

Introduction

1.1 Overview

Plantwide process control involves the systems and strategies required to control an entire chemical plant consisting of many interconnected unit operations.

One of the most common, important, and challenging control tasks confronting chemical engineers is: How do we design the control loops and systems needed to run our process? We typically are presented with a complicated process flowsheet containing several recycle streams, energy integration, and many different unit operations: distillation columns, reactors of all types, heat exchangers, centrifuges, dryers, crystallizers, liquid-liquid extractors, pumps, compressors, tanks, absorbers, decanters, etc. Given a complex, integrated process and a diverse assortment of equipment, we must devise the necessary logic, instrumentation, and strategies to operate the plant safely and achieve its design objectives.

This is, in essence, the realm of control system synthesis for an entire plant. What issues do we need to consider? What is of essential importance within this immense amount of detail? How does the dynamic behavior of the interconnected plant differ from that of the individual unit operations? What, if anything, do we need to model or test? How do we even begin?

This book addresses each of these questions and explains the fundamental ideas of control system synthesis. As its core, the book presents a general heuristic design procedure that generates an effective plantwide base-level regulatory control structure for an *entire, complex* process flowsheet and not simply individual units.

The nine steps of the design 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.

We first review in Part 1 the basics of plantwide control. We illustrate its importance by highlighting the unique characteristics that arise when operating and controlling complex integrated processes. The steps of our design procedure are described. In Part 2, we examine how the control of individual unit operations fits within the context of a plantwide perspective. Reactors, heat exchangers, distillation columns, and other unit operations are discussed. Then, the application of the procedure is illustrated in Part 3 with four industrial process examples: the Eastman plantwide control process, the butane isomerization process, the HDA process, and the vinyl acetate monomer process.

1.2 HDA Process

Let's begin with an example of a real industrial process to highlight what we mean by *plantwide process control*. The hydrodealkylation of toluene (HDA) process is used extensively in the book by Douglas (1988) on conceptual design, which presents a hierarchical procedure for generating steady-state flowsheet structures. Hence the HDA process should be familiar to many chemical engineering students who have had a course in process design. It also represents a flowsheet topology that is similar to many chemical plants, so practicing engineers should recognize its essential features.

The HDA process (Fig. 1.1) 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 kinetic rate expressions are functions of the partial pressures of toluene p_T , hydrogen p_H , benzene p_B , and diphenyl p_D , with an Arrhenius temperature dependence. By-product diphenyl is produced in an equilibrium reaction.

$$r_1 = k_{1(T)} p_T p_H^{1/2} \quad (1.3)$$

$$r_2 = k_{2(T)} p_B^2 - k_{2(T)} p_D p_H \quad (1.4)$$

The two fresh reactant makeup feed streams (one gas for hydrogen

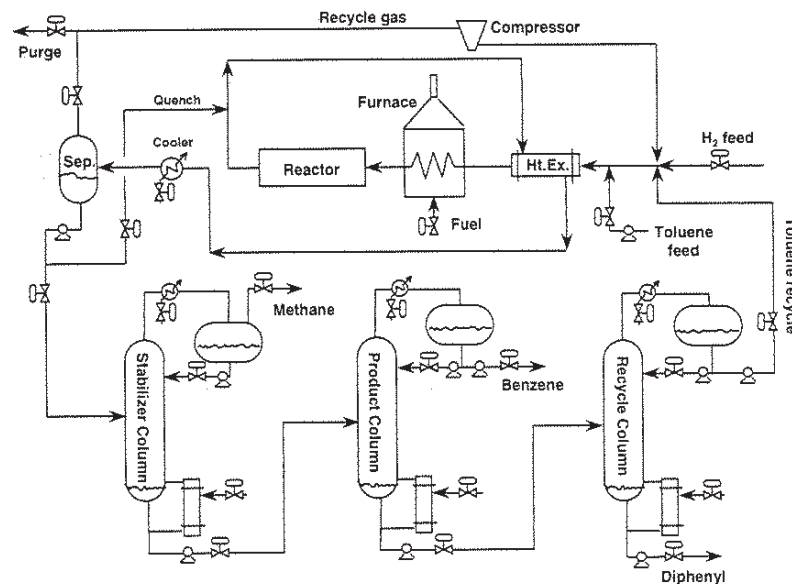


Figure 1.1 HDA process flowsheet.

and one liquid for toluene) are combined with the gas and liquid recycle streams. This combined stream is the cold inlet feed to the process-to-process heat exchanger, where the hot stream is the reactor effluent after the quench. The cold outlet stream is heated further, via combustion of fuel in the furnace, up to the required reactor inlet temperature. 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 hot outlet stream from the process-to-process heat exchanger goes to a partial condenser and then to a vapor-liquid separator. The gas stream from the overhead of the separator recycles unconverted hydrogen plus methane back to the reactor via a compressor. Since methane enters as an impurity in the hydrogen feed stream and is further produced in the reactor, it will accumulate in the gas recycle loop. Hence a purge stream is required to remove methane from the process. Part of the liquid from the separator serves as the reactor quench stream.

The remainder of the liquid from the separator is fed to the stabilizer column to remove any of the remaining hydrogen and methane gas from the aromatic liquids. The bottoms stream from the stabilizer column feeds the product column, which yields the desired product ben-

zene in the distillate. The by-product diphenyl exits from the process in the bottoms stream from the recycle column, which is fed from the bottoms of the product column. The liquid distillate stream from the recycle column returns unconverted toluene to the reactor.

Given this process flowsheet, we'd like to know how we can run this process to make benzene. We naturally have a lot of questions we want answered about operating this plant:

- How do we control the reactor temperature to prevent a runaway?
- How can we increase or decrease the production rate of benzene depending upon market conditions?
- How do we ensure the benzene product is sufficiently pure for us to sell?
- How do we know how much of the fresh hydrogen and toluene feed streams to add?
- How do we determine the flowrate of the gas purge stream?
- How can we minimize the raw material yield loss to diphenyl?
- How do we prevent overfilling any liquid vessels and overpressuring any units?
- How do we deal with units tied together with heat integration?
- How can we even test any control strategy that we might develop?

Answering these questions is not at all a trivial matter. But these issues lie at the foundation of control system synthesis for an entire plant. The plantwide control problem is extremely complex and very much open-ended. There are a combinatorial number of possible choices and alternative strategies. And there is no unique "correct" solution.

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. However, it is possible to employ a systematic strategy (or engineering method) to get a feasible solution. Our framework in tackling a problem of this complexity is based upon heuristics that account for the unique features and concerns of integrated plants. This book presents such a general plantwide control design procedure.

The scope embraces continuous processes with reaction and separation sections. Because our approach in this book is based upon a plantwide perspective, we cover what is relevant to this particular area. We omit much basic process control material that constitutes the framework and provides the tools for dynamic analysis, stability, system identification, and controller tuning. But we refer the interested reader

to Luyben and Luyben (1997) and other chemical engineering textbooks on process control.

1.3 History

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. And 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 flowsheet 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. 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 capital investment, working capital, and operating cost and to respond to safety and environmental concerns. This has prompted design engineers to start eliminating many surge tanks, increasing recycle streams, and introducing heat integration for both existing and new plants. Often this is done without a complete understanding of their effects 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. Im-

proved product yields, which reduce raw material costs, are achieved via lower reactant per-pass conversion and higher material recycle rates through the process. Better product quality, energy integration, and higher yields are all economically attractive in the steady-state flowsheet, but they present significant challenges to smooth dynamic plant operation. Hence an effective control system regulating the entire plant operation and a process designed with good dynamic performance play critical parts in achieving the business objectives of reducing operating and capital costs.

Buckley (1964) proposed a control design procedure for the plantwide control problem that consisted of two stages. The first stage 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, the two-stage 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.

We believe that chemical process control must move beyond the sphere of unit operations into the realm of viewing the plant as a whole system. The time is ripe in the chemical and petroleum industry for the development of a plantwide control design procedure. The technology, insight, and understanding have reached a state where general guidelines can be presented. The computer software needed for plantwide dynamic simulations is becoming commercially available. While linear methods are very useful to analyze control concepts, we strongly believe that the final evaluation of any plantwide control structure requires rigorous nonlinear dynamic simulations, not linear transfer function analysis.

1.4 Model-Based and Conventional Control

Some people claim that the plantwide control problem has already been solved by the application of several commercial forms of model predictive control (MPC). MPC rests on the idea that we have a fair amount of knowledge about the dynamic behavior of the process and that this knowledge can be incorporated into the *controller* itself. The controller uses past information and current measurements to predict

the future response and to adjust its control valves so that this anticipated response is optimal in some sense.

Model predictive control is particularly useful when several control valves (or manipulators) affect an output of interest (what is called *interaction*) and also when some sort of constraint comes into play either on the inputs or on some measured variable. Since the controller itself *knows* about these interactions and constraints, it can in theory avoid those perils. It is important to remember that MPC merely suggests that the controller can predict the process response into the future, only to be checked (and corrected) by the next round of measurements.

On the other hand, conventional control approaches also rely on models, but they are usually not built into the controller itself. Instead the models form the basis of simulations and other analysis methods that guide in the selection of control loops and suggest tuning constants for the relatively simple controllers normally employed [PI, PID, I-only, P-only, lead-lag compensation, etc. (P = proportional, PI = proportional-integral, PID = proportional-integral-derivative)]. Conventional control approaches attempt to build the *smarts* into the *system* (the process and the controllers) rather than only use complex control algorithms.

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 book.

Very few unbiased publications have appeared in the literature comparing control effectiveness using MPC versus a well-designed conventional control system. Most of the MPC applications reported have considered fairly simple processes with a small number of manipulated variables. There are no published reports that discuss the application of MPC to an entire complex chemical plant, with one notable exception. That is the work of Ricker (1996), who compared MPC with conventional PI control for the Eastman process (TE problem). His conclusion was “there appears to be little, if any, advantage to the use of nonlinear model predictive control (NMPC) in this application. In particular, the decentralized strategy does a better job of handling constraints—an area in which NMPC is reputed to excel.”

One of the basic reasons for his conclusion ties into the plantwide context that our procedure explicitly addresses, namely the need to regulate all chemical inventories. MPC gives no guidance on how to make the critical decisions of what variables need to be controlled. As Ricker states, “the naive MPC designer might be tempted to control only variables having defined setpoints, relying on optimization to make appropriate use of the remaining degrees of freedom. This fails in the

TE problem. As discussed previously, all chemical inventories must be regulated; it cannot be left to chance. Unless setpoints for key internal concentrations are provided, MPC allows reactant partial pressures to drift to unfavorable values." Our design procedure considers the concept of component balances as an explicit step in the design.

Another reason is related to the issue of constraints and priorities, which we address in the sequence of steps for our design procedure. Ricker says that "the TE problem has too many competing goals and special cases to be dealt with in a conventional MPC formulation." Normally this is addressed within MPC by the choice of weights, but for the Eastman process the importance of a variable changes depending upon the situation. "Ricker and Lee found that no single set of weights and constraints could provide the desired performance in all cases."

While we use conventional control systems here, our plantwide control design procedure does not preclude the use of MPC at a certain level. Our focus is on the issues arising from the operation of an integrated process. We find that a good control structure provides effective control, independent of any particular controller algorithm, while a poor one cannot be greatly improved with any algorithm (MPC or PID controllers).

1.5 Process Design

The traditional approach to developing a new process has been to perform the design and control analyses sequentially. First, the design engineer constructs a steady-state process flowsheet, with particular structure, equipment, design parameters, and operating conditions. The objective is to optimize the economics of the project in evaluating the enormous number of alternatives. The hierarchical design procedure proposed by Douglas (1988) is a way to approach this task. Little attention is given to dynamic controllability during the early stages of the design.

After completion of the detailed design, the control engineer then must devise the control strategies to ensure stable dynamic performance and to satisfy the operational requirements. The objective is to operate the plant in the face of potentially known and unknown disturbances, production rate changes, and transitions from one product to another.

While this staged approach has long been recognized as deficient, it is defensible from a certain perspective. For example, it would be difficult for the control engineers to specify the instrumentation and the distributed control system (DCS) without knowing exactly what process it was intended for. Similarly, it would make no sense for the process engineers to request a control system design for all those flowsheets

that were considered but rejected on the basis of steady-state economics alone. However, this staged approach can result in missed opportunities because of the close connection between process design and controllability. How a process is designed fundamentally determines its inherent controllability, which means qualitatively how well the process rejects disturbances and how easily it moves from one operating condition to another. In an ideal project system, dynamics and control strategies would be considered during the process synthesis and design activities.

This issue grows increasingly important as plants become more highly integrated with complex configurations, recycle streams, and energy integration. Competitive economic pressures, safety issues, and environmental concerns have all contributed to this. However, if a control engineer becomes involved early enough in the process design, he or she may be able to show that it would be better in the long run to build a process with higher capital and utility costs if that plant provides more stable operation and less variability in the product quality.

We believe that process design impacts controllability far more than control algorithms do. We base our opinion on many years of experience. We have participated as control engineers in many design projects. Some involved building new plants with new process technology, some involved new plants with existing technology, and some projects were modernizations of the control system on an existing plant. We have found that a consideration of dynamics and control strategies for new process designs has a much larger positive economic impact (when the design can potentially be modified) compared with control strategy upgrades on an existing process (with a fixed design). However, we stress that for those new plants and technologies we became involved before the process design was fixed. We performed dynamic simulations and undertook control system design as soon as the process engineers had an economically viable flowsheet. Most importantly, by working together with the process engineers and plant engineers, we changed the flowsheet until we were all satisfied that we had developed the most profitable process when viewed over the entire life time of the project. This inevitably involved making trade-offs between steady-state investment economics and dynamic performance measured in uptime, throughput, product quality, and yield.

One of the important themes weaving through this book is the central role we place on the process design. Good control engineers need also to be good process engineers!

1.6 Spectrum of Process Control

We can view the field of process control as five parts of a continuous spectrum (Fig. 1.2). Each part is important, can be economically signifi-

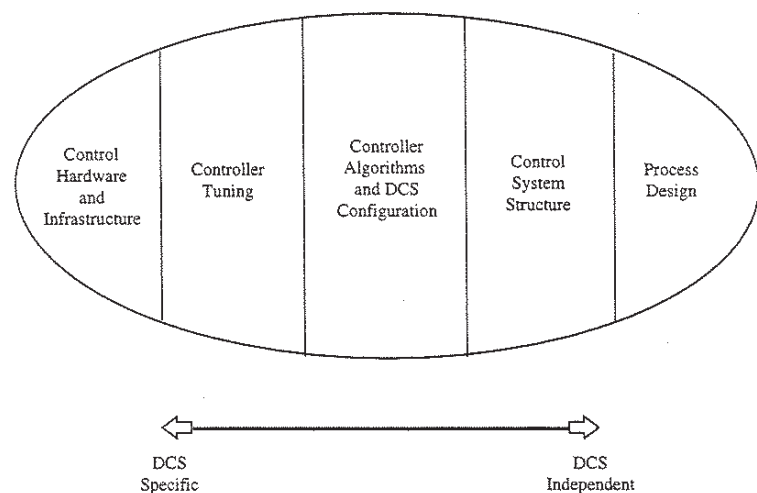


Figure 1.2 Spectrum of process control.

cant, and interacts in some manner with the others. Moving toward the left on the spectrum means dealing with more detailed issues on the level of the distributed control system (DCS). Moving toward the right means operating on a more general level with issues that are independent of the DCS.

