To the Graduate Council:

I am submitting herewith a dissertation written by Erin Scott Percell entitled "Designing Control Strategies from a Plant-Wide Perspective." I have examined the final copy of this dissertation for form and content and recommend that it be accepted in partial fulfillment of the requirements for the degree of Doctor of Philosophy, with a major in Chemical Engineering.

Charles F. Moore

Charles F. Moore, Major Professor

We have read this dissertation
and recommend its acceptance:

[signature]

[signature]

Duane D Bruns

Vera Williams

Kenneth E. Kirby

Accepted for the Council:

C W Minkel

Associate Vice Chancellor and
Dean of The Graduate School

Designing Control Strategies from a Plant-Wide Perspective

A Dissertation
Presented for the
Doctor of Philosophy
Degree
The University of Tennessee, Knoxville

Erin Scott Percell
August 1996

Copyright © Erin Scott Percell, 1996
All rights reserved

ACKNOWLEDGEMENTS

I would like to express sincere thanks to Dr. Charlie Moore, my major professor, for his invaluable guidance and insights on this project. I also would like to thank Dr. John Prados, Dr. Ken Kirby, Dr. Duane Bruns, Dr. Don Bogue, and Dr. Vera Williams for their efforts as members of my doctoral committee, and for their help and suggestions in preparing this dissertation. I am especially grateful for the patience and perseverance of my wife Sheri, who has waited longer than she expected to move on to greater things.

I would also like to acknowledge the financial support for this project, which was provided by the Measurement and Control Engineering Center at the University of Tennessee.

ABSTRACT

Chemical manufacturing plants nearly always consist of many manufacturing stages, made up of unit operations such as mixers, reactors, and separators. Each unit operation requires at least a minimum control system to insure proper operation. Typically, these unit control systems are designed separately, if not independently. As higher demands are placed on manufacturing for increased quality, productivity, and efficiency, it is increasingly apparent that more attention needs to be given to the plant-wide context in the design of the control system for each unit operation.

This study presents a straightforward approach to designing control strategies for multi-unit plants. The work follows a plant-wide approach, which emphasizes process knowledge, plant objectives, and management of inventories and variation.

Current steady-state tools for control strategy design do not adequately address the first-level material balance control decisions (selection of throughput manipulator, design of level controls) that must be made early in the process of designing control strategies. This work uses a simple plant module to develop heuristics for the design of these material balance controls.

The current literature also does not describe a systematic procedure for selecting sets of manipulated and controlled variables from large variable sets. This work describes a mathematical procedure for identifying variables that do not contribute to the control of the plant, and for grouping variables that represent essentially the same information.

The material balance heuristics, the variable selection procedure, and existing literature tools were evaluated by using them to design a multi-loop control strategy for the Eastman Challenge Plant. The control strategies suggested by this procedure compare favorably to other published solutions to this test problem.

Table of Contents

List of Tables

TABLE

List of Figures

FIGURE

Chapter 1: Introduction

1.1 The Plant-Wide Control Study

Chemical manufacturing plants usually consist of many manufacturing stages, each made up of unit operations such as mixers, reactors, and separators. Each of these unit operations requires some kind of control system to maintain proper operation. As the complexity of a plant increases and the unit operations become more integrated, the control problem will require a systematic approach to control strategy design.

Much of the process control research in recent years has been directed toward developing new algorithms to find optimum plant conditions or to improve the quality of multivariable control. Less attention has been given to the design of the first-level material and energy balance controls that provide the backbone of more complicated strategies, even though good first-level controls are crucial to the plant's performance. Optimization procedures can save money, but they will be useful only if the plant can be efficiently made to maintain the desired operation; poorly designed material and energy balance controls will limit the effectiveness of even the best high-level algorithms. On the other hand, a well designed first-level control system might not require multivariable or optimization analysis to perform well.

Unfortunately, there is no simple method for designing the best first-level control strategy for any plant. The current approach is often to pick a unit operation, design a control strategy for that unit, and move on to the next one. In contrast, plant-wide control is an effort to design control strategies by considering the whole plant as a single integrated operating unit, not a series of individual operations. The focus of this research is on defining tools that can help sort out the complicated interactions between process units.

An organized approach to plant-wide control strategy design should address the following issues:

- What are the objectives of this plant?
- What are the quality and performance measures for the plant?
- Where are the production bottlenecks and constraints?
- What variable should control the production rate?
- How should bulk inventories be controlled?
- Where are the recycle systems?
- How are component inventories controlled in these recycle systems?
- What are the primary sources of variation?
- How does variation propagate through the plant?
- Can this variation be redirected to less harmful locations?
- Would process design changes improve the operation and control?

