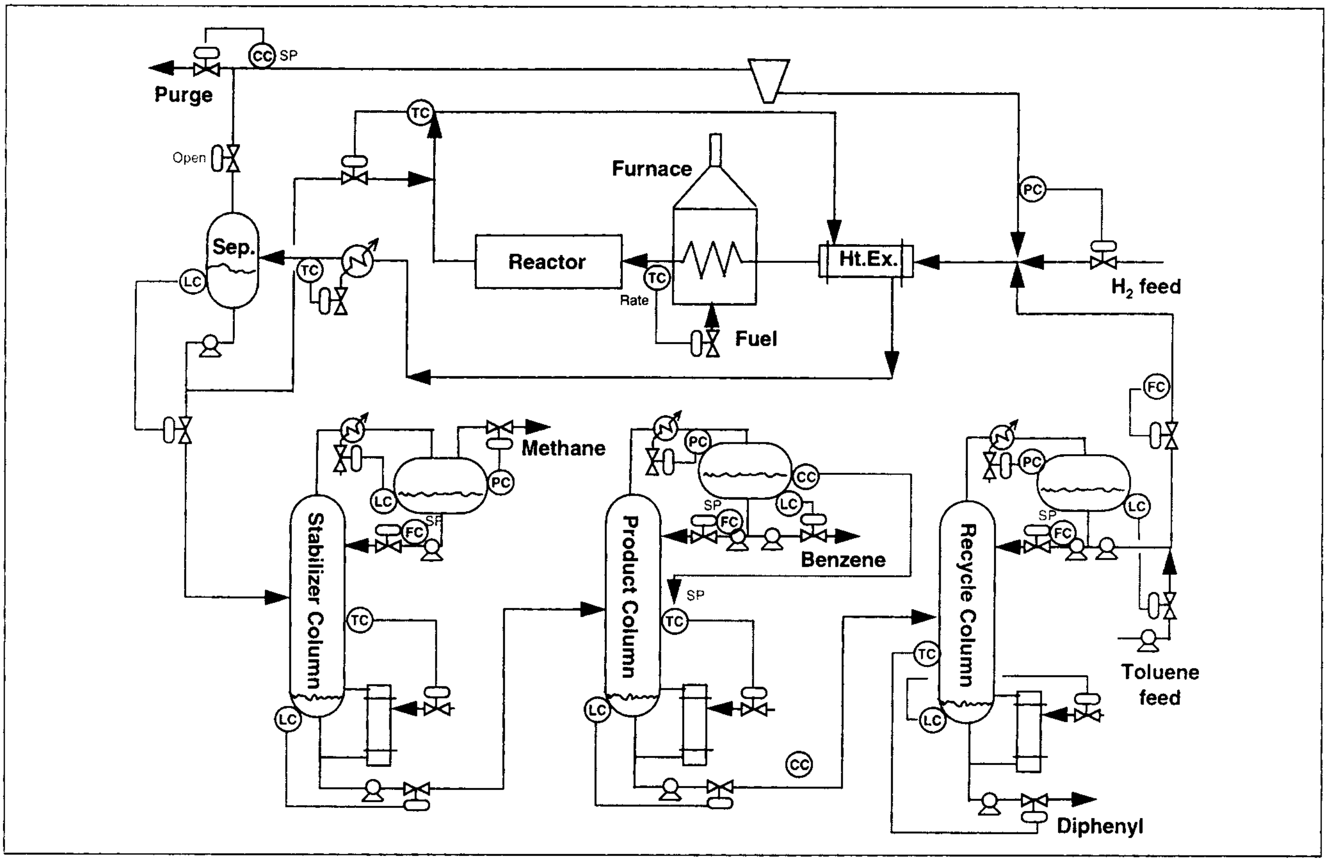

Figure 7. Plantwide control strategy for HDA process.

we choose to control base level with the steam flow because it has a much larger effect. If we use distillate flow from the recycle column to control overhead receiver level, then we see that all of the flows around the liquid recycle loop are set based upon level. This strategy would work only in the case where we had chosen to control reactor feed toluene composition directly with the fresh toluene feed flow as the production rate handle. This direct composition control prevents the buildup or depletion of toluene and avoids large changes in toluene recycle flow. For all other choices of production rate control, we want to control the total toluene flow to the reactor. This leads us to consider overhead receiver level control with the fresh makeup toluene feed, since this level represents the toluene inventory in the process. Such a scheme limits large flow-rate changes to the refining section and automatically ensures the component balance for toluene.