The far left part of the spectrum deals with the control hardware and infrastructure required to operate a plant. We need to assemble the proper types of control valves and process measurements (for temperature, flow, pressure, composition, etc.). These are the sensory devices of the plant and are essential for any control system to function. Any control strategy, no matter how clever, will have severe difficulties without the right measurements and valves in the process. An Instrument Society of America (ISA) publication catalog (67 Alexander Drive, P.O. Box 12277, Research Triangle Park, NC 27709) contains many references that deal with control hardware.

The next part involves controller tuning. We must determine the tuning constants for the controllers in the plant. While this task is often performed by using heuristics and experience, it can sometimes be a nontrivial exercise for certain loops. We recommend using a relay-feedback test that determines the ultimate gain and period for the control loop, from which controller settings can be calculated (Luyben and Luyben, 1997).

The middle of the spectrum deals with the controller algorithms and DCS configuration. We must decide the type of controller to use (proportional, integral, derivative, multivariable, nonlinear, model pre-

dictive, etc.). We must also determine whether we need dynamic elements (lead/lags, feedforward, etc.) and how to handle overrides and interlocks. In addition, input and output variables must be assigned loop numbers, displays must be created, alarms must be specified, instrument groupings must be determined, etc.

The next part is the determination of the control system structure. We must decide what variables to control and manipulate and how these should be paired. The control structure is vitally important because a poor strategy will result in poor performance no matter what type of control algorithm we use or how much we tune it. There is little information or guidance in the literature or in process control textbooks (both introductory and advanced) on how to develop an effective control structure for an entire complex chemical plant. This is the main subject of this book.

The far right part of the spectrum is the design of the process itself. We sometimes can change the flowsheet structure, use different design parameters, and employ different types of process equipment to produce a plant that can be controlled more easily than other alternatives. At this level, a good process control engineer can potentially have an enormous economic impact. Most companies in the chemical and petroleum industries have had the unfortunate and unwelcome experience of building a plant that could not easily be started up because of operational difficulties arising from the plant design. Fixing these kinds of problems after the plant is built can often require large amounts of additional capital expense in addition to the lost sales opportunities.

In this book, we focus primarily on control structure selection. Interactions between design and control are illustrated by examples, and the effects of design parameters on control are discussed. However, we do not present a synthesis procedure for process design that is capable of generating the most controllable flowsheet for a given chemistry. This is still very much an open area for further research.

1.7 Conclusion

In this first chapter we have defined the plantwide process control problem. This was illustrated by using the HDA process, which will figure prominently in later parts of the book. We have provided a historical perspective and context. Finally we explained where the material in this book fits into the spectrum of process control activities.

1.8 References

- Buckley, P. S. *Techniques of Process Control*, New York: Wiley (1964).
 Douglas, J. M. *Conceptual Design of Chemical Processes*, New York: McGraw-Hill (1988).

- Luyben, W. L., and Luyben M. L. *Essentials of Process Control*, New York: McGraw-Hill (1997).
- Ricker, N. L. "Decentralized Control of the Tennessee Eastman Challenge Process," *J. Proc. Cont.*, **6**, 205–221 (1996).
- Stephanopoulos, G. "Synthesis of Control Systems for Chemical Plants—A Challenge for Creativity," *Comput. Chem. Eng.*, **7**, 331–365 (1983).

Plantwide Control Fundamentals

2.1 Introduction

In this chapter we examine some of the fundamental features and properties of the plantwide control problem. Our goal is to explain why we must design a control system from the viewpoint of the entire plant and not just combine the control schemes of each individual unit.

A typical chemical plant flowsheet has a mixture of multiple units connected both in series and in parallel. As noted in the previous chapter, the common topology consists of reaction sections and separation sections. Streams of fresh reactants enter the plant by being fed into the reaction section (or sometimes into the separation section) through a heat exchanger network. Here the chemical transformations occur to produce the desired species in one or more of a potentially wide array of reactor types: continuous stirred tank, tubular, packed bed, fluidized bed, sparged, slurry, trickle bed, etc.

The reactor effluent usually contains a mixture of reactants and products. It is fed into a separation section where the products are separated by some means from the reactants. Because of their economic value, reactants are recycled back to upstream units toward the reactor. The products are transported directly to customers, are fed into storage tanks, or are sent to other units for further processing. The separation section uses one or more of the fundamental unit operations: distillation, evaporation, filtration, crystallization, liquid-liquid extraction, adsorption, absorption, pressure-swing adsorption, etc. In this book we typically use distillation as the separation method because of its widespread use and our considerable experience with it. Everyone is a victim of his or her experience. Our backgrounds are in petroleum processing

and chemical manufacturing, where distillation, despite frequently occurring predictions to the contrary, remains the premier separation method. However, the general principles also apply to processes with other separation units.

In addition to recycle streams returned back to upstream units, thermal integration is also frequently done. Energy integration can link units together in locations anywhere in the flowsheet where the temperature levels permit heat transfer to occur. The reaction and separation sections are thus often intimately connected. If conditions are altered in the reaction section, the resulting changes in flowrates, compositions, and temperatures affect the separation section and vice versa.

Changes in temperatures and thermal conditions can propagate into the separation section and significantly degrade dynamic performance. Changes in flowrates create load disturbances that can be recycled around a material loop. Changes in stream compositions fed into the separation section are also troublesome disturbances because they alter separation requirements (the work of separation is often a strong function of the feed mixture composition). Significant shifts in the compositions and flowrates within the separation section are needed to achieve the desired purities of product and recycle streams. Achieving a composition change can sometimes take a long time because the component inventories within the separation section must be varied and this inherently governs the system's dynamic behavior.

So we must pay particular attention to the effects of the reaction section on the separation section. In this chapter we strip away all of the confusing factors associated with complex physical properties and phase equilibrium so that we can concentrate on the fundamental effects of flowsheet topology and reaction stoichiometry. Therefore, in the processes studied here, we use such simplifying assumptions as constant relative volatilities, equimolal overflow, and constant densities.

These "ideal" physical property assumptions may appear to represent an overly simplistic view of the problem. Our experience, however, is that we can often gain significant insight into the workings and interactions of processes with recycle streams by not confusing the picture with complexities such as azeotropes. Considering the complexities of a real chemical system is, of course, vital at some stage. But we attempt in this chapter to focus on the "forest" and not on the individual "trees."

For example, suppose there is a stream in the process that is a binary mixture of chemical components A and B . If these components obey ideal vapor-liquid equilibrium behavior, we can use a single distillation column to separate them. If they form an azeotrope, we may have to use a two-column separation scheme. If the azeotropic composition

changes significantly with pressure, we can use a two-column sequence with each column operating at different pressures. If the azeotrope is homogeneous and minimum boiling, the two fairly pure product streams can be produced as bottoms products from the two columns. So there are two columns in the nonideal case instead of one column in the ideal case. But the reaction section and the recycle streams really don't care if we have one column or two. The reactor sees the same types of disturbances coming from the separation section, perhaps with different dynamics but with similar steady-state effects. Since many of the important plantwide and recycle effects are really steady-state phenomena, the idealized single-column separation section yields results that are similar to those of the complex two-column separation section.

2.2 Integrated Processes

Three basic features of integrated chemical processes lie at the root of our need to consider the entire plant's control system: (1) the effect of material recycle, (2) the effect of energy integration, and (3) the need to account for chemical component inventories. If we did not have to worry about these issues, then we would not have to deal with a complex plantwide control problem. However, there are fundamental reasons why each of these exists in virtually all real processes.

2.2.1 Material recycle

Material is recycled for six basic and important reasons.

1. *Increase conversion:* For chemical processes involving reversible reactions, conversion of reactants to products is limited by thermodynamic equilibrium constraints. Therefore the reactor effluent by necessity contains both reactants and products. Separation and recycle of reactants are essential if the process is to be economically viable.
2. *Improve economics:* In most systems it is simply cheaper to build a reactor with incomplete conversion and recycle reactants than it is to reach the necessary conversion level in one reactor or several in series. The simple little process discussed in Sec. 2.6 illustrates this for a binary system with one reaction $A \rightarrow B$. A reactor followed by a stripping column with recycle is cheaper than one large reactor or three reactors in series.
3. *Improve yields:* In reaction systems such as $A \rightarrow B \rightarrow C$, where B is the desired product, the per-pass conversion of A must be kept low to avoid producing too much of the undesirable product C . Therefore

the concentration of B is kept fairly low in the reactor and a large recycle of A is required.

4. *Provide thermal sink:* In adiabatic reactors and in reactors where cooling is difficult and exothermic heat effects are large, it is often necessary to feed excess material to the reactor (an excess of one reactant or a product) so that the reactor temperature increase will not be too large. High temperature can potentially create several unpleasant events: it can lead to thermal runaways, it can deactivate catalysts, it can cause undesirable side reactions, it can cause mechanical failure of equipment, etc. So the heat of reaction is absorbed by the sensible heat required to raise the temperature of the excess material in the stream flowing through the reactor.
5. *Prevent side reactions:* A large excess of one of the reactants is often used so that the concentration of the other reactant is kept low. If this limiting reactant is not kept in low concentration, it could react to produce undesirable products. Therefore the reactant that is in excess must be separated from the product components in the reactor effluent stream and recycled back to the reactor.
6. *Control properties:* In many polymerization reactors, conversion of monomer is limited to achieve the desired polymer properties. These include average molecular weight, molecular weight distribution, degree of branching, particle size, etc. Another reason for limiting conversion to polymer is to control the increase in viscosity that is typical of polymer solutions. This facilitates reactor agitation and heat removal and allows the material to be further processed.

2.2.2 Energy integration

The fundamental reason for the use of energy integration is to improve the thermodynamic efficiency of the process. This translates into a reduction in utility cost. For energy-intensive processes, the savings can be quite significant. We can illustrate the use and benefits of energy-integration by considering again the HDA process introduced in the previous chapter (Fig. 1.1). Here energy is required to heat up the reactants in the furnace and to provide boilup in the three distillation columns. Heat must be removed in the separator condenser and in the three column condensers. Heat is generated in the exothermic reactor that normally would be removed through the plant utility system. However, by using a feed/effluent heat exchanger we can recover some of that energy. This reduces the amount of fuel required in the furnace to heat up the reactants and the duty required to cool the reactor effluent stream.

In fact we could theoretically introduce considerably more energy

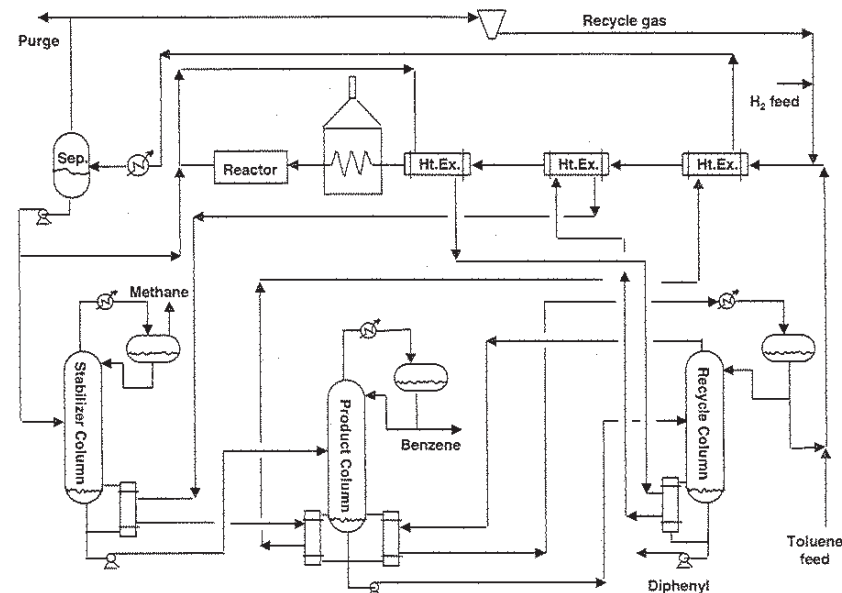


Figure 2.1 HDA process flowsheet with complex heat integration.

integration into the HDA process (Fig. 2.1). This is alternative 6 from the paper by Terrill and Douglas (1987). Heat from the reactor is used in reboilers of all three distillation columns. In addition, condensation of the overhead vapor from the recycle column provides heat input to the base of the product column. This is a good illustration of how units anywhere in the process can be linked together thermally. Figure 2.1 also shows how complex heat-integrated processes can quickly become, creating nontrivial control issues. This highlights why we cannot combine the control systems of individual unit operations in such processes.

2.2.3 Chemical component inventories

We can characterize a plant's chemical species into three types: reactants, products, and inerts. A material balance for each of these components must be satisfied. This is typically not a problem for products and inerts. However, the real problem usually arises when we consider reactants (because of recycle) and account for their inventories within the entire process. Every molecule of reactants fed into the plant must either be consumed via reaction or leave as an impurity or purge. Because of their value, we want to minimize the loss of reactants exiting the process since this represents a yield penalty. So we prevent

reactants from leaving. This means we must ensure that every mole of reactant fed to the process is consumed by the reactions.

This is an important concept and is generic to many chemical processes. From the viewpoint of individual units, chemical component balancing is not a problem because exit streams from the unit automatically adjust their flows and compositions. However, when we connect units together with recycle streams, the entire system behaves almost like a pure integrator in terms of the reactants. If additional reactant is fed into the system without changing reactor conditions to consume the reactant, this component will build up gradually within the plant because it has no place to leave the system.

Plants are not necessarily self-regulating in terms of reactants. We might expect that the reaction rate will increase as reactant composition increases. However, in systems with several reactants (e.g., $A + B \rightarrow \text{products}$), increasing one reactant composition will decrease the other reactant composition with an uncertain net effect on reaction rate. Section 2.7 contains a more complete discussion of this phenomenon. Eventually the process will shut down when manipulated variable constraints are encountered in the separation section. Returning again to the HDA process, the recycle column can easily handle changes in the amount of (reactant) toluene inventory within the column. However, unless we can somehow account for the toluene inventory within the entire process, we could feed more fresh toluene into the process than is consumed in the reactor and eventually fill up the system with toluene.

The three features outlined in this section have profound implications for a plant's control strategy. Simple examples in this chapter will illustrate the effects of material recycle and component balancing. Chapter 5 contains more details of the effects created by energy integration on the entire plant.

2.3 Units in Series

If process units are arranged in a purely series configuration, where the products of each unit feed downstream units and there is no recycle of material or energy, the plantwide control problem is greatly simplified. We do not have to worry about the issues discussed in the previous section and we can simply configure the control scheme on each individual unit operation to handle load disturbances.