The answers to these questions will often suggest the framework for a plant-wide control strategy. A control strategy, regardless of its sophistication, will perform poorly if it is designed with no attention to these issues.

The purpose of the plant-wide control study is to provide mathematical and heuristic tools that can be used to design control strategies for multi-unit plants. This research will investigate existing tools and apply new ones to the analysis of two chemical plants.

1.2 Organization of Dissertation

Chapter 2 of this dissertation is a review of literature related to the design of plant-wide control strategies. The literature review gives particular attention to the steady-state tools used to design control strategies. Several test plants are also reviewed, with a detailed look at efforts to address the Eastman Challenge Plant.

Chapter 3 provides technical details about the two plants that will be examined in this research. This chapter also provides a brief preview of the research work for each plant.

Chapter 4 is a detailed look at a series of simulation studies on Module 2, a fairly simple plant module. The chapter concludes with some recommendations for the design of effective material balance strategies.

Chapters 5-7 address the design of a control strategy for the Eastman Challenge Plant. Chapter 5 uses the results of the Module 2 studies to design the initial material balance controls for this plant. Further progress is not possible without some systematic method of choosing appropriate controlled and manipulated variables, so Chapter 6 describes a procedure for reducing large sets of variables and choosing control sets. Chapter 7 uses this variable set reduction method to generate four candidate control strategies for the Eastman Challenge Plant. Chapter 7 also includes results of simulation studies for these four proposed strategies.

Chapter 8 summarizes the results of the study, including an outline of a systematic approach to control strategy design. This method includes both heuristic and mathematical procedures. Chapter 8 also recommends some directions for future work.

Appendix A is an introduction to some of the concepts of plant-wide control. The appendix includes examples of simple processes for which plant-wide analysis is useful.

Appendices B and D show graphs of the dynamic responses of the control strategies tested on Module 2 and the Eastman Challenge Plant.

Appendix C presents some mathematical details of the variable set reduction methods described conceptually in Chapter 6. The appendix includes some simple numerical examples.

Chapter 2: Literature and Concepts

2.1 Terminology

One of the difficulties in reviewing process control literature is that different authors and researchers within the control community use very different terminology to describe their work. The following terminology will be used throughout this work:

An input is a plant variable that is set by operators. Valve positions, motor speed settings, and flow loop setpoints may be some of the inputs available to the operators of a plant. Not all possible inputs may be used for control purposes. Any input chosen to be used by a controller is a manipulated variable.

A measurement is a plant variable that can be monitored but not directly set. Temperatures, pressures, flow rates, and compositions may be some of the measurements available to monitor the state of a plant. Not all measurements in a plant are associated with control loops. Any measurement that is controlled by a loop is a controlled variable.

PID refers to the standard proportional-integral-derivative control algorithm. Such an algorithm may or may not use all of these three modes; a controller which uses only proportional and integral modes is a special case of a PID controller.

SISO ("single-input, single-output") refers to a control algorithm, such as a PID algorithm, which uses one manipulated variable to control one controlled variable.

M-SISO ("multiple single-input, single-output") refers to control schemes which consist of a series of single loops, each connecting one controlled variable with one manipulated variable.

Multivariable control refers to control schemes which associates controlled and manipulated variables in some way other than the one-to-one pairing of m-SISO controllers; the control algorithm uses calculations with several controlled variables to set each manipulated variable. A full multivariable controller uses every controlled variable to set each manipulated variable. A block multivariable controller divides the variables into sets, with all the controlled variables within a set used to set the manipulated variables in the set.

2.2 Definition of the Control Problem

The problem of control system design can be broken down into at least these six tasks:

- defining the system to be controlled (the "plant");
- defining control objectives for the plant;
- selecting a set of controlled variables whose control will meet the plant objectives;
- selecting a set of manipulated variables that will reject disturbances and maintain the controlled variables at desired values;
- creating connections between the manipulated and controlled variables (choosing between m-SISO and multivariable controllers, pairing the variables, etc.); and
- choosing and tuning the control laws that will set the manipulated variables.

Unfortunately, there is no standard method to undertake any of these steps for large plants. The control engineer must rely on his knowledge of plant behavior plus some set of tools to direct the design process. The approach to control system design will also depend somewhat on the engineer's view of the purpose of a control system.

Garcia and Prett (24) state that the Fundamental Control Problem is to "online update the manipulated variables to satisfy multiple, changing performance criteria based on a process representation which includes a description of the uncertainties." The model, uncertainties, performance criteria, and manipulated variables must be defined for each plant. The differences among the various approaches lie in the assumptions and simplifications made so that solution of the problem is possible using current hardware, software, and mathematical tools.