Step 7. Methane is purged from the gas recycle loop to prevent it from accumulating and its composition can be controlled with purge flow. Diphenyl is removed in the bottoms stream from the recycle column, where steam flow controls base level. Here we control composition (or temperature) with the bottoms flow. The inventory of benzene is accounted for by temperature and overhead receiver level control in the product column. Toluene inventory is accounted for by level control in the recycle column overhead receiver. Purge flow and stabilizer column temperature and pressure control establish the methane component balance. Gas loop pressure control accounts for hydrogen inventory.

Step 8. We now can assign control loops within individual units. Cooling water flow to the cooler controls process temperature to the separator. Reflux to the stabilizer, product, and recycle columns can be flow controlled because there is no requirement at the unit operations level to do anything beyond this. Step 9 discusses how the reflux flows should be set.

Step 9. The basic regulatory strategy has now been established (Figure 7). We have some freedom to select several controller setpoints to optimize economics and plant performance. If reactor inlet temperature sets production rate, the setpoint of the total toluene flow controller can be selected to optimize reactor yield. However, there is an upper limit on this toluene flow to maintain at least a 5:1 hydrogen to aromatic ratio in the reactor feed since hydrogen recycle rate is maximized. The setpoint for the methane composition controller in the gas recycle loop must balance the trade-off between yield loss and reactor performance. Reflux flows to the stabilizer, product, and recycle columns must be determined based upon column energy requirements and potential yield losses of benzene (in the overhead of the stabilizer and recycle columns) and toluene (in the base of the recycle column). Since the separations are easy, in this system economics indicate that the reflux flows would probably be constant.

We have constructed a rigorous nonlinear dynamic model of the HDA process with TMODS. We have used the model to demonstrate that we have developed a workable control strategy for various disturbances including changes in production rate. Other control strategies have been previously proposed for this process that are significantly different than ours (Stephanopoulos, 1984; Ponton and Laing, 1993). However, to our knowledge no simulation results or verifications have been reported.

Conclusions

The problem of plantwide control is to develop a control strategy for an entire *complex* and *integrated* process that satisfies the plant's design objectives. The presence of recycle streams and energy integration in chemical processes creates unique features for plantwide control because of the potential for disturbance propagation and the alteration of the system's dynamic behavior. In this article we have presented a general heuristic design procedure that generates a workable plantwide control structure for processes involving reaction and separation sections. Since this is an open-ended design problem, there is no unique correct solution.

Five of the nine steps in the procedure deal with plantwide issues that would not be addressed by simply combining the control systems from all of the plant's individual units: establishing an energy management system to dissipate exothermic heats of reaction and prevent the propagation of thermal disturbances via energy integration; setting production rate; controlling product quality and handling safety, operational, and environmental constraints; controlling pressure and level inventories for total mass balance; and accounting for nonconserved chemical components.

Application of the procedure is illustrated in this article with three industrial process examples: the vinyl acetate monomer process, Eastman process, and HDA process. The procedure produces a workable plantwide control strategy for a given process design, which is one part of the problem. Since the design of a process profoundly affects its dynamic controllability, another part of the problem's open-ended nature is the opportunity to change the *process* design. The design-and-control interaction problem remains as yet an open research area in terms of the plantwide control problem.

Literature Cited

Banerjee, A., and Y. Arkun, "Control Configuration Design Applied to the Tennessee Eastman Plant-Wide Control Problem," *Comput. Chem. Eng.*, **19**, 453 (1995).

Buckley, P. S., *Techniques of Process Control*, Wiley, New York (1964).