If production rate is set at the front end of the process, each unit will only see load disturbances coming from its upstream neighbor. If the plant is set up for "on-demand" production, changes in throughput will propagate back through the process. So any individual unit will see load disturbances coming from both its downstream neighbor (flowrate changes to achieve different throughputs) and its upstream neighbor

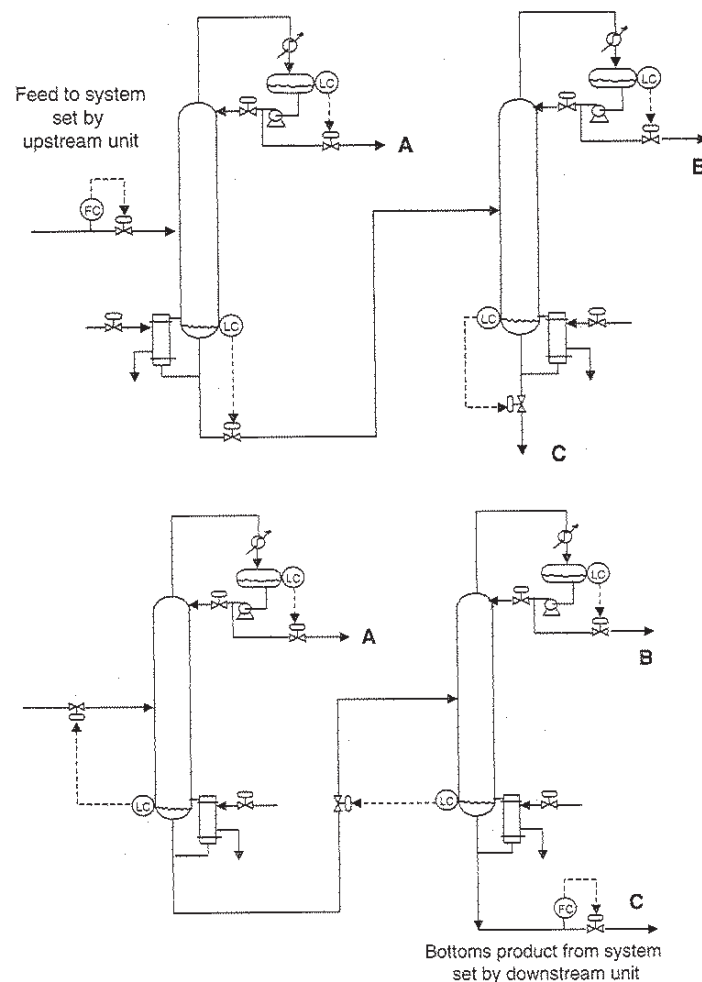


Figure 2.2 Units in series. (a) Level control in direction of flow; (b) level control in direction opposite flow.

(composition changes as the upstream units adjust to the load changes they see).

Figure 2.2 compares these two possible configurations for a simple plant. A fresh feed stream containing a mixture of chemical components A , B , and C is fed into a two-column distillation train. The relative volatilities are $\alpha_A > \alpha_B > \alpha_C$, and we select the "direct" (or "light-out-first") separation sequence: A is taken out the top of the first column and B out the top of the second column.

Figure 2.2a shows the situation where the fresh feed stream is flow-controlled into the process. The inventory loops (liquid levels) in each unit are controlled by manipulating flows leaving that unit. All disturbances propagate from unit to unit down the series configuration. The only disturbances that each unit sees are changes in its feed conditions.

Figure 2.2b shows the on-demand situation where the flowrate of product C leaving the bottom of the second column is set by the requirements of a downstream unit. Now some of the inventory loops (the base of both columns) are controlled by manipulating the feed into each column.

When the units are arranged in series with no recycles, the plantwide control problem can be effectively broken up into the control of each individual unit operation. There is no recycle effect, no coupling, and no feedback of material from downstream to upstream units. The plant's dynamic behavior is governed by the individual unit operations and the only path for disturbance propagation is linear along the process.

2.4 Effects of Recycle

Most real processes contain recycle streams. In this case the plantwide control problem becomes much more complex and its solution is not intuitively obvious. The presence of recycle streams profoundly alters the plant's dynamic and steady-state behavior. To gain an understanding of these effects, we look at some very simple recycle systems. The insight we obtain from these idealized, simplistic systems can be extended to the complex flowsheets of typical chemical processes. First we must lay the groundwork and have some feel for the complexities and phenomena that recycle streams produce in a plant.

In this section we explore two basic effects of recycle: (1) Recycle has an impact on the dynamics of the process. The overall time constant can be much different than the sum of the time constants of the individual units. (2) Recycle leads to the "snowball" effect. This has two manifestations, one steady state and one dynamic. A small change in throughput or feed composition can lead to a large change in steady-state recycle stream flowrates. These disturbances can lead to even larger dynamic changes in flows, which propagate around the recycle loop. Both effects have implications for the inventory control of components.

2.4.1 Time constants in recycle systems

Figure 2.3 gives a block-diagram representation of a simple process with recycle. The input to the system is u . We can think of this input as a flowrate. It enters a unit in the forward path that has a transfer

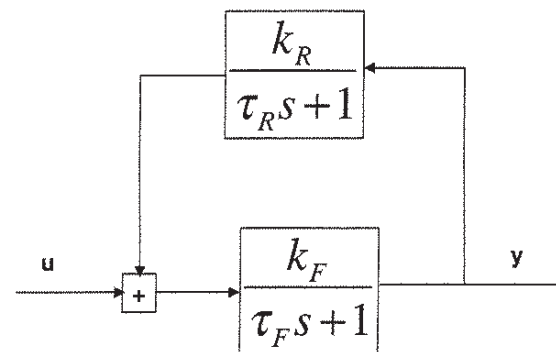


Figure 2.3 Simple block diagram of process with recycle.

function $G_{F(s)}$ that relates dynamically the input to the output of the unit. This transfer function consists of a steady-state gain K_F and a first-order lag with a time constant τ_F :

$$G_{F(s)} = \frac{K_F}{\tau_F s + 1} \quad (2.1)$$

The output of $G_{F(s)}$ is y , which also recycles back through a second transfer function $G_{R(s)}$ in the recycle path. This recycle transfer function also consists of a steady-state gain and a time constant.

$$G_{R(s)} = \frac{K_R}{\tau_R s + 1} \quad (2.2)$$

The output of the recycle block is added to the original input to the process u , and the sum of these two signals enters the forward block $G_{F(s)}$. It is important to note that the recycle loop in this process features *positive* feedback, not *negative* feedback that we are used to dealing with in feedback control. Most recycles produce this positive feedback behavior, which means that an increase in the recycle flowrate causes an increase in the flowrates through the process.

Some simple algebra gives the overall relationship for this system between input and output.

$$\begin{aligned} \frac{y(s)}{u(s)} &= \frac{\frac{K_F}{\tau_F s + 1}}{1 - \left(\frac{K_F}{\tau_F s + 1}\right)\left(\frac{K_R}{\tau_R s + 1}\right)} \\ &= \frac{K_F(\tau_R s + 1)}{\tau_F \tau_R s^2 + (\tau_F + \tau_R)s + (1 - K_F K_R)} \end{aligned} \quad (2.3)$$




Figure 2.4 Effect of recycle loop gain on overall dynamic response.

The denominator of the transfer function is the characteristic equation of any system, so the characteristic equation of this recycle system is

$$\tau_F \tau_R s^2 + (\tau_F + \tau_R)s + (1 - K_F K_R) = 0 \quad (2.4)$$

$$\frac{\tau_F \tau_R}{(1 - K_F K_R)} s^2 + \frac{(\tau_F + \tau_R)}{(1 - K_F K_R)} s + 1 = 0 \quad (2.5)$$

This is the standard form of a second-order system, whose time constant is $\sqrt{\tau_F \tau_R / (1 - K_F K_R)}$. As the loop gain in the system $K_F K_R$ (the product of the gains in all units in the forward and recycle path) approaches unity, the time constant of the overall process becomes large. Hence the time constant of an entire process with recycle can be much larger than any of the time constants of its individual units. Figure 2.4 illustrates this for several values of $K_F K_R$. The value of K_F is constant at unity for these plots, as are the values of τ_F and τ_R . We can see that the effective time constant of the overall process is 25 minutes when $K_R = 0.9$, while the time constants of the individual units are equal to 1 minute. The steady-state gain of the process is $K_F / (1 - K_F K_R)$, so the steady-state effect of the recycle stream also becomes larger as the loop gain approaches unity.

What are the implications of this phenomenon for the plantwide

control problem? It means that any change in a recycle process can take a long time to line out back to steady state. We are then tempted not to automate the control loops that handle inventories in recycle loops but rather let the operators manage them. Because the recycle effects are so slow, it is hard to recognize that there is a growing problem in the system inventory. It also takes an equally long time to rectify the situation. Intermediate vessel inventories may overflow or go empty. An imbalance may develop in the inventories of intermediate components. Whenever we do not account for this in the control strategy, the plant's separation section may be subjected to ramplike load disturbances. If the final product column sees this type of disturbance, the product quality controller has difficulty maintaining setpoint. To handle ramp disturbances, special low-frequency-compensated controllers can be used. But these types of controllers are not typically implemented either in conventional control or MPC systems (Belanger and Luyben, 1997). Morud and Skogestad (1996) present a more detailed analysis of the effect of material recycle and heat integration on the dynamic behavior of integrated plants.

2.4.2 Snowball effects

Another interesting observation that has been made about recycle systems is their tendency to exhibit large variations in the magnitude of the recycle flows. Plant operators report extended periods of operation when very small recycle flows occur. It is often difficult to turn the equipment down to such low flowrates. Then, during other periods when feed conditions are not very different, recycle flowrates increase drastically, usually over a considerable period of time. Often the equipment cannot handle such a large load.

We call this high sensitivity of the recycle flowrates to small disturbances the *snowball effect*. We illustrate its occurrence in the simple example below. It is important to note that this is *not* a dynamic effect; it is a *steady-state* phenomenon. But it does have dynamic implications for disturbance propagation and for inventory control. It has nothing to do with closed-loop stability. However, this does not imply that it is independent of the plant's control structure. On the contrary, the extent of the snowball effect is very strongly dependent upon the control structure used.

The large swings in recycle flowrates are undesirable in a plant because they can overload the capacity of the separation section or move the separation section into a flow region below its minimum turndown. Therefore it is important to select a plantwide control structure that avoids this effect. As the example below illustrates and as

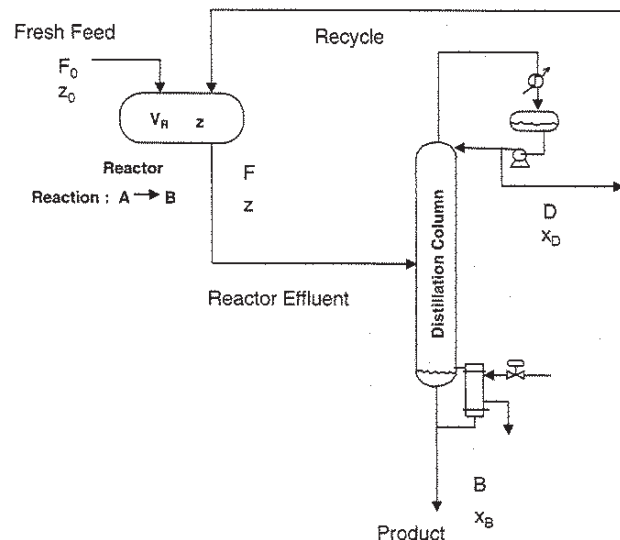


Figure 2.5 Flowsheet of binary recycle process.

more complex processes discussed in later chapters also show, a very effective way to prevent the snowball effect is to apply the following plantwide control heuristic:

A stream somewhere in each liquid recycle loop should be flow controlled.

Let us consider one of the simplest recycle processes imaginable: a continuous stirred tank reactor (CSTR) and a distillation column. As shown in Figure 2.5, a fresh reactant stream is fed into the reactor. Inside the reactor, a first-order isothermal irreversible reaction of component A to produce component B occurs $A \rightarrow B$. The specific reaction rate is k (h^{-1}) and the reactor holdup is V_R (moles). The fresh feed flowrate is F_0 (moles/h) and its composition is z_0 (mole fraction component A). The system is binary with only two components: reactant A and product B . The composition in the reactor is z (mole fraction A). Reactor effluent, with flowrate F (moles/h) is fed into a distillation column that separates unreacted A from product B .

The relative volatilities are such that A is more volatile than B , so the bottoms from the column is the product stream. Its flowrate is B (moles/h) and its composition is x_B (mole fraction A). The amount of A impurity in this product stream is an important control objective and must be maintained at some specified level to satisfy the product quality requirements of the customer.

The overhead distillate stream from the column contains almost all of component A that leaves the reactor because of the purity specifica-

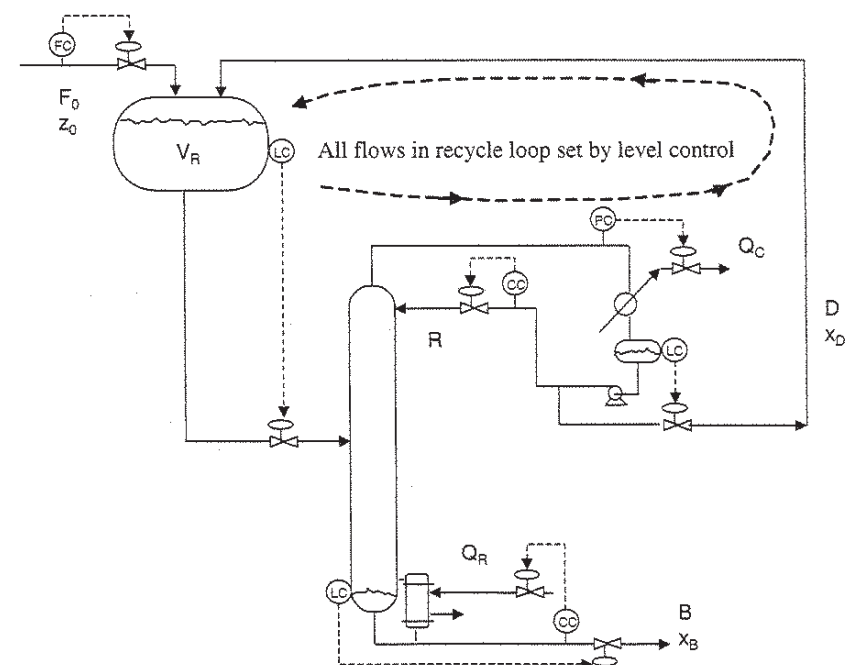


Figure 2.6 Conventional control structure with fixed reactor holdup.

tion on the bottoms stream. It is recycled back to the reactor at a flowrate D and with a composition x_D (mole fraction A). The column has N_T trays and the feed tray is N_F (counting from the bottom). The reflux flowrate is R and the vapor boilup is V (moles/h).

We now explore two alternative control structures for this process.

Conventional control structure. As shown in Fig. 2.6, the following control loops are chosen:

1. Fresh feed flow is controlled.
2. Reactor level is controlled by manipulating reactor effluent flow.
3. Bottoms product purity is controlled by manipulating heat input to the reboiler.
4. Distillate purity is controlled by manipulating reflux flow. Note that we have chosen to use dual composition control (controlling both distillate and bottoms purities) in the distillation column, but there is no *a priori* reason for holding the composition of the recycle stream constant since it does not leave the process. It may be useful to control the composition of this recycle stream for reactor yield pur-

poses or for improved dynamic response. We are often free to find the "best" recycle purity levels in both the design and operation of the plant.

5. Reflux drum level is held by distillate flow (recycle).
6. Base level is held by bottoms flow.
7. Column pressure is controlled by manipulating coolant flowrate to the condenser.

This control scheme is probably what most engineers would devise if given the problem of designing a control structure for this simple plant. Our tendency is to start with setting the flow of the fresh reactant feed stream as the means to regulate plant production rate. We would then work downstream from there as if looking at a steady-state flow-sheet and simply connect the recycle stream back to the reactor based upon a standard control strategy for the column.