In their brief survey of control structure design methods, Van de Wal and Jager (96) state that the goal of the design methodology is to "minimize the number of measured and manipulated variables selected and the interconnections between them, subject to the achievement of robust performance." There may be many different ways to satisfy some set of performance criteria, but a control system ought to be as simple as possible to help insure that operators can really understand it, especially if something goes wrong. If all else is equal, operators prefer smaller, less complex control systems.

Moore (62) describes a control system as a means of transforming variation from one process variable to another, more benign, location. Without a functioning control system, disturbances cause variation in the important process variables (flow rates, purities, etc.), while the potential manipulated variables are constant. The function of the control system, then, is to remove variation in the controlled variables by creating variation in the manipulated variables. This view of process variation is particularly useful in a modern business climate which often defines quality in terms of product variability, not just product purity (Downs and Doss [13]). The control systems which make good business sense are those which move variability from critical measurements (such as product quality) into less critical ones (such as utility streams).

2.3 Mathematical Tools for Control System Design

Control engineers often design controls using mathematical descriptions of the plant. One simple mathematical representation is the steady-state gain matrix, which describes the steady-state relationship between the controlled and manipulated variables. Sometimes a complete transfer function model or state-space model is available in addition to the steady-state gains. Control engineers base their control system designs on these simple models; whether those control systems work on the real plant often depends on the model accuracy. Assuming that a reasonably accurate model is available, there are many mathematical tools that can simplify the control design process.

2.3.1 Interaction: Good, Bad, or Indifferent?

The simplest control problems are single-loop problems, for which the entire plant consists of one manipulated variable, one controlled variable, and perhaps some disturbance. The objective for a PID controller is generally to minimize some error measurement. However, there are several different possible error measurements (integral of absolute error, integral of squared error, integral of time-weighted absolute error, etc.), and several different conditions for which this measurement can be minimized (step disturbance, sinusoidal disturbance, setpoint change, etc.). These problems are generally well defined, and standard control texts give some tuning formulas appropriate for several different criteria.

The control problems are much more difficult when the plant consists of more than one controlled and manipulated variable. The characteristics of the best control system become more difficult to quantify. One of the approaches to m-SISO control system design is to minimize the interaction between the control loops (that is, design loops so that the action of one loop does not disturb another). In theory, if all interaction could be removed, then the multivariable problem would be reduced to a set of single-loop tuning problems. In practice it is impossible to design completely non-interacting loops, but one of the popular ideas about control system design is that the best systems are those which minimize interaction. Various mathematical measures of interaction have been introduced under the assumption that the control schemes with the lowest interaction measures will be easiest to operate. While this approach is not perfect, it has been fairly successful in a number of applications.

Although many steady-state tools are designed to indicate which control schemes exhibit the least interaction between loops, some kinds of interaction may actually be useful. For example, Niederlinski (71) describes a system in which control loop interaction contributes to disturbance rejection. Thus, reducing interaction might not be the proper goal for a control system, and large interaction measures might not always imply poor control.

The following discussion of the relative gain array demonstrates that minimizing interaction is not always a well-understood goal. Part of the reason is that there is no standard definition of interaction.

2.3.2 The Relative Gain Array

The Relative Gain Array (RGA) was one of the first mathematical tools used to analyze multi-loop control systems. Although it was originally formulated with little theoretical background, a number of authors have presented theoretical justifications and interpretations to support its use as a tool for control system design. On the other hand, many authors have pointed out its significant limitations. Despite those limitations, though, authors such as Thurston (93) and Witcher and McAvoy (102) illustrate that the RGA has found widespread industrial use.

2.3.2.1 Definition of the RGA

Bristol (6) introduced the RGA in 1966 as a measure of the degree of steady-state interaction among control loops. The individual elements of the RGA are defined by

$$\Lambda_{ij} = \frac{\left.\frac{\partial c_i}{\partial m_j}\right|_{m_{k,k \neq j}}}{\left.\frac{\partial c_i}{\partial m_j}\right|_{c_{k,k \neq i}}}$$

where

Λ_{ij}	=	RGA element for the pairing c_i-m_j;
c_i	=	controlled variable i;
m_j	=	manipulated variable j.

When the steady-state gain matrix is known, there is a compact way to calculate the RGA matrix:

$$\Lambda = K \cdot (K^{-1})^T$$

where

$\cdot$ denotes a term-by-term product (not a matrix product);

Λ	=	RGA matrix;
K	=	steady-state gain matrix.