Douglas, J. M., *Conceptual Design of Chemical Processes*, McGraw-Hill, New York (1988)

Downs, J. J., "Distillation Control in a Plantwide Control Environment," *Practical Distillation Control*, Chap. 20, W. L. Luyben, ed., Van Nostrand Reinhold, New York (1992).

Downs, J. J., and E. F. Vogel, "A Plant-Wide Industrial Process Control Problem," *Comput. Chem. Eng.*, **17**, 245 (1993).

Fisher, W. R., M. F. Doherty, and J. M. Douglas, "The Interface Between Design and Control. 3. Selecting a Set of Controlled Variables," *Ind. Eng. Chem. Res.*, **27**, 611 (1988).

Kanadibhotla, R. S., and J. B. Riggs, "Nonlinear Model Based Control of a Recycle Reactor Process," *Comput. Chem. Eng.*, **19**, 933 (1995).

Luyben, M. L., and B. D. Tyreus, "An Industrial Design/Control Study for the Vinyl Acetate Monomer Process," *Comput. Chem. Eng.*, submitted (1997).

Luyben, M. L., and W. L. Luyben, "Design and Control of a Complex Process Involving Two Reaction Steps, Three Distillation Columns, and Two Recycle Streams," *Ind. Eng. Chem. Res.*, **34**, 3885 (1995).

Luyben, M. L., B. D. Tyreus, and W. L. Luyben, "Analysis of Control Structures for Reaction/Separation/Recycle Processes with Second-Order Reactions," *Ind. Eng. Chem. Res.*, **35**, 758 (1996).

Lyman, P. R., and C. Georgakis, "Plant-Wide Control of the Tennessee Eastman Problem," *Comput. Chem. Eng.*, **19**, 321 (1995).

McAvoy, T. J., and N. Ye, "Base Control for the Tennessee Eastman Problem," *Comput. Chem. Eng.*, **18**, 383 (1994).

McAvoy, T. J., N. Ye, and C. Gang, "Nonlinear Inferential Parallel Cascade Control," *Ind. Eng. Chem. Res.*, **35**, 130 (1996).

Narraway, L. T., and J. D. Perkins, "Selection of Process Control Structure Based on Linear Dynamic Economics," *Ind. Eng. Chem. Res.*, **32**, 2681 (1993).

Ponton, J. W., and D. M. Laing, "A Hierarchical Approach to the Design of Process Control Systems," *Chem. Eng. Res. Des.*, **71**, 181 (1993).

Price, R. M., P. R. Lyman, and C. Georgakis, "Throughput Manipulation in Plantwide Control Structures," *Ind. Eng. Chem. Res.*, **33**, 1197 (1994).

Price, R. M., and C. Georgakis, "Plantwide Regulatory Control Design Procedure Using a Tiered Framework," *Ind. Eng. Chem. Res.*, **32**, 2693 (1993).

Ricker, N. L., "Decentralized Control of the Tennessee Eastman Challenge Process," *J. Process Control*, **6**, 205 (1996).

Ricker, N. L., and J. H. Lee, "Nonlinear Model Predictive Control of the Tennessee Eastman Challenge Process," *Comput. Chem. Eng.*, **19**, 961 (1995).

Shinskey, F. G., *Process Control Systems*, 3rd ed., McGraw-Hill, New York (1988).

Stephanopoulos, G., *Chemical Process Control*, Prentice Hall, New York (1984).

Stephanopoulos, G., "Synthesis of Control Systems for Chemical Plants—A Challenge for Creativity," *Comput. Chem. Eng.*, **7**, 331 (1983).

Tyreus, B. D., and W. L. Luyben, "Dynamics and Control of Recycle Systems: 4. Ternary Systems with One or Two Recycle Streams," *Ind. Eng. Chem. Res.*, **32**, 1154 (1993).

Wolff, E. A., and S. Skogestad, "Operability of Integrated Plants," *Proc. of PSE*, Kyongju, Korea (1994).

Yi, C. K., and W. L. Luyben, "Evaluation of Plantwide Control Structures by Steady-State Disturbance Sensitivity Analysis," *Ind. Eng. Chem. Res.*, **34**, 2393 (1995).

Manuscript received Dec. 23, 1996, and revision received Aug. 22, 1997.