However, we see in this strategy that there is no flow controller anywhere in the recycle loop. The flows around the loop are set based upon level control in the reactor and reflux drum. Given what we said above, we expect to find that this control structure exhibits the snowball effect. By writing the various overall steady-state mass and component balances around the whole process and around the reactor and column, we can calculate the flow of the recycle stream, at steady state, for any given fresh reactant feed flow and composition. The parameter values used in this specific numerical case are in Table 2.1.

With the control structure in Fig. 2.6 and the base-case fresh feed flow and composition, the recycle flowrate is normally 260.5 moles/h. However, the recycle flow must decrease to 205 moles/h when the fresh feed composition is 0.80 mole fraction A. It must increase to 330 moles/h when the fresh feed composition changes to pure A. Thus a 25 percent change in the disturbance (fresh feed composition) results in a 60 percent change in recycle flow. With this same control structure and the base-case fresh reactant feed composition, the recycle flow drops to 187 moles/h if the fresh feed flow changes to 215 moles/h. It

TABLE 2.1 Process Data

Base-case fresh feed composition	0.9	mole fraction A
Base-case fresh feed flowrate	239.5	moles/h
Reactor holdup	1250	moles
Reactor effluent flowrate	500	moles/h
Recycle flowrate	260.5	moles/h
Specific reaction rate	0.34086	h^{-1}
Bottoms composition	0.0105	mole fraction A
Recycle composition	0.95	mole fraction A

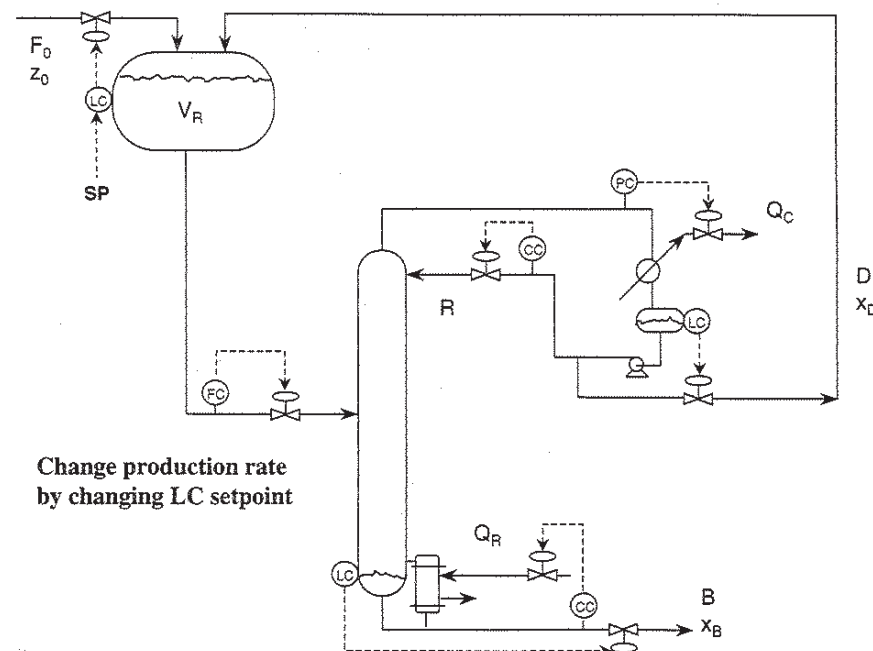


Figure 2.7 Control structure with variable reactor holdup.

must increase to 362 moles/h when the fresh feed flowrate is changed to 265 moles/h. Thus a 23 percent change in fresh feed flowrate results in a 94 percent change in recycle flowrate. These snowball effects are typical for many recycle systems when control structures such as that shown in Figure 2.6 are used and there is no flow controller somewhere in the recycle loop.

Variable reactor holdup structure. An alternative control structure is shown in Figure 2.7. This strategy differs from the previous one in two simple but important ways.

1. Reactor effluent flow is controlled.
2. Reactor holdup is controlled by manipulating the fresh reactant feed flowrate.

All other control loops are the same. We see here that we cannot change production rate directly by manipulating the fresh feed flow, because it is used to control reactor level. However, we must have some means to set plant throughput, which can be achieved indirectly in this scheme by changing the setpoint of the reactor level controller. Using the same

numerical case considered previously, the recycle flowrate does not change at all when the fresh feed composition changes. To alter production rate from 215 moles/h to 265 moles/h (a 23 percent change), the reactor holdup must be changed from 1030 moles/h to 1520 moles/h (a 48 percent change). Recycle flow also changes, but only from 285 to 235 moles/h. This is an 18 percent change in recycle flow compared with 94 percent in the alternative strategy.

What are the implications of this phenomenon for the plantwide control problem, when a small disturbance produces a proportionally larger change in recycle flow within the process? Although it is caused by steady-state issues, the snowball effect typically manifests itself in wide dynamic swings in stream flowrates that propagate around the recycle loop. This shows the strong connection between the reaction and separation sections. Whenever all flows in a recycle loop are set by level controllers, wide dynamic excursions occur in these flows because the total system inventory is not regulated. The control system is attempting to control the inventory in each individual vessel by changing the flowrate to its downstream neighbor. In a recycle loop, all level controllers see load disturbances coming from the upstream unit. This causes the flowrate disturbances to propagate around the recycle loop. Thus any disturbance that tends to increase the total inventory in the process (such as an increase in the fresh feed flowrate) will produce large increases in all flowrates around the recycle loop.

2.5 Reaction/Separation Section Interaction

For the process considered in the previous section where the reaction is $A \rightarrow B$, the overall reaction rate depends upon reactor holdup, temperature (rate constant), and reactant composition (mole fraction A) $\mathcal{R} = V_R k z$. The two control structures considered above produce fundamentally different behavior in handling disturbances. In the first, the separation section must absorb almost all of the changes. For example, to increase production rate of component B by 20 percent, the overall reaction rate must increase by 20 percent. Since both reactor temperature (and therefore k) and reactor holdup V_R are held constant, reactor composition z must increase 20 percent. This translates into a very significant change in the composition of the feed stream to the separation section. This means the load on the separation section changes significantly, producing large variations in recycle flowrates.

In the second structure, *both* reactor holdup V_R and reactor composition z can change, so the separation section sees a smaller load disturbance. This reduces the magnitude of the resulting change in recycle flow because the effects of the disturbance can be distributed between the reaction and separation sections.

If the tuning of the reactor level controller in the conventional struc-

ture (Fig. 2.6) is modified from normal PI to P only, then changes in production rate also produce changes in reactor holdup. This tends to compensate somewhat for the required changes in overall reaction rate and lessens the impact on the separation section. So both control system structure and the algorithm used in the inventory controller of the reactor affect the amount of this snowball phenomenon.

This example has a liquid-phase reactor, where volume can potentially be varied. If the reactor were vapor phase, reactor volume would be fixed. However, we now have an additional degree of freedom and can vary reactor pressure to affect reaction rate.

We can draw a very useful general conclusion from this simple binary system that is applicable to more complex processes: changes in production rate can be achieved only by changing conditions in the reactor. This means something that affects reaction rate in the reactor must vary: holdup in liquid-phase reactors, pressure in gas-phase reactors, temperature, concentrations of reactants (and products in reversible reactions), and catalyst activity or initiator addition rate. Some of these variables affect the conditions in the reactor more than others. Variables with a large effect are called *dominant*. By controlling the dominant variables in a process, we achieve what is called *partial control*. The term partial control arises because we typically have fewer available manipulators than variables we would like to control. The setpoints of the partial control loops are then manipulated to hold the important economic objectives in the desired ranges.

The plantwide control implication of this idea is that production rate changes should preferentially be achieved by modifying the setpoint of a partial control loop in the reaction section. This means that the separation section will not be significantly disturbed. Using the control structure in Fig. 2.6, changes in production rate require large changes in reactor composition, which disturb the column. Using the control structure shown in Fig. 2.7, changes in production rate are achieved by altering the setpoint of a controlled dominant variable, reactor holdup, with only small changes in reactor composition. This means that the column is not disturbed as much as with the alternative control scheme.

Hence a goal of the plantwide control strategy is to handle variability in production rate and in fresh reactant feed compositions while minimizing changes in the feed stream to the separation section. This may not be physically possible or economically feasible. But if it is, the separation section will perform better to accommodate these changes and to maintain product quality, which is one of the vital objectives for plant operation. Reactor temperature, pressure, catalyst/initiator activity, and holdup are preferred dominant variables to control compared to direct or indirect manipulation of the recycle flows, which of course affect the separation section.

In Chaps. 4 and 6 we discuss specific control issues for chemical reactors and distillation columns. We shall then have much more to say about the important concepts of dominant variables and partial control. Much of the material in those chapters centers on the control of the units individually. However, we also try to show how plantwide control considerations may sometimes alter the control strategy for the unit from what we would normally have in an isolated system.

Some of our previous discussion provides selected clues about why the “best” control structure for an isolated reactor or column may not be the best control strategy when plantwide dynamics are considered.

Let’s look again at the simple reactor/column process in Fig. 2.5. In Sec. 2.4.2 we proposed two control structures where both the bottoms composition x_B (the plant product) and the distillate composition x_D (the recycle stream) are controlled, i.e., dual composition control. Bottoms composition must be controlled because it is the product stream leaving the plant and sold to our customers. However, there is *a priori* no reason to control the composition of the recycle stream since this is an internal flow within the plant.

From the perspective of an isolated column, we can achieve better performance in bottoms product composition control by using simple single-end control. Dual composition control means two interacting control loops that normally must be detuned to achieve closed-loop stability. Single-end composition control means one SISO (single-input–single-output) loop that can be tuned up as tightly as the performance/robustness trade-off permits. If we look at just the operation of this distillation column with the control objective to do the best job we can to achieve on-aim product quality, then we would select a single-end control structure for the column.

However, our column is connected via material flow with a reactor. In Chap. 4 we show that reactor control often boils down to two issues: (1) managing energy (temperature control) and (2) keeping as constant as possible the composition and flowrate of the total reactor feed stream (fresh feed plus recycle streams). The latter goal implies that it may in fact be desirable to control the composition of the recycle stream. This minimizes the variability in recycle impurity composition back into the reactor. This recycle composition is dictated by the economic trade-offs between yield, conversion, energy consumption in the separation section, and reactor size.

Our plantwide control perspective may push us to use a dual composition control system on the column. We would have to loosen up the bottoms composition loop tuning. But smoother reactor operation may reduce disturbances to the column and result in better product quality control.

These are the issues discussed in Part 2: the control of unit operations individually and as part of a plantwide flowsheet.

2.6 Binary System Example

Our simple process considered previously was arbitrarily specified to contain a flowsheet with a reactor, column, and recycle stream. If we step farther back and consider the design of this process, we have many alternative ways to accomplish our objective, which is to take a fresh feed stream containing mostly reactant *A* and convert it into a stream of mostly product *B*. In addition to the reactor/column/recycle configuration, we could accomplish the same task by using one large CSTR or by using several CSTRs in series. In this section we analyze these alternatives quantitatively by comparing their steady-state economics (that is, which flowsheet gives the minimum total annual cost considering capital plus energy cost). Then we discuss the dynamic controllability of these alternative flowsheets.

2.6.1 Steady-state design

We neglect the energy cost of cooling the reactor because this will be essentially the same for all alternative flowsheets. Therefore designs with only reactors have to consider just the capital cost of the reactor. Designs with a reactor and column have both energy costs (heat input to the reboiler) and capital costs (reactor, column, reboiler, condenser, and trays). We use here the installed capital costs correlations given by Douglas (1988). The cost of the reactor is assumed to be 5 times the cost of a plain tank. We use a payback period of 3 years to calculate the annual cost of capital.

$$\text{Annual capital cost} = \frac{\text{total capital cost}}{3} \quad (2.6)$$

Table 2.2 gives equipment sizes and cost data for several alternative designs. Molecular weights are assumed for simplicity to be 50 lb/mole and density is 50 lb/ft³. An aspect ratio (diameter/length) of 0.5 is used.

TABLE 2.2 Economic Data for CSTRs

Number of CSTRs	1	2	3	4	5
Holdup per vessel, ft ³	59,523	5,802	2,395	1,435	1,009
Diameter, ft	33.6	15.5	11.5	9.7	8.63
Capital cost 10 ⁶ \$	11.8	5.56	4.81	4.66	4.68
Annual capital cost, 10 ⁶ \$/yr	3.95	1.86	1.60	1.55	1.56

TABLE 2.3 Economic Data for CSTR and Stripper

Reactor size, ft ³	800	1000	2500	5000
Reactor diameter, ft	7.98	8.6	11.7	14.7
Trays in stripper	14	14	17	16
Recycle composition, mole fraction A	0.873	0.761	0.391	0.215
Column diameter, ft	8.24	6.04	3.71	3.21
Reboiler energy, 10 ⁶ Btu/h	25.1	13.5	5.09	3.80
Area condenser, ft ²	8360	4496	1697	1267
Area reboiler, ft ²	5020	2698	1018	760
Capital cost, \$1000:				
Reactor	810	1090	1645	2535
Column	304	218	152	124
Reboiler	396	264	140	116
Condenser	552	369	196	162
Trays	13	8	5	4
Total capital cost (10 ⁶ \$)	2.075	1.949	2.138	2.941
Annual costs, 10 ⁶ \$/yr:				
Energy	1.099	0.591	0.223	0.167
Capital	0.692	0.650	0.713	0.980
Total annual cost, 10 ⁶ \$/yr	1.79	1.24	<u>0.936</u>	1.15

Additional details of the economic and sizing calculations can be found in Luyben (1993). Notice that the flowsheet with the smallest annual cost has four CSTRs. Now let's compare this system with a process that has one CSTR and a column whose overhead product is recycled back to the reactor. Economic studies of this system have shown that a simple stripping column is cheaper than a full column. Table 2.3 gives size and cost data over a range of reactor sizes.

This simplistic economic evaluation shows that the reactor/stripper process is more economical than the reactors-in-series process. A 2500 ft³ reactor followed by a stripping column can achieve the same result that would require four 1435 ft³ reactors in series with no recycle.

In the simple binary process considered above, the 2500 ft³ reactor with a 17-tray stripper gives the process with the smallest total annual cost: \$936,000/yr versus \$1,550,000/yr for the best of the CSTR-in-series flowsheets. Thus this process with recycle is more economical, from the viewpoint of steady state, than the alternative process consisting of reactors in series. This is the point we made in Sec. 2.2 about the economic advantage for recycle.

2.6.2 Dynamic controllability

Dynamic simulations of two alternative processes provide a quantitative comparison of their dynamic controllabilities. To strike a balance between simplicity and the economic optimum, we selected the three-CSTR process to compare with the reactor/stripper process. The scheme




Figure 2.8 Three-CSTR control structure.

used for the three-CSTR process controls the composition of the final product leaving the third reactor (z_3) by changing the setpoint signal to three level controllers for the three vessels (Fig. 2.8). The composition controller has PI action with $K_c = 1$ and $\tau_I = 10.2$ min. A composition transmitter deadtime of 3 minutes is used. Fresh feed is flow-controlled. Level controllers are proportional-only with gains of 10.

The scheme for the reactor/stripper process uses a PI controller to hold product composition (x_B) by manipulating vapor boilup in the stripper. The same analyzer deadtime is used. Proportional level controllers are used for the stripper base (manipulating bottoms flow), the overhead receiver (manipulating recycle flow), and the reactor (manipulating reactor effluent flow) with gains of 2.