In simple terms, Λ_{ij} is the ratio of a loop's steady-state gain with all other control loops open to its steady-state gain with all other control loops closed. Thus, the size of each element Λ_{ij} is a measure of how much the gain of one control loop is affected by the action of the other loops. A value close to 1.0 indicates that the controlled variable has the same sensitivity to the manipulated variable whether or not the other controllers are active.

The following properties listed by Bristol are among those that make his the RGA a useful tool for evaluating a control strategy:

- The sum of the elements in each row or column of the RGA is 1.
- The RGA is independent of variable scaling.
- The controlled/manipulated variable pairings suggested by the RGA are independent of the arrangement of rows and columns in the steady-state gain matrix.
- RGA values much greater than 1 indicate a near-singular gain matrix.
- A loosely-tuned 2x2 control system with negative feedback is always stable if the controllers use variable pairs with positive RGA values. If the RGA value is negative for the pairs, stability is possible only under certain restrictive conditions.

2.3.2.2 Interpreting the RGA

Bristol (6) originally described the RGA as a measure of the interaction between controllers. By "interaction," he meant that the action of one control loop may change the steady-state gain of another loop. Unfortunately, there are many other definitions of interaction and many other interpretations of the RGA.

Skogestad and Morari (89) derive several mathematical relationships showing that plants with large elements in the RGA are always difficult to control in some sense. Bristol (6) first stated that large RGA elements imply a near-singular gain matrix, but Skogestad and Morari provide a mathematical basis for such a statement. They show that plants with large RGA elements can become unstable when there are even small uncertainties in the process inputs (the manipulated variables). Because of this sensitivity to input uncertainty, multivariable controllers controllers and decouplers based on the inverse of the plant matrix will perform poorly if the plant has large RGA elements. In fact, under certain conditions, inverse-based controllers may perform poorly even for a plant whose RGA elements are small. Thus, there are situations where a control system made up of many SISO loops is preferable to a multivariable control system.

Yu and Luyben (108) demonstrate that the RGA is also a measure of the amount of error in one model element that will render the model unusable for designing a control system. The amount of permissible error in an element of a plant model is inversely related to the size of the RGA. If the RGA is large for a particular pair of controlled and manipulated variables, then the model relating those two variables must be very accurate in order to correctly predict closed-loop stability.

Although it is widely described as an interaction measure, the RGA is a measure of only one kind of control loop interaction--it measures how much the gain of one loop is affected by the action of the others. Suppose, for example, a 2x2 system in which λ_{11} = 1, indicating that the sensitivity of controlled variable C_1 to manipulated variable M_1 is the same whether the second loop is open or closed. A setpoint change in loop 1 will require the same amount of control action in M_1 whether loop 2 is open or closed. There are other kinds of interaction, however. With control loop 2 closed, it is possible that a setpoint change in loop 1 will require some change in M_2 as well as M_1 in order to reach the new steady state. The RGA does not measure this kind of interaction. Thus, favorable RGA values are no guarantee that a plant will perform well, because it does not measure every kind of interaction.

Friedly (21) describes interaction in terms of the amount of steady-state offset that would be encountered if only proportional controllers were used. He demonstrates that there are some systems in which the RGA elements are close to 1 but the dynamic behavior demonstrates considerable interaction between the loops.

The exchange between Jensen (29) and McAvoy (54) illustrates that the RGA can be misinterpreted when "interaction" is not defined completely. For example, Jensen (29) points out that any system with a triangular gain matrix has an RGA equal to the identity matrix. Such an RGA indicates that, at least in some sense, the control loops are "non-interacting." In practice, however, the dynamic performance of some of the loops in a triangular system does depend on whether the other loops are closed or not. Thus, while there is no interaction in the sense that is measured by the RGA, the loops do affect each other.

Although the RGA is a steady-state measure, it does provide some information about the expected dynamic performance of a control system. McAvoy (52) states some equations relating the dynamic performance of a 2x2 control system to the value of the RGA and the relative speed of the loops. In general, 2x2 systems with large RGA values are sluggish, and systems with small RGA's (positive but less than 1) are oscillatory. Of course, closed-loop behavior depends on controller tuning; McAvoy shows that large RGA values also imply that the controller tuning derived from open-loop step responses may be very unsatisfactory when all loops are closed. The RGA can provide some indication of the amount of detuning necessary for acceptable closed-loop performance.

2.3.2.3 Limitations of the Original RGA

Although the RGA is widely used to analyze control system designs, it has some significant limitations. For example, the RGA requires specification of an equal number of controlled and manipulated variables; when more manipulated variables are available, the RGA does not directly suggest which ones should be chosen for control.