Figures 2.9 and 2.10 show the dynamic responses of the two alternative processes for step changes in the fresh feed composition z_0 and fresh feed flowrate F_0 . Note the differences in the time scales. The three-CSTR process takes much longer to settle out after the disturbance occurs. However, the maximum deviation of product purity is about half that experienced with the reactor/stripper process. The large holdups in the three reactors filter the disturbances but also slow the process response.

Because the reactor/stripper process is much more attractive economically, it may be the flowsheet of choice despite its larger short-term variability in product quality. This illustrates how plants with recycle are generally more difficult to control than units in series.

2.7 Ternary System Example

We now move on to study another simple process, but again we gain a considerable amount of insight into some important generic concepts for both process design and control (Tyreus and Luyben, 1993). Here

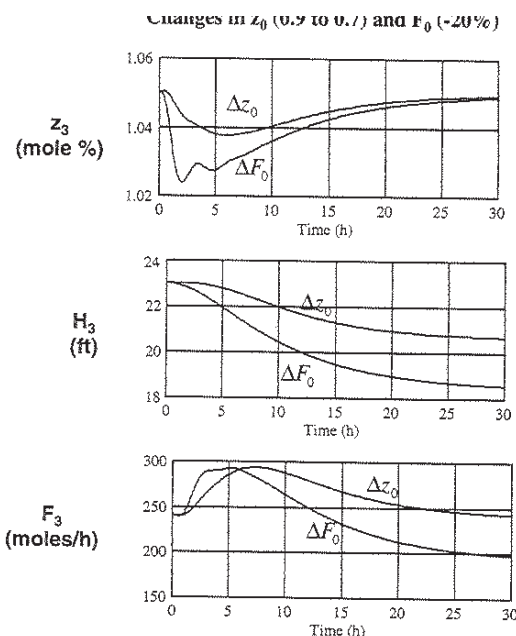


Figure 2.9 Dynamic responses of three-CSTR process.

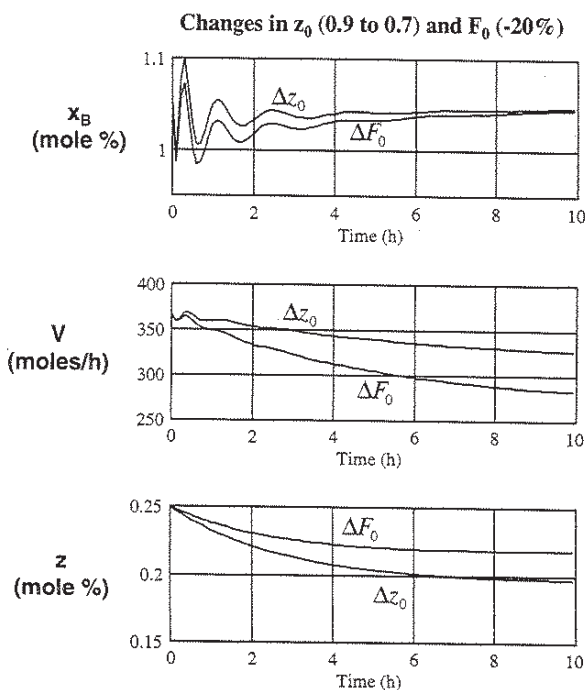


Figure 2.10 Dynamic responses of reactor/stripper process.

we consider a reactor where two reactants A and B form product C : $A + B \rightarrow C$. Since there are three components, we call this system a *ternary example*.

Two kinetic cases will be considered. In the first we assume the reaction rate is so large that the limiting reactant B is completely consumed in the reactor, i.e., there is 100 percent per pass conversion of B . The reactor effluent contains only excess reactant A and product C , so the separation section deals with these two components and recycles A back to the reaction section. In industrial processes, this type of system is typically encountered with extremely hazardous reactants, which we want to be completely consumed in the reactor.

In the second case, which is more general for industrial processes, the reaction rate is not large, so complete one-pass conversion of one reactant would require an excessively large reactor. Economics dictate that reactant concentrations must be significant and recycling of reactants is required. Now the separation section must recover both reactants for recycle.

2.7.1 Complete one-pass reactant conversion

Figure 2.11 shows the ternary process where no B is in the reactor effluent. The size of the reactor and concentrations in the reactor are arbitrary because the consumption of B is independent of these variables. We assume that the separation section consists of a single distillation column. If A is more volatile than C , the overhead product from the column is recycled back to the reactor. If the volatilities are reversed, the bottoms from the column is the recycle stream. Figure 2.11 illustrates the first case.

Two control structures are shown in Fig. 2.11a and b. In both, the composition of component A in the product stream $x_{B,A}$ is controlled by manipulating vapor boilup in the column. This prevents component A from leaving the system. Except for this small amount of A impurity in the product, all A that enters the system must be consumed in the reactor. This illustrates the point we made in Sec. 2.2 about the need to change conditions in the reactor so that the additional reactant is consumed and will not accumulate.

In the first control structure (Fig. 2.11a), both fresh reactant feeds are flow-controlled into the system, with one of the reactants ratioed to the other. This type of control structure is seen quite frequently because we want to set production rate with a reactant feed flow and we know that a stoichiometric ratio of reactants is needed. Unfortunately this strategy does *not* work! It is not possible to feed exactly the stoichiometric amounts of the two reactants. Inaccuracies in flow measurement prevent this from occurring in practice with real instru-




Figure 2.11 Ternary process with complete one-pass conversion of reactant B. (a) Ratio control structure with fixed reactant feed (unworkable); (b) reactant makeup control based on component inventory (workable).

mentation. But even if the flow measurements were perfect, the slightest change in fresh feed compositions would cause the same component imbalance problem. Unless the amounts of the two reactants are perfectly balanced, a gradual buildup will occur of whichever component is in excess. This phenomenon may take hours, days, or weeks. The time depends upon the amount of mismatch between A and B feeding the system.

In the second control structure (Fig. 2.11b), which *does* work, the fresh feed makeup of the limiting reactant (F_{OB}) is flow-controlled. The other fresh feed makeup stream (F_{OA}) is brought into the system to control the liquid level in the reflux drum of the distillation column. The inventory in this drum reflects the amount of A inside the system. If more A is being consumed by reaction than is being fed into the process, the level in the reflux drum will go down. Thus this control structure employs knowledge about the amount of component A in the system to regulate this fresh reactant feed makeup to balance exactly the amount of B fed into the process.

Notice that the total rate of recycle plus fresh feed of A is flow-controlled. There is a flow controller in the recycle loop, which prevents the snowball effect. Sometimes the fresh feed of A is added directly into the reflux drum, making the effect of its flow on reflux drum level more obvious. The piping system where it is not added directly to the drum still gives an immediate effect of makeup flow on drum level because the flowrate of the total stream (recycle plus fresh feed) is held constant. If the fresh feed flow increases, flow from the drum decreases, and this immediately begins to raise the drum level.

2.7.2 Incomplete conversion of both reactants

Now let us consider what is the more common situation where both reactants are present in the reactor effluent. The reaction rate in the reactor \mathcal{R} depends upon the holdup in the reactor V_R , the temperature (through the specific reaction rate k), and the concentrations of both reactants (z_A and z_B):

$$\mathcal{R} = kV_R z_A z_B \quad (2.7)$$

An infinite number of operating conditions in the reactor give exactly the same reaction rate but have different reactor compositions. The only requirement is that the product of the two concentrations (z_A times z_B) be constant. For a given reactor size and temperature, we can have any number of different reactor compositions, and these reactor compositions have a strong impact on the separation system. If z_A is large and z_B is small, there must be a large recycle of A and a small recycle

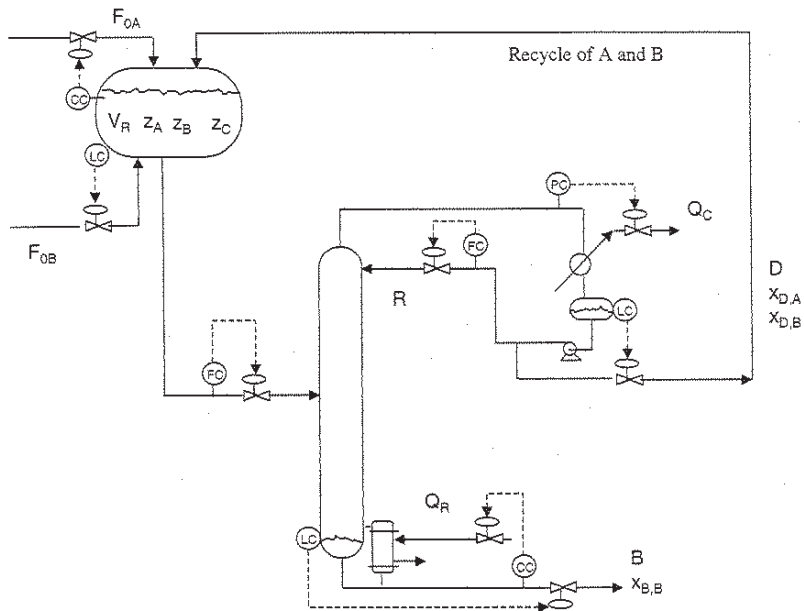


Figure 2.12 Ternary process flowsheet with incomplete conversion of both reactants and one recycle stream.

of B . If the compositions are reversed, the recycle flows are reversed in magnitude. We examine these alternatives later to see how they affect both steady-state economic design and dynamic controllability.

The separation section required to achieve reactant recycle depends upon the relative volatilities of the three components. We consider two cases: (1) the volatility of the product C is heavier or lighter than both of the reactants and (2) the volatility of the product C is intermediate between the reactants. In the first case, we need only one distillation column. In the second, we require two columns if we are limited to a simple two-product configuration.

Single-column case. Let us assume that the relative volatilities are $\alpha_A > \alpha_B > \alpha_C$, so the flowsheet shown in Figure 2.12 is appropriate. Product C is removed from the bottom of the column and contains a small amount of B impurity. It typically has no A because this is the most volatile component. Thus all the A and essentially all the B fed into the process must be consumed in the reactor. The recycle stream is a mixture of mostly B with a modest amount of A and some C . Economics dictate whether this recycle stream should be fairly pure (reducing reactor size but increasing separation costs) or impure.

The control structure shown in Figure 2.12 controls reactor effluent

flow to prevent the snowball effect, controls reactor composition by manipulating fresh feed F_{0A} , and controls reactor level with the fresh feed F_{0B} . Both controlled variables are dominant so we have effective partial control of the reactor. This control strategy works. It satisfies the stoichiometry by adjusting the fresh reactant feed flows.

We might be tempted to control reflux drum level with one of the fresh reactant feeds, as done above. The problem with this is that the material in the drum can contain a little of component C mixed with either A or B . Simply looking at the level doesn't tell us anything about component inventories within the process and which might be in excess. The system can fill up with either. Some measure of the composition of at least one of the reactants is required to make this system work. Compositions in the reactor or the recycle stream indicate an imbalance in the amounts of reactants being fed and being consumed. If direct composition measurement is not possible, inferential methods using multiple trays temperatures in the column are sometimes feasible (Yu and Luyben, 1984).

Two-column case. If the relative volatility of the product C is intermediate between the two reactants, a two-column distillation system is typically used. Either the light-out-first (LOF), direct separation sequence, or the heavy-out-first (HOF), indirect separation sequence, can be used. The former is more common because the lightest component only has to be taken overhead once (in the first column) and not twice (as would be the case in the HOF configuration). However there are processes in which the HOF is preferred because it sometimes has the advantage of reducing the exposure of temperature-sensitive components to high base temperatures.

Assuming we use cooling water in both column condensers, the pressure in the first column of the LOF system (with mostly A) will be higher than the pressure in the second column (with mostly B). The base of the first column contains a mixture of B and C , and the base temperature can sometimes be too high for thermally sensitive components. Using the HOF system gives a lower pressure in the first column, and even though the base is now mostly B , the base temperature is sometimes lower than in the LOF system. In addition, component B is being held at high temperature in the base of both columns in the LOF system, and this may be undesirable if B is thermally sensitive.

Whatever separation sequence is chosen, the control structures that work well are quite similar. We will choose the HOF system to illustrate this type of process. Figure 2.13 gives a sketch of a ternary process with two recycle streams. The heaviest component B is recycled back to the reactor from the base of the first column. The lightest component A is recycled back to the reactor from the top of the second column.

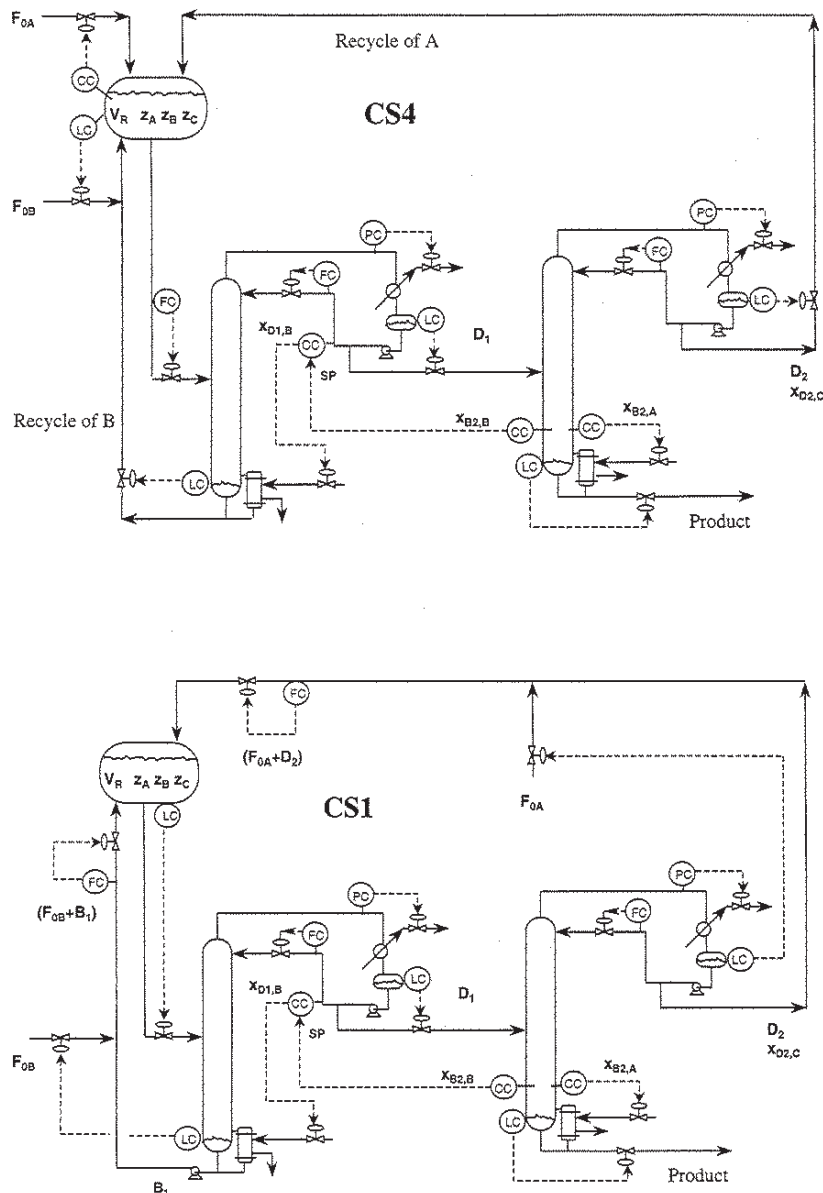


Figure 2.13 Ternary process flowsheet with incomplete conversion and two recycle streams (heavy-out-first sequence). (a) Control structure CS4: reactor composition and level control (workable); (b) control structure CS1: reactant makeup control based on component inventories (workable).