Luyben (46) describes a problem in which accurate RGA values may be derived only by using very small increments of the manipulated variables to derive a process gain matrix. Although such increments are usually possible in computer models, they often cannot be implemented in actual plant equipment. Luyben demonstrates that nonlinearities in a plant may profoundly limit the usefulness of an RGA calculated from actual plant data.

Although Bristol (6) claims that controller pairings corresponding to positive RGA elements will give a stable system in 2x2 cases, the same is not always true for larger systems. Koppel (39) presents a 3x3 system in which the variable pairings indicated by the RGA give a closed-loop system which is unstable under PI control. Koppel recommends using Niederlinski's (70) index to eliminate the unstable pairing from consideration.

Papadourakis et al. (73) demonstrate that RGA values calculated for isolated process units are not necessarily valid when those process units interact with others in a larger plant. Recycle streams in a plant often change the steady-state gains enough to invalidate RGA values calculated for individual units. As a result, controllers tuned with standard open-loop methods may be highly oscillatory or even unstable if they are connected using variable pairings recommended by the RGA. These authors also demonstrate how different inventory control policies can affect the pairings recommended by the RGA.

2.3.2.4 Dynamic Versions of the RGA

McAvoy (55) shows one way in which the concept of the Relative Gain Array (originally a steady-state tool) may be extended to dynamic analysis. The steady-state RGA gives information about the system behavior at zero frequency, but many control loops operate at relatively high frequencies where the steady-state RGA may not describe the system performance very well. McAvoy uses a matrix of transfer functions instead of steady-state gains to calculate a dynamic RGA as a function of frequency. This dynamic RGA is useful when the RGA is a strong function of frequency and when the steady-state RGA is very different from the RGA at the natural frequency of the control loop. McAvoy applies the dynamic RGA to several examples based on linear transfer function models.

Tung and Edgar (94) use a state-space model to derive a dynamic version of the RGA for setpoint changes. They show that Bristol's original RGA is a limiting case (at zero frequency) of their dynamic version. Their derivation gives a theoretical basis for the steady-state RGA, which was originally developed on an intuitive basis. Results may be interpreted in either the frequency domain or the time domain. Tung and Edgar do not, however, present guidelines to indicate when a control system will exhibit closed-loop difficulties. They provide an example of a system in which "considerable interactions" occur between two control loops, but their accompanying graphs do not suggest that such interactions will make the control any worse.

Witcher and McAvoy (102) define a time-varying relative gain array to measure dynamic process interactions. This Relative Dynamic Array (RDA) uses a "dynamic potential," defined as the integral of the open-loop step response over some time interval (Θ). The RDA is calculated exactly the same as the RGA, except that it uses the matrix of these dynamic potentials instead of steady-state gains. The choice of Θ is left to the control engineer; Gagnepain and Seborg (22) recommend calculating an Average Relative Gain Matrix, the average of the RDA over a time interval equal to the largest time constant of the process. Controller pairings based on the RDA and Average Relative Gain Matrix are determined using the same rules as for the RGA.

2.3.2.5 Block Versions of the RGA

The RGA may be extended to include controllers that are not completely decentralized. Manousiouthakis et al. (50) introduce the Block Relative Gain (BRG) to define steady-state interaction for plants with block multivariable control systems (control systems which divide the plant into several subsystems, each with a multivariable controller). The properties, usefulness, and limitations of the BRG are analogous to those of the RGA, though it is somewhat scale-sensitive. The BRG for a particular subsystem is derived assuming perfect control in all other susbystems. Arkun and Manousiouthakis (2) developed a dynamic version of the BRG that contains information about dynamic behavior. Unlike the original BRG, the dynamic BRG does not assume perfect control of other subsystems; while steady-state control may be perfect, dynamic control is generally not.

2.3.2.6 Integrating Variables and the RGA

In general, the RGA includes only variables that are self-regulating, not integrating, although dynamic versions of the RGA are able to include integrating transfer functions. In practice, though, many methods of calculating steady-state gains require a stable plant. If there is an integrating variable or other non-self-regulating variable, some control must be applied before the gains can be evaluated. This problem is not unique to the RGA. Garcia and Morari (23), for example, cite the same problem in applying internal model control--a plant must first be stabilized by conventional control before IMC can be used.

Woolverton (105) introduced a simple method to incorporate integrating variables into the RGA. The open-loop gains for a non-integrating controlled variable are usually calculated by changing manipulated variables one at a time and measuring the corresponding change in the final steady-state value of the controlled variable. Woolverton suggests changing the manipulated variables and measuring the change in the time derivative of the integrating variable. In mathematical terms, he calculates $\Delta(dX/dt)/\Delta m$, not $\Delta X/\Delta m$. The gain calculated by this procedure is the gain of a process with a transfer function of $G = K/s$.

McAvoy (55) provides some basis for Woolverton's approach. He uses a transfer-function model to demonstrate that the dynamic RGA derived by using the integrating variables is the same as the dynamic RGA derived by using the derivatives of the integrators. He gives examples, though, which demonstrate that using integrating transfer functions in RGA analysis can give erroneous results. These examples show that not all the difficulties with integrating variables have been eliminated.

Arkun and Downs (1) derive a method for using the state-space description of a plant to calculate its RGA. This method also allows the RGA to be calculated for plants with integrators; the results are equivalent to other dynamic versions of the RGA.

It is still not clear what is a proper way to incorporate integrating variables into the RGA or other kinds of gain analysis. These approaches eliminate some of the problems, but not all of them. Despite some theoretical progress in the treatment of integrators, authors such as McAvoy (55) and Papadourakis et al. (73), recommend that control loops for tank levels be closed before calculating the steady-state gains necessary for the RGA. This recommendation assumes, however, that a manipulated variable has already been selected to control each level. Such a choice is not arbitrary--Papadourakis et al. (73) show that different inventory control policies can lead to very different RGA recommendations.

2.3.2.7 Summary of the RGA

Although the RGA has been used to design control systems for a number of multivariable applications, good RGA values are not sufficient to insure good control performance. The RGA may be used as a screening tool to eliminate some control schemes (those which require pairing on a negative or very large RGA value, for example), but the RGA does not address every control design problem. Other methods must be used to evaluate alternative sets of controlled and manipulated variables, to close loops for integrating variables, to evaluate different kinds of disturbances, and to suggest beneficial ways of using ratio and cascade control.

2.3.3 Singular Value Decomposition and Principal Component Analysis

The Singular Value Decomposition (SVD) is a method of decomposing a matrix into three matrices such that

$$K = U\Sigma V^T$$

where

K	=	the original matrix (dimension mxn);
U	=	an orthonormal mxm matrix whose columns u_i are the "left singular vectors" of K;
Σ	=	an mxn diagonal matrix whose entries σ_i, called the "singular values" of K, are arranged in descending order;
V	=	an orthonormal nxn matrix whose columns v_i are the "right singular vectors" of K;

Klema and Laub (38) discuss some of the numerical properties of the SVD and suggest (in 1980) that it should be a valuable tool for control analysis. One useful property is that the rank of a matrix is equal to the number of non-zero singular values. That is, the number of non-zero singular values of a process gain matrix is a measure of the degrees of freedom in the space spanned by that set of controlled and manipulated variables. Klema and Laub point out that the uncertainty in the data produces a "zero threshold," a number below which singular values should be assumed to be zero. The condition number of a matrix, given by the ratio of the largest to the smallest non-zero singular value, is a measure of how numerical errors may be magnified in the solution of the system Ax = b. The SVD also can provide information about the row and column dependencies of a matrix. The SVD is particularly useful in analyzing non-square systems.

Lau et al. (42) and Moore (63) discuss a process-control-oriented interpretation of the SVD. If K is the matrix of steady-state gains, then U, Σ, and V have physical interpretations that can be useful in defining a control system. The vectors of V are "input directions"; that is, they are combinations of the manipulated variables. The vectors of U are "output directions"; that is, they are combinations of the controlled variables. The elements of Σ are scaling factors that relate the columns of U and V. Furthermore, the order of the vectors indicates the "ease" of control: u_1 indicates the direction in which the plant will move with the least effort; v_1 indicates the manipulated variable combination that moves the plant in that direction; and σ_1 indicates how far the plant moves in the direction of u_1 for a manipulated variable change of v_1.

The condition number of a plant is given by

$$\gamma = \frac{\sigma_1}{\sigma_n}$$

where

γ	=	condition number of K;
σ_1	=	largest singular value of K;
σ_n	=	smallest singular value of K.

The condition number and the size of the smallest singular value are two indicators of how easily the matrix may be inverted. In control terms, these numbers indicate how easily the plant can match setpoint changes. In general, given appropriate scaling, a large condition number and a small σ_n both indicate that it will be difficult to move the plant in some directions.

Principal component analysis (PCA) is a technique very closely related to SVD. While much SVD work concentrates on the sizes of the singular values, PCA is also concerned with the structure of the singular vectors. Wold et al. (103) provide some historical background on the use of PCA, which was first developed as a statistical tool with many uses, including data reduction, modeling, and selection of subsets of variables.