Figure 2.13a and b shows two control structures that work (CS4 and CS1). Both of these provide a mechanism for adjusting the fresh feed reactant flowrates so that the overall stoichiometry can be satisfied. In CS4 this is accomplished by measuring reactor composition. In CS1 it is accomplished by deducing the amounts of the reactants in the process from two levels in the two recycle loops.

In both strategies the control of the separation section is similar:

1. In both columns, reflux flows are fixed (or ratioed to feedrates) and pressures are controlled by condenser cooling.
2. The impurity of A ($x_{B2,A}$) in the product stream B_2 from the second column is controlled by vapor boilup in the second column.
3. The impurity of B ($x_{B2,B}$) in the product stream B_2 from the second column is controlled by vapor boilup in the first column through a composition-composition cascade control system. Any B that goes overhead in the first column comes out the bottom of the second column. So the first column must be operated to prevent B from going overhead. The impurity of B in the first column distillate ($x_{D1,B}$) is controlled by a composition controller that manipulates the vapor boilup in the first column. The setpoint of this composition controller is changed by a second composition controller looking at the impurity of B in the product stream ($x_{B2,B}$).

Control structure CS4 (Fig. 2.13a) controls reactor effluent flow, brings fresh A in to hold reactor composition z_A , and brings fresh B in to control reactor level. In both columns, the base levels are controlled by manipulating bottoms flowrates and the reflux drum levels are controlled by manipulating distillate flowrates.

Control structure CS1 (Fig. 2.13b) controls the flowrates of the two total light and heavy recycle streams; i.e., the sum of the fresh feed and recycle of A ($F_{0A} + D_2$) is flow-controlled and the sum of the fresh feed and recycle of B ($F_{0B} + B_1$) is flow-controlled. The fresh reactant A feed controls the level in the reflux drum of the second column, which reflects this component's inventory within the process. Similarly, the fresh reactant B feed controls the level in the base of the first column.

Both of these control structures have the slight disadvantage of lacking a single direct handle to set production rate, i.e., a one-to-one relationship with product flow. Desired throughput must be achieved by changing the setpoint of the reactor concentration controller, the reactor level controller, the reactor effluent flow controller, and/or the recycle flow controllers (one or both). Structure CS4 has another disadvantage since it requires a composition measurement, which can be very expensive and unreliable in many systems.

We could easily propose many other control structures for this pro-

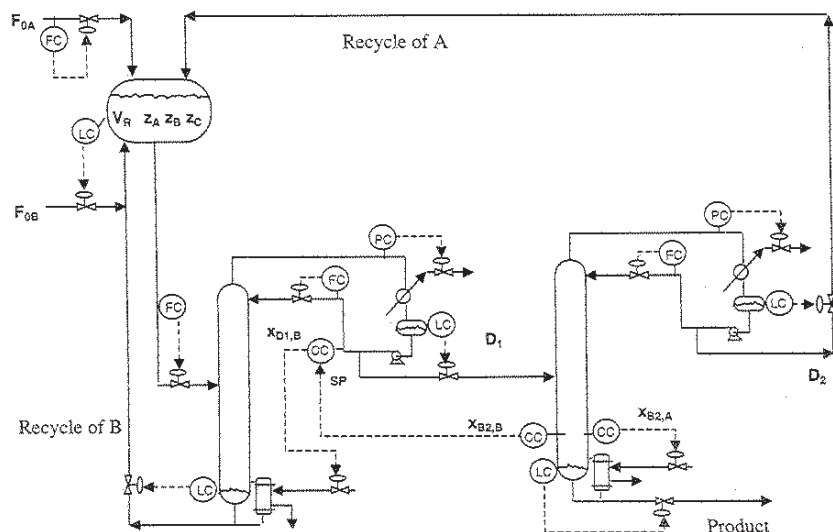


Figure 2.14 Ternary process flowsheet with incomplete conversion and two recycle streams (heavy-out-first sequence): control structure CS2 with fixed flow of one reactant (unworkable).

cess, but most *do not work* in these types of systems. Schemes where one of the reactant fresh feeds is simply flow-controlled into the process do not work unless the per-pass conversion of this limiting component is quite high; i.e., the concentration of this component in the reactor effluent is very small. An analysis of this problem is given in Luyben et al. (1996).

For example, consider the control system shown in Figure 2.14. Here there is a direct handle on production: the flow of fresh A into the system. However, this scheme does not work. Figure 2.15a illustrates that the system is able to handle a very small (2 percent) change in fresh feed flow. But if the change in fresh feed flow is increased to 5 percent, the system fills up with A and shuts down after 150 hours (see Fig. 2.15b). If the increase is +10 percent (Fig. 2.15c), the system shuts down in 70 hours. Thus this control structure can handle only very small disturbances. The imbalance in chemical components and the long time period over which the problem occurs highlight the importance of these phenomena in the plantwide control problem.

2.7.3 Stability analysis

To gain some understanding of what is happening in the results shown in Fig. 2.15 and to explain why the process shuts down, it is useful to

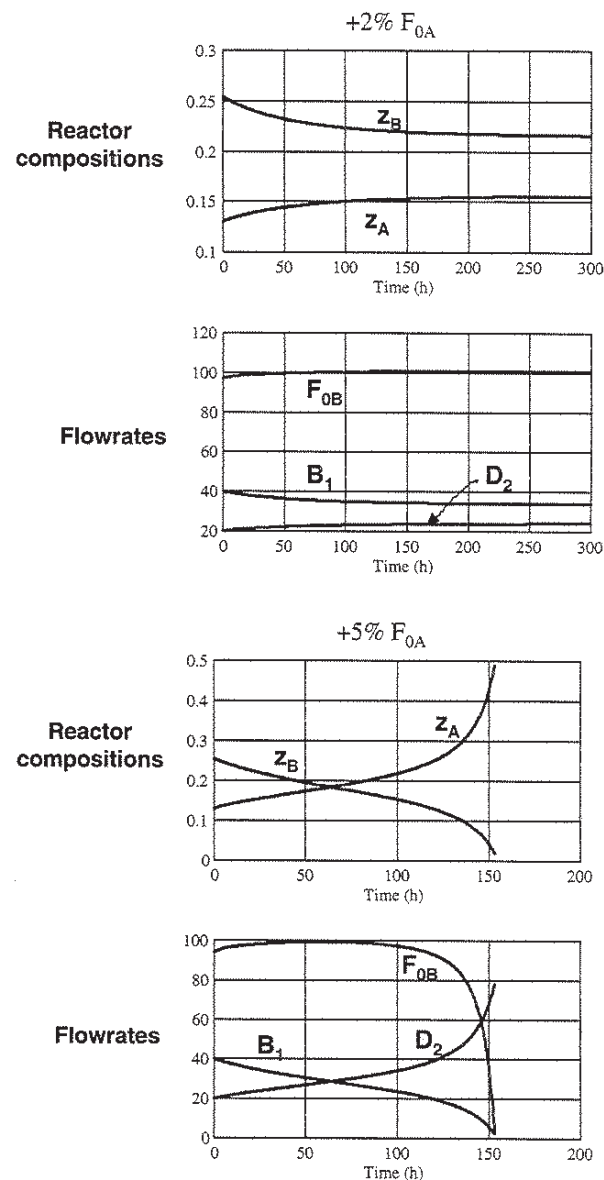


Figure 2.15 Dynamic response of ternary process with CS2 for change in fixed reactant feed rate. (a) 2 percent increase; (b) 5 percent increase; (c) 10 percent increase.

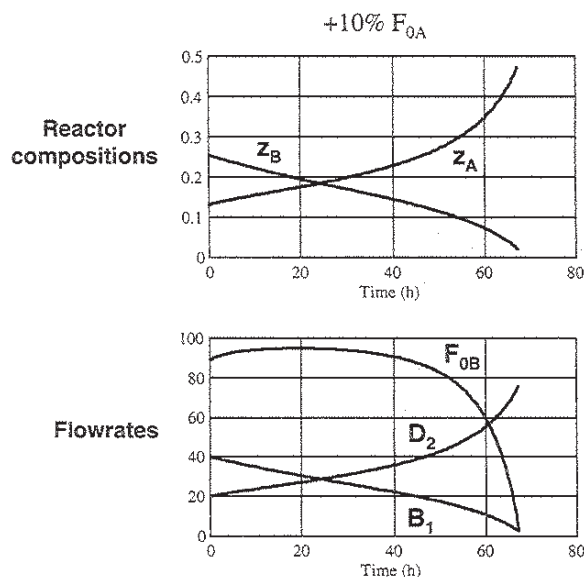


Figure 2.15 (Continued)

look at a simple model of the process and to see what such a model predicts concerning the stability of the process. The results in Fig. 2.15 show that the disturbance in F_{0A} drives the reactor compositions into a region where z_A becomes larger than z_B and then a shutdown eventually occurs after several hours.

Let us derive a dynamic model of the process with control structure CS2 included. A rigorous model of the reactor and the two distillation columns would be quite complex and of very high order. Because the dynamics of the liquid-phase reactor are much slower than the dynamics of the separation section in this process, we can develop a simple second-order model by assuming the separation section dynamics are instantaneous. Thus the separation section is always at steady state and is achieving its specified performance, i.e., product and recycle purities are at their setpoints. Given a flowrate F and the composition z_A/z_B of the reactor effluent stream, the flowrates of the light and heavy recycle streams D_2 and B_1 can be calculated from the algebraic equations

$$D_2 = \frac{Fz_A - A_{\text{loss}}}{x_{D2,A}} \quad (2.8)$$

$$B_1 = \frac{Fz_B - B_{\text{loss}}}{x_{B1,B}} \quad (2.9)$$

where $A_{\text{loss}}, B_{\text{loss}}$ = small molar flowrates of components A and B in product stream B_2 leaving base of second column (assumed to be constant)

$x_{D2,A}, x_{B1,B}$ = purities of recycle streams

Perfect reactor level control is assumed. The reactor effluent flowrate F is fixed in control structure CS2. The two state variables of the system are the two reactor compositions z_A and z_B . The two nonlinear ordinary differential equations describing the system are

$$V_R \frac{dz_A}{dt} = F_{0A} + D_2 x_{D2,A} - Fz_A - V_R k z_A z_B \quad (2.10)$$

$$V_R \frac{dz_B}{dt} = F_{0B} + B_1 x_{B1,B} - Fz_B - V_R k z_A z_B \quad (2.11)$$

At any point in time we know z_A and z_B . The variables $F, F_{0A}, k, V_R, A_{\text{loss}}, B_{\text{loss}}, x_{D2,A}$, and $x_{B1,B}$ are constant. At each point in time Eqs. (2.8) and (2.9) can be used to find the recycle flowrates. A total molar balance around the reactor can be used to calculate the makeup flowrate of component B, F_{0B} . Remember that the reaction is $A + B \rightarrow C$, so moles are not conserved.

$$F_{0A} + F_{0B} + D_2 + B_1 = F + V_R k z_A z_B \quad (2.12)$$

The two nonlinear ordinary differential equations can be linearized around the steady-state values of the reactor compositions \bar{z}_A and \bar{z}_B . Laplace transforming gives the characteristic equation of the system. It is important to remember that we are looking at the closed-loop system with control structure CS2 in place. Therefore Eq. (2.13) is the closed-loop characteristic equation of the process:

$$s^2 + s \left[k\bar{z}_B + \frac{F}{V_R x_{B1,B}} \right] + \frac{kF}{V_R} \left[\frac{\bar{z}_B}{x_{B1,B}} - \frac{\bar{z}_A}{x_{D2,A}} \right] = 0 \quad (2.13)$$

Thus the linear analysis predicts that the system will be closed-loop unstable when

$$\frac{\bar{z}_B}{x_{B1,B}} < \frac{\bar{z}_A}{x_{D2,A}} \quad (2.14)$$

If the two recycle purities are about the same ($x_{B1,B} \cong x_{D2,A}$), which is the case in the numerical example considered earlier in the chapter, the linear analysis predicts that instability will occur when z_A is bigger than z_B . This is exactly what we observed in Fig. 2.15.

The physical reason for this instability is the lack of some mechanism

in the process or in the control structure to ensure that the A and B component balances are satisfied in this integrating plantwide process. Both reactant components are prevented from leaving the system by the impurity controllers that are looking at the product stream. Thus essentially all of the reactants fed into the system must be consumed by chemical reaction. And the stoichiometry must be satisfied down to the last molecule: every mole of A requires exactly one mole of B to react with. The flowrates of the fresh feed cannot be controlled in an open-loop fashion anywhere nearly accurately enough to match the molecules of the two reactants exactly. This is why we need some information about the amounts of the two components in the system. This knowledge can be used in a feedback control system to make some adjustments so that the component in excess does not continue to build up in the system.

2.7.4 Modification of CS2

Both of the control structures discussed in Sec. 2.7.2 (CS1 and CS4) work because they detect the inventories of the reactant components A and B in the system and bring in fresh feed streams to balance the consumption of the two components. Structure CS1 does this by using the liquid level in the reflux drum of the second column as an indicator of the amount of A in the system and the liquid level in the base of the first column as an indicator of the amount of B in the system. Structure CS4 uses a composition analyzer to measure directly the concentration of one of the reactants in the reactor. But both of these structures lack a direct handle on production rate.

Control structure CS2 has such a direct handle, but this structure does not work. However, a modification can be made to CS2 that will make it work. The basic idea is to recognize that the separation section acts like an on-line analyzer. Any component B in the reactor effluent gets recycled in B_1 . Any component A in the reactor effluent gets recycled in D_2 . Therefore, the flowrates of these two streams give a direct indication of the amounts of the two reactants in the system.

Figure 2.16 shows a control scheme in which the ratio of the two recycle flowrates is controlled by adjusting the flowrate of the reactor effluent. The dynamics of the separation system must be considered because a change in the amount of A in the reactor effluent has to work its way through two columns before showing up as a change in the flowrate of D_2 . Thus a lag is added to the measurement of B_1 before it is used to calculate the ratio. This control structure works.