Wise et al. (101) demonstrate that PCA can be used with multivariate statistical process control to monitor dynamic process data in plants with more measurements than states. In such plants, PCA helps distinguish measurement noise from real changes in the state of the process.

Kresta et al. (40) describe how PCA is used in statistical work to represent the variability of a data set using only a few dimensions. The "variability" of the data set is analogous to the feasible operating region spanned by a plant's gain matrix; therefore, PCA might be a valuable tool in representing that operating region using only a few variables. McCabe (58) and Krzanowski (41) describe some criteria for selecting a few significant variables that will adequately describe a large data set; the analogous control problem is selecting a few necessary controlled and manipulated variables that will span the whole operating region.

Moore (61) and Keller and Bonvin (37) discuss the role of principal components in reducing the number of variables necessary for modeling. Morari (66) describes a feedback controller as an approximate model of the inverse of the plant transfer function matrix; if PCA can reduce the complexity of a plant model, we might expect it to be able to suggest simple control structures also.

Bruns and Smith (7) use the SVD to pair manipulated and controlled variables by choosing the largest values in the vectors of U and V. If vector u_i has a single component u_{ji} that is much larger than the other components, and if the corresponding v_i has a single component v_{ki} that is much larger than the other components, then controlled variable j and manipulated variable k should be paired together. The dominance of one term in each singular vector may be interpreted to mean that a step change in manipulated variable k produces a significant change in controlled variable j and very little change in any other controlled variable. Thus, at least in a steady-state sense, such loops are practically non-interacting.

Similarly, Lau et al. (42) define "natural loops" of a system to be SISO control loops created by observing pairs of singular vectors that are closely aligned with the standard basis vectors. Their pairing procedure is equivalent to that of Bruns and Smith (7); it essentially identifies singular vectors in which one term dominates the others. Lau et al., however, do not propose that all loops be paired in this way. Singular vectors not closely aligned with the standard basis vectors require some other pairing method, or perhaps some form of multivariable control.

Moore (63) discusses how the SVD may be used to select the controlled and manipulated variables for a distillation column. He also addresses using SVD for multivariable controllers in which there are more measurements than controlled variables. This work hints at the usefulness of creating a "reduced-rank" vector space, using only a few principal components to describe the operation of a plant.

Some control system design methods choose sets of variables based on the condition number of the gain matrix, which is a measure of the degree of interaction which may be expected from the closed-loop system. Other methods rely on the size of the smallest singular value, which is a measure of the smallest process sensitivity. Moore et al. (65) use the "intersivity index" (σ_2/γ) as a compromise between minimizing interaction and maximizing sensitivity. A large intersivity index indicates that the proposed control scheme is fairly robust.

Lau and Jensen (43) discuss the condition number of a matrix as a measure of sensitivity to parameter perturbation (such as noise or model error, for example). They also interpret the smallest non-zero singular value as a measure of the invertibility of the matrix and, therefore, a measure of the physical control effort required to achieve control.

Skogestad and Morari (88) use the SVD to define the directions in which a set of disturbances will move the operating point of a plant. They define a disturbance condition number to measure the amount of manipulated variable action necessary to reject a disturbance in a particular direction. In general, a large disturbance condition number implies a difficult control problem.

Grosdidier (26) uses the SVD to analyze the closed-loop performance of a m-SISO control system. Given a pairing of manipulated and controlled variables, the SVD can indicate which disturbances and setpoint changes might be likely to cause dynamic interaction problems.

Price (78) uses the maximum singular value as a rough measure of the sensitivity of the system. He identifies control strategies based on the singular values of their manipulated variable gain matrices (mSV) and the singular values of their disturbance gain matrices (dSV). Ideally, a control structure will have a large mSV and small dSV.

One of the difficulties with SVD-based methods is that, unlike the RGA, they are scaling dependent. The size of the singular values and the order and composition of the singular vectors all depend on the scaling of the gain matrix. The scaling method, therefore, figures prominently in control structure selection. Mijares et al. (60) demonstrate a case for which improper scaling leads to a very bad pairing recommendation using an SVD criterion.

Lau and Jensen (43) discuss several methods of scaling based on mathematical criteria, including equilibration (scaling so that the maximum value in each row and column is of the same magnitude) and geometric scaling (scaling each row or column by the geometric mean of its largest and smallest entries). They conclude that these methods are not reliable at producing SVD decompositions that accurately predict process behavior. They present a two-step scaling method in which process gains are normalized and then the rows are equilibrated.