In this modified CS2, the feedback adjustment that is made to adjust for any imbalance in the amounts of the two reactants in the system is a change in the reactor effluent flowrate to achieve a constant ratio

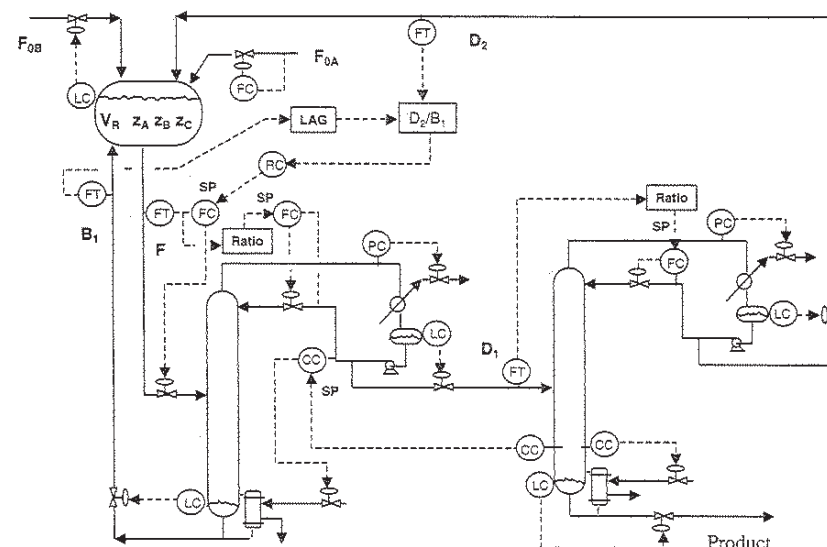


Figure 2.16 Ternary process flowsheet with incomplete conversion and two recycle streams (heavy-out-first sequence): control structure CS2C using separation as analyzer for control of D_2/B_1 ratio (workable).

of the two recycle flowrates. This works because these flowrates give good indications of the concentrations of the two reactants in the reactor. The two columns act like composition analyzers, separating the A and B components from the product C .

In the numerical case studied in this chapter, we considered a liquid-phase reactor with dynamics that were slower than the dynamics of the separation system. Suppose we have a process with a vapor-phase reactor whose dynamics are much faster than those of the separation section. Will the modified CS2 control structure work in this process?

Luyben et al. (1996) explored this question in detail by developing a rigorous simulation of such a process. Their results demonstrate that the proposed control structure does provide effective control for processes with fast reactor dynamics. The time constant of the separation section is about 30 minutes. The reactor time constant was reduced to 3 minutes, and control was still good.

2.7.5 Reactor composition trade-offs

As discussed earlier, if the concentration of A (or B) in the reactor is essentially zero, we can flow-control the fresh feed of A (or B) into the system, and large disturbances can be handled. In the numerical case

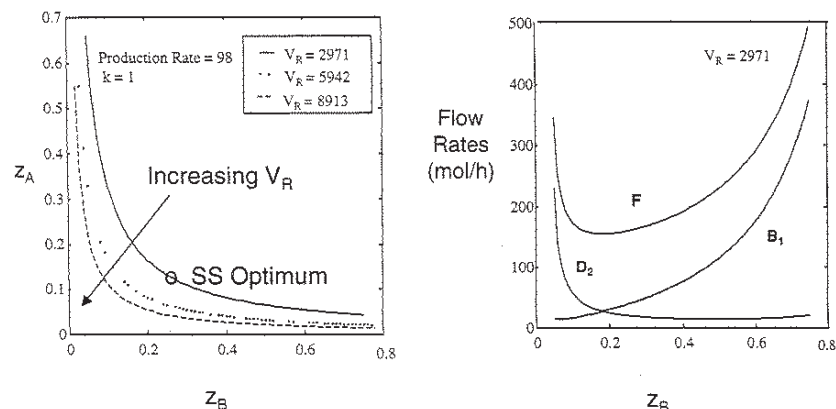


Figure 2.17 Steady-state design for ternary process with incomplete conversion and two recycle streams (heavy-out-first sequence).

presented in the previous section, the steady-state economic design of the process yielded reactor compositions that are $z_A = 0.15$ mole fraction and $z_B = 0.25$ mole fraction. It is cheaper to recycle B than A because B comes out the bottom of the first column and does not have to be vaporized. Component A , on the other hand, must be vaporized twice as it is taken overhead in both columns. Therefore the steady-state separation design favors smaller z_A and larger z_B . But remember that if reactor temperature and holdup are constant (fixed k and V_R), the product of the two concentrations must be fixed to achieve a given production rate of C .

Figure 2.17 illustrates that we must lie somewhere on the hyperbolic line in the $z_A - z_B$ plane. At any position on one of the constant reactor volume lines, the production rate is constant. The concentrations fed to the separation section vary with our choice of location on this curve. For large z_A and small z_B , the recycle of A (D_2) is large. For large z_B and small z_A , the recycle of B (B_1) is large.

Since we are dealing with the product of the two reactant concentrations, making them approximately equal is the best way to minimize reactor holdup. Thus steady-state reactor design favors compositions that are somewhat similar. From a dynamic viewpoint, the system can handle disturbances more easily if the concentrations of the two reactants are very different (very small z_A and large z_B). We saw an indication of this in the ternary process considered earlier. Control structure CS2 worked when the concentration of the limiting reactant was very low, but failed when the concentration of the limiting reactant was in the 0.15 mole fraction region.

So this simple process provides another nice example of the very

common situation where a conflict exists between steady-state economic design and dynamic controllability.

2.8 Conclusion

In this chapter we have looked at the steady-state and dynamic implications of integrated processes with recycle (as compared with units connected only in series). From a dynamic standpoint we found that recycles provide positive feedback that alters the overall time constant of the process. From a steady-state viewpoint, recycles introduce the possibility of the snowball effect, where a small change in throughput or feed composition can produce a large change in recycle flowrates. These features restrict the set of workable control structures for an integrated process. Several simple processes were used to illustrate the interaction between the reaction and separation sections. The generic conclusion was to control dominant variables using local manipulators in the reaction section. We then achieve production rate changes by manipulating the setpoints so that disturbances to the separation section are minimized, thereby reducing product quality variability. Another point that was highlighted involved the need for the control strategy to account for the chemical component balances, i.e., to keep track of the inventory of components within the system.

2.9 References

- Belanger, P. W., and Luyben, W. L. "Design of Low-Frequency Compensators for Improvement of Plantwide Regulatory Performance," *Ind. Eng. Chem. Res.*, **36**, 5339-5347 (1997).
- Douglas, J. M. *Conceptual Design of Chemical Processes*, New York: McGraw-Hill (1988).
- Luyben, W. L. "Dynamics and Control of Recycle Systems: 2. Comparison of Alternative Process Designs," *Ind. Eng. Chem. Res.*, **32**, 476-486 (1993).
- Luyben, M. L., Tyreus, B. D., Luyben, W. L. "Analysis of Control Structures for Reaction/Separation/Recycle Processes with Second-Order Reactions," *Ind. Eng. Chem. Res.*, **35**, 758-771 (1996).
- Morud, J., and Skogestad, S. "Dynamic Behavior of Integrated Plants," *J. Proc. Cont.*, **6**, 145-156 (1996).
- Terrill, D. L., and Douglas, J. M. "Heat-Exchanger Network Analysis. 1. Optimization," *Ind. Eng. Chem. Res.*, **26**, 685-691 (1987).
- Tyreus, B. D., and Luyben, W. L. "Dynamic and Control of Recycle Systems: 4. Ternary Systems with One or Two Recycle Streams," *Ind. Eng. Chem. Res.*, **32**, 1154-1162 (1993).
- Yu, C. C., and Luyben, W. L. "Use of Multiple Temperatures for the Control of Multicomponent Distillation Columns," *Ind. Eng. Chem. Proc. Des. Dev.*, **23**, 590-597 (1984).

Plantwide Control Design Procedure

3.1 Introduction

In an industrial environment, a plant's control strategy should be simple enough, at least conceptually, so that everyone from the operator to the plant manager can understand how it works. Our governing philosophy is *it is always best to utilize the simplest control system that will achieve the desired objectives*. The more complex the process, the more desirable it is to have a simple control strategy. This view differs radically from much of the current academic thinking about process control, which suggests that a complex process demands complex control. Our viewpoint is a result of many years of working on practical plant control problems, where it is important to be able to identify whether an operating problem has its source in the process or in the control system.

The goals for an effective plantwide process control system include (1) safe and smooth process operation; (2) tight control of product quality in the face of disturbances; (3) avoidance of unsafe process conditions; (4) a control system run in automatic, not manual, requiring minimal operator attention; (5) rapid rate and product quality transitions; and (6) zero unexpected environmental releases.

As illustrated in the previous chapter, the need for a plantwide control perspective arises from three important features of integrated processes: the effects of material recycle, of chemical component inventories, and of energy integration. We have shown several control strategies that highlight important general issues. However, we did not describe how we arrived at these strategies, and many of our choices may seem mysterious at this point. Why, for instance, did we choose

to use fresh liquid reactant feed streams in the control of liquid inventories? What prompted us to have a reactor composition analyzer? Why were we concerned with a single direct handle to set production rate?

In this chapter we outline the nine basic steps of a general heuristic plantwide control design procedure (Luyben et al., 1997). After some preliminary discussion of the fundamentals on which this procedure is based, we outline each step in general terms. We also summarize our justification for the sequence of steps. The method is illustrated in applications to four industrial process examples in Part 3.

The procedure essentially decomposes the plantwide control problem into various levels. It forces us to focus on the unique features and issues associated with a control strategy for an entire plant. We highlighted some of these questions in Chap. 1 in discussing the HDA process. How do we manage energy? How is production rate controlled? How do we control product quality? How do we determine the amounts of fresh reactants to add?

Our plantwide control design procedure (Fig. 3.1) satisfies the two fundamental chemical engineering principles, namely 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 and that the propagation of thermal disturbances is prevented. In Steps 4 and 5 we

1. Establish Control Objectives
2. Determine Control Degrees of Freedom
3. Establish Energy Management System
4. Set Production Rate
5. Control Product Quality and Handle Safety,
Environmental, and Operational Constraints
6. Fix a Flow in Every Recycle Loop and Control Inventories
(Pressures and Liquid Levels)
7. Check Component Balances
8. Control Individual Unit Operations
9. Optimize Economics and Improve Dynamic Controllability

Figure 3.1 Nine steps of plantwide control design procedure.

satisfy the key 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.

The plantwide control design procedure presented here was developed after many years of work and research in the fields of process control and process design. Research efforts by a number of people in industry and at universities have contributed essential ideas and concepts. We have assembled, analyzed, and processed this prior work to reach a logical, coherent, step-by-step procedure. We want to acknowledge these previous contributions and state that we are indeed fortunate to stand upon the shoulders of many giants. Listed below are some of the fundamental concepts and techniques that form the basis of the procedure.

3.2 Basic Concepts of Plantwide Control

3.2.1 Buckley basics

Page Buckley (1964), a true pioneer with DuPont in the field of process control, was the first to suggest the idea of separating the plantwide control problem into two parts: material balance control and product quality control. He suggested looking first at the flow of material through the system. A logical arrangement of level and pressure control loops is established, using the flowrates of the liquid and gas process streams. No controller tuning or inventory sizing is done at this step. The idea is to establish the inventory control system by setting up this "hydraulic" control structure as the first step.

He then proposed establishing the product-quality control loops by choosing appropriate manipulated variables. The time constants of the closed-loop product-quality loops are estimated. We try to make these as small as possible so that good, tight control is achieved, but stability constraints impose limitations on the achievable performance.

Then the inventory loops are revisited. The liquid holdups in surge volumes are calculated so that the time constants of the liquid level loops (using proportional-only controllers) are a factor of 10 larger than the product-quality time constants. This separation in time constants permits independent tuning of the material-balance loops and the prod-

uct-quality loops. Note that most level controllers should be proportional-only (P) to achieve flow smoothing.

3.2.2 Douglas doctrines

Jim Douglas (1988) of the University of Massachusetts has devised a hierarchical approach to the conceptual design of process flowsheets. Although he primarily considers the steady-state aspects of process design, he has developed several useful concepts that have control structure implications.

Douglas points out that in the typical chemical plant the costs of raw materials and the value of the products are usually much greater than the costs of capital and energy. This leads to the two *Douglas doctrines*:

1. Minimize losses of reactants and products.
2. Maximize flowrates through gas recycle systems.

The first idea implies that we need tight control of stream compositions exiting the process to avoid losses of reactants and products. The second rests on the principle that yield is worth more than energy. Recycles are used to improve yields in many processes, as was discussed in Chap. 2. The economics of improving yields (obtaining more desired products from the same raw materials) usually outweigh the additional energy cost of driving the recycle gas compressor.

The control structure implication is that we do not attempt to regulate the gas recycle flow and we do not worry about what we control with its manipulation. We simply maximize its flow. This removes one control degree of freedom and simplifies the control problem.

3.2.3 Downs drill

Jim Downs (1992) of Eastman Chemical Company has insightfully pointed out the importance of looking at the chemical component balances around the entire plant and checking to see that the control structure handles these component balances effectively. The concepts of overall component balances go back to our first course in chemical engineering, where we learned how to apply mass and energy balances to any system, microscopic or macroscopic. We did these balances for individual unit operations, for sections of a plant, and for entire processes.

But somehow these basics are often forgotten or overlooked in the complex and intricate project required to develop a steady-state design for a large chemical plant and specify its control structure. Often the design job is broken up into pieces. One person will design the reactor and its control system and someone else will design the separation

section and its control system. The task sometimes falls through the cracks to ensure that these two sections operate effectively when coupled together. Thus it is important that we perform the *Downs drill*.

We must ensure that all components (reactants, products, and inerts) have a way to leave or be consumed within the process. The consideration of inerts is seldom overlooked. Heavy inerts can leave the system in the bottoms product from a distillation column. Light inerts can be purged from a gas recycle stream or from a partial condenser on a column. Intermediate inerts must also be removed in some way, for example in sidestream purges or separate distillation columns.

Most of the problems occur in the consideration of reactants, particularly when several chemical species are involved. All of the reactants fed into the system must either be consumed via reaction or leave the plant as impurities in the exiting streams. Since we usually want to minimize raw material costs and maintain high-purity products, most of the reactants fed into the process must be chewed up in the reactions. And the stoichiometry must be satisfied *down to the last molecule*.

Chemical plants often act as pure integrators in terms of reactants. This is due to the fact that we prevent reactants from leaving the process through composition controls in the separation section. Any imbalance in the number of moles of reactants involved in the reactions, no matter how slight, will result in the process gradually filling up with the reactant component that is in excess. The ternary system considered in Chap. 2 illustrated this effect. There must be a way to adjust the fresh feed flowrates so that exactly the right amounts of the two reactants are fed in.

3.2.4 Luyben laws

Three laws have been developed as a result of a number of case studies of many types of systems:

1. A stream somewhere in all recycle loops should be flow controlled. This is to prevent the snowball effect and was discussed in Chap. 2.
2. A fresh reactant feed stream cannot be flow-controlled unless there is essentially complete one-pass conversion of one of the reactants. This law applies to systems with reaction types such as $A + B \rightarrow \text{products}$ and was discussed in Chap. 2. In systems with consecutive reactions such as $A + B \rightarrow M + C$ and $M + B \rightarrow D + C$, the fresh feeds can be flow-controlled into the system because any imbalance in the ratios of reactants is accommodated by a shift in the amounts of the two products (M and D) that are generated. An excess of A will result in the production of more M and less D . An excess of B results in the production of more D and less M .

3. If the final product from a process comes out the top of a distillation column, the column feed should be liquid. If the final product comes out the bottom of a column, the feed to the column should be vapor (Cantrell et al., 1995). Changes in feed flowrate or feed composition have less of a dynamic effect on distillate composition than they do on bottoms composition if the feed is saturated liquid. The reverse is true if the feed is saturated vapor: bottoms is less affected than distillate. If our primary goal is to achieve tight product quality control, the basic column design should consider the dynamic implications of feed thermal conditions. Even if steady-state economics favor a liquid feed stream, the profitability of an operating plant with a product leaving the bottom of a column may be much better if the feed to the column is vaporized. This is another example of the potential conflict between steady-state economic design and dynamic controllability.

3.2.5 Richardson rule

Bob Richardson of Union Carbide suggested the heuristic that the largest stream should be selected to control the liquid level in a vessel. This makes good sense because it provides more *muscle* to achieve the desired control objective. An analogy is that it is much easier to maneuver a large barge with a tugboat than with a life raft. We often use the expression that you can't make a garbage truck drive like a Ferrari. But this is not necessarily true. If you put a 2000-hp engine in the garbage truck (and redesigned the center of gravity), you could make it handle just like a sports car. The point is that the bigger the handle you have to affect a process, the better you can control it. This is why there are often fundamental conflicts between steady-state design and dynamic controllability.

3.2.6 Shinsky schemes

Greg Shinsky (1988), over the course of a long and productive career at Foxboro, has proposed a number of "advanced control" structures that permit improvements in dynamic performance. These schemes are not only effective, but they are simple to implement in basic control instrumentation. Liberal use should be made of ratio control, cascade control, override control, and valve-position (optimizing) control. These strategies are covered in most basic process control textbooks.

3.2.7 Tyreus tuning

One of the vital steps in developing a plantwide control system, once both the process and the control structure have been specified, is to

determine the algorithm to be used for each controller (P, PI, or PID) and to tune each controller. We strongly recommend the use of P-only controllers for liquid levels (even in some liquid reactor applications). Tuning of a P controller is usually trivial: set the controller gain equal to 1.67. This will have the valve wide open when the level is at 80 percent and the valve shut when the level is at 20 percent (assuming the stream flowing out of the vessel is manipulated to control liquid level; if the level is controlled by the inflowing stream the action of the controller is reverse instead of direct).

For other control loops, we suggest the use of PI controllers. The relay-feedback test is a simple and fast way to obtain the ultimate gain (K_u) and ultimate period (P_u). Then either the Ziegler-Nichols settings (for very tight control with a closed-loop damping coefficient of about 0.1) or the Tyreus-Luyben (1992) settings (for more conservative loops where a closed-loop damping coefficient of 0.4 is more appropriate) can be used:

$$K_{ZN} = K_u/2.2 \quad \tau_{ZN} = P_u/1.2$$

$$K_{TL} = K_u/3.2 \quad \tau_{TL} = 2.2P_u$$

The use of PID controllers should be restricted to those loops where two criteria are both satisfied: the controlled variable should have a very large signal-to-noise ratio and tight dynamic control is really essential from a feedback control stability perspective. The classical example of the latter is temperature control in an irreversible exothermic chemical reactor (see Chap. 4).

3.3 Steps of Plantwide Process Control Design Procedure

In this section we discuss each step of the design procedure in detail.

Step 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 control objectives lead to different control structures. There is an old Persian saying "If you don't know where you are going, any road will get you there!" This is certainly true in plantwide control. The "best" control structure for a plant depends upon the design and control criteria established.

These objectives include reactor and separation yields, product qual-

ity specifications, product grades and demand determination, environmental restrictions, and the range of safe operating conditions.

Step 2: Determine control degrees of freedom

Count the number of control valves available.

This is the number of degrees of freedom for control, i.e., the number of variables that can be controlled to setpoint. The valves must be legitimate (flow through a liquid-filled line can be regulated by only one control valve). The placement of these control valves can sometimes be made to improve dynamic performance, but often there is no choice in their location.

Most of these valves will be used to achieve basic regulatory control of the process: (1) set production rate, (2) maintain gas and liquid inventories, (3) control product qualities, and (4) avoid safety and environmental constraints. Any valves that remain after these vital tasks have been accomplished can be utilized to enhance steady-state economic objectives or dynamic controllability (e.g., minimize energy consumption, maximize yield, or reject disturbances).

During the course of the subsequent steps, we may find that we lack suitable manipulators to achieve the desired economic control objectives. Then we must change the process design to obtain additional handles. For example, we may need to add bypass lines around heat exchangers and include auxiliary heat exchangers.

Step 3: Establish energy management system

Make sure that energy disturbances do not propagate throughout the process by transferring the variability to the plant utility 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 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 flowrate of the limiting fresh reactant to the flowrate of a recycle stream acting as a thermal sink). More details of reactor control are discussed in Chap. 4.

Heat transfer between process streams can create significant interaction. In the case of 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.

For example, suppose the temperature of a stream being fed to a distillation column is controlled by manipulating steam flowrate to a feed preheater. And suppose the stream leaving the preheater is partially vaporized. Small changes in composition can result in very large changes in the fraction of the stream that is vaporized (for the same pressure and temperature). The resulting variations in the liquid and vapor rates in the distillation column can produce severe upsets.

Heat integration of a distillation column with other columns or with reactors is widely used in chemical plants to reduce energy consumption. While these designs look great in terms of steady-state economics, they 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. Chapter 5 deals with heat exchanger systems.

Step 4: Set production rate

Establish the variables that dominate the productivity of the reactor and determine the most appropriate manipulator to control production rate.

Throughput changes can be achieved only 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).

Our first choice for setting production rate should be to alter one of these variables in the reactor. The variable we select must be dominant for the reactor. Dominant reactor variables always have significant effects on reactor performance. For example, temperature is often a

dominant reactor variable. In irreversible reactions, specific rates increase exponentially with temperature. As long as reaction rates are not limited by low reactant concentrations, we can *increase* temperature to increase production rate in the plant. In reversible exothermic reactions, where the equilibrium constant decreases with increasing temperature, reactor temperature may still be a dominant variable. If the reactor is large enough to reach chemical equilibrium at the exit, we can *decrease* reactor temperature to increase production.

There are situations where reactor temperature is not a dominant variable or cannot be changed for safety or yield reasons. In these cases, we must find another dominant variable, such as the concentration of the limiting reactant, flowrate of initiator or catalyst to the reactor, reactor residence time, reactor pressure, or agitation rate.

Once we identify the dominant variables, we must also identify the manipulators (control valves) that are most suitable to control them. The manipulators are used in feedback control loops to hold the dominant variables at setpoint. The setpoints are then adjusted to achieve the desired production rate, in addition to satisfying other economic control objectives.

Whatever variable we choose, we would like it to provide smooth and stable production rate transitions and to reject 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.

When the setpoint of a dominant variable is used to establish plant production rate, the control strategy must ensure that the right amounts of fresh reactants are brought into the process. This is often accomplished through fresh reactant makeup control based upon liquid levels or gas pressures that reflect component inventories. We must keep these ideas in mind when we reach Steps 6 and 7.

However, design constraints may limit our ability to exercise this strategy concerning fresh reactant makeup. An upstream process may establish the reactant feed flow sent to the plant. A downstream process may require on-demand production, which fixes the product flowrate from the plant. In these cases, the development of the control strategy becomes more complex because we must somehow adjust the setpoint of the dominant variable on the basis of the production rate that has been specified externally. We must balance production rate with what has been specified externally. This cannot be done in an open-loop sense. Feedback of information about actual internal plant conditions is required to determine the accumulation or depletion of the reactant components. This concept was nicely illustrated by the control strategy in Fig. 2.16. In that scheme we fixed externally the flow of fresh reactant A feed. Also, we used reactor residence time (via the effluent flowrate)

as the controlled dominant variable. Feedback information (internal reactant composition information) is provided to this controller by the ratio of the two recycle stream flows.

Step 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 deadtimes 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).

It should be noted that establishing the product-quality loops first, before the material balance control structure, is a fundamental difference between our plantwide control design procedure and Buckley's procedure. Since product quality considerations have become more important in recent years, this shift in emphasis follows naturally.

The magnitudes of various flowrates 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 flowrate. However, in columns with a large boilup ratio and small bottoms flowrate, these loops should be reversed because boilup has a larger effect on base level than bottoms flow (Richardson rule). 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.

Step 6: Fix a flow in every recycle loop and control inventories (pressures and levels)

Fix a flow in every recycle loop and then select the best manipulated variables to control inventories.

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, as illustrated by the simple process examples in Chap. 2. Steady-state and dynamic benefits result from this flow control strategy. From a steady-state viewpoint, the plant's separation

section is not forced to operate at significantly different load conditions, which could lead to turndown or flooding problems.

From a dynamic viewpoint, whenever all flows in a recycle loop are set by level controllers, wide dynamic excursions can occur in these flows because the total system inventory is not regulated. The control system is attempting to control the inventory in each individual vessel by changing the flowrate to its downstream neighbor. In a recycle loop, all level controllers see load disturbances coming from the upstream unit. This causes the flowrate disturbances to propagate around the recycle loop. Thus any disturbance that tends to increase the total inventory in the process (such as an increase in the fresh feed flowrate) will produce large increases in all flowrates around the recycle loop.

Fixing a flowrate in a recycle stream does not conflict with our discussion of picking a dominant reactor variable for production rate control in Step 4. Flow controlling a stream somewhere in all recycle loops is an important simple part of any plantwide control strategy.

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

Once we have fixed a flow in each recycle loop, we then determine what valve should be used to control each inventory variable. This is the material balance step in the Buckley procedure. Inventories 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 (Richardson rule). Because we have fixed a flow in each recycle loop, our choice of available valves has been reduced for inventory control in some units. Sometimes this actually eliminates the obvious choice for inventory control for that unit. This constraint forces us to look outside the immediate vicinity of the holdup we are considering.

For example, suppose that the distillate flowrate from a distillation column is large compared to the reflux. We normally would use distillate to control level in the reflux drum. But suppose the distillate recycles back to the reactor and so we want to control its flow. What manipulator should we use to control reflux drum level? We could potentially use condenser cooling rate or reboiler heat input. Either choice would have implications on the control strategy for the column, which would ripple through the control strategy for the rest of the plant. This would lead to control schemes that would never be considered if one looked only at the unit operations in isolation.

Inventory may also be controlled with fresh reactant makeup streams as discussed in Step 4. Liquid fresh feed streams may be added to a location where level reflects the amount of that component in the pro-

cess. 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 flowrate disturbances to the downstream separation system. There is nothing necessarily sacred about holding reactor level constant.

Step 7: Check component balances

Identify how chemical components enter, leave, and are generated or consumed in the process.

Component balances can often be quite subtle, but they are particularly important in processes with recycle streams because of their integrating effect. 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 reaction rate in the reactor. The buildup of chemical components in recycle streams must be prevented by keeping track of chemical component inventories (reactants, products, and inerts) inside the system.

We must identify the specific mechanism or control loop to guarantee that there will be no uncontrollable buildup of any chemical component within the process (Downs drill).

What are the methods or loops to ensure that the overall component balances for all chemical species are satisfied at steady state? We can limit their intake, control their reaction, or adjust their outflow from the process.

As we noted in Chap. 2, we can characterize a plant's chemical components into reactants, products, and inerts. We don't want reactant components to leave in the product streams because of the yield loss and the desired product purity specifications. Hence we are limited to the use of two methods: consuming the reactants by reaction or adjusting their fresh feed flow. Product and inert components all must have an exit path from the system. In many systems inerts are removed by purging off a small fraction of the recycle stream. The purge rate is adjusted to control the inert composition in the recycle stream so that an economic balance is maintained between capital and operating costs.

We recommend making a Downs drill table that lists each chemical component, its input, its generation or consumption, and its output. This table should specify how the control system will detect an imbalance in chemical components and what specific action it will take if an imbalance is detected.

Step 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 individual chemical units (Shinsky, 1988). For example, a tubular reactor usually requires control of inlet temperature. High-temperature endothermic reactions typically have a control system to adjust the fuel flowrate 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. We deal with the control of various unit operations in Chaps. 4 through 7.

Step 9: Optimize economics or improve dynamic controllability

Establish the best way to use the remaining control degrees of freedom.

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 utilized either to optimize steady-state economic process performance (e.g., minimize energy, maximize selectivity) or to improve dynamic response.

For example, suppose an exothermic chemical reactor may be cooled with both jacket cooling water and brine (refrigeration) to a reflux condenser. For fast reactor temperature control, manipulation of brine is significantly better than cooling water. However, the utility cost of brine is much higher than cooling water. Hence we would like the control system to provide tight reactor temperature control while minimizing brine usage. This can be achieved with a valve position control strategy. Reactor temperature is controlled by manipulating brine. A valve position controller looks at the position of the brine control valve and slowly adjusts jacket cooling water flow to keep the brine valve approximately 10 to 20 percent open under steady-state operation (Fig. 3.2).

Additional considerations

Certain quantitative measures from linear control theory may help at various steps to assess relationships between the controlled and manipulated variables. These include steady-state process gains, open-loop time constants, singular value decomposition, condition numbers, eigenvalue analysis for stability, etc. These techniques are described in




Figure 3.2 Illustration of valve position control strategy.

detail in most process control textbooks. The plantwide control strategy should ultimately be tested on a nonlinear dynamic model that captures the essential process behavior.

Since the design of a chemical 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.

3.4 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 via 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 a different variable to set production rate. Determining the *best* choice at Step 4 can only be done via 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.

This, then, is a general and straightforward method for tackling the control system design problem for an entire process. Using the procedure as a framework, we should be able to transform an initially complex and seemingly intractable problem into one that can be solved. Before we illustrate the application of the procedure to four industrial processes, we analyze and summarize the control systems for individual unit operations. We also discuss how they fit into the plantwide perspective.

3.5 Conclusion

We have discussed in detail each of the nine steps in our plantwide control design procedure. The first two steps establish the control objectives and control degrees of freedom for the plant. In the third step we discuss how the plantwide energy management problem can be converted to a local unit operation energy management problem by using the plant utility system.

The heart of the plantwide control problem lies in Steps 4 through 7, where we establish how to set production rate, maintain product

quality, prevent excessive changes in recycle flowrates, control inventories, and balance chemical components. These steps demand a plantwide perspective that often leads to control strategies differing significantly from those devised by looking at isolated unit operations.

In Part 3 we illustrate the application of these steps in four industrial processes.

3.6 References

- Buckley, P. S. *Techniques of Process Control*, New York: Wiley (1964).
- Cantrell, J. G., Elliott, T. R., and Luyben, W. L. "Effect of Feed Characteristics on the Controllability of Binary Distillation Columns," *Ind. Eng. Chem. Res.*, **34**, 3027-3036 (1995).
- Douglas, J. M. *Conceptual Design of Chemical Processes*, New York: McGraw-Hill (1988).
- Downs, J. J. "Distillation Control in a Plantwide Control Environment," Chap. 20 in *Practical Distillation Control*, W. L. Luyben (ed.), New York: Van Nostrand Reinhold (1992).
- Luyben, M. L., Tyreus, B. D., Luyben, W. L. "Plantwide Control Design Procedure," *AIChE J.*, **43**, 3161-3174 (1997).
- Luyben, M. L., Tyreus, B. D., Luyben, W. L. "Analysis of Control Structures for Reaction/Separation/Recycle Processes with Second-Order Reactions," *Ind. Eng. Chem. Res.*, **35**, 758-771 (1996).
- Shinskey, F. G. *Process Control Systems*, 3d ed., New York: McGraw-Hill (1988).
- Tyreus, B. D., Luyben, W. L. "Dynamic and Control of Recycle Systems: 4. Ternary Systems with One or Two Recycle Streams," *Ind. Eng. Chem. Res.*, **32**, 1154-1162 (1993).
- Tyreus, B. D., and Luyben, W. L. "Tuning of PI Controllers for Integrator/Deadtime Processes," *Ind. Eng. Chem. Res.*, **31**, 2625-2628 (1992).