Johnston and Barton (31) suggest that the basic function of scaling is to prevent some subset of variables from dominating the results. (They do not address whether some subset actually dominates plant behavior, in which case we might want the scaled gain matrix to reflect that fact.) They suggest these scaling guidelines:

- Scale outputs so that a change of a given magnitude in each output has the same significance in the control of the plant. Normalizing by dividing an output by its steady state value often achieves this result, but not always.
- Scale inputs so that a change of a given magnitude in each input represents an equivalent amount of control action. Representing all inputs as fractional changes from the steady-state value will usually be sufficient.
- Scale disturbances according to the expected magnitude and likelihood of each disturbance.

These scaling methods usually incorporate some subjective information that ought accurately to reflect plant behavior.

Waller et al. (98) demonstrate how the scaling of a gain matrix can yield misleading condition numbers; it is quite possible that a plant could be well-conditioned even though a particular scaling method gives a model that is poorly conditioned. They propose a distinction between different condition numbers:

- The stability of a control structure is predicted by the minimum condition number of Grosdidier et al. (27), that is, the smallest condition number that can be achieved by any scaling method.
- The performance of a control structure is predicted by the condition number of a gain matrix scaled by some physically meaningful method, a method that reflects measurement capabilities or differences in the perceived importance of certain variables.

Moore (63) recommends physical scaling by the range of the sensor and the range of the manipulator, and SVD work at the University of Tennessee has generally followed that recommendation.

2.3.4 Other Steady-State Gain Analyses

Niederlinski (70) defines the Niederlinski Index (NI) as

$$NI = \frac{\det(K)}{\prod_{i=1}^{n} K_{ii}}$$

where

NI	=	the Niederlinski Index for a set of pairings;
K	=	the steady-state gain matrix of the plant arranged so that the diagonal elements represent a proposed pairing scheme;
K_{ii}	=	the i^{th} diagonal element of K;
n	=	the dimension of the control problem.

The Niederlinski Index is thus defined for a particular set of pairings of the manipulated and controlled variables. If $NI < 0$, then that set of pairings represents a control system that will be unstable regardless of the controller tunings. Niederlinski also suggests a pairing procedure based on a critical gain and frequency for the plant, but it requires dynamic tests for every possible set of pairings.

Mijares et al. (60) introduce the Jacobi Eigenvalue Criterion (JEC) as a measure of the ease with which a control system will respond to setpoint changes. The JEC is based on the idea that the best pairing between controlled and manipulated variables is the pairing which makes the controlled system most resemble a diagonal plant. The effort of the control system to reach a new set of setpoints is similar to the iterative numerical procedure to invert the plant matrix by the Jacobi method. It is a well-known result in numerical analysis that the speed of convergence is related to the eigenvalues of the Jacobi iteration matrix, so Mijares et al. (60) suggest that the best pairing is the one whose corresponding Jacobi iteration matrix has the smallest maximum eigenvalue. The analogy between the numerical technique and the control system is not perfect, but it is most accurate when the integral control is relatively slow and all loops have approximately the same speed. For 2x2 systems, the JEC is equivalent to selecting the pairing with RGA closest to 1.0. For larger systems, the JEC is often simpler to interpret than the RGA. Unfortunately, the JEC must be calculated for every possible variable pairing, which can be cumbersome for large plants.

Although a great deal of control literature focuses on designing controllers that are non-interacting, Witcher and McAvoy (102) describe another approach: creating new manipulated variables that are naturally non-interacting. Thus, in a blending system where two flows are available, one might choose the total flow and flow ratio as non-interacting variables. In general, some physical model of the system must be available for mathematical derivation of these new manipulated variables.

Jafarey and McAvoy (28) apply the method of noninteracting variables to a 2x2 distillation column and find that the two new manipulated variables are almost identical--in other words, there is really only a single manipulated variable, a composite function of the column operating parameters, that controls both ends of the column. In some sense, then, their distillation column has only a single degree of freedom, despite the fact that two controllers must be used to find each new operating point. This result is true to some degree for all binary distillation; the lower the relative volatility, the more closely the column operation approaches a single degree of freedom. As will be shown later, the same concept may be useful in describing plant-wide control problems--in many cases, multiple measurements or multiple valves do not represent multiple degrees of freedom.

Morari and Zafiriou (68) define decentralized integral controllability (DIC) as the ability of a control system with integral action to stabilize the plant with each individual loop stable, even when the controller gains are reduced by any amount. They suggest pairing controlled and manipulated variables based on a DIC criterion.

Chiu and Arkun (11) define decentralized closed-loop integrity (DCLI) as the ability of a control system with integral action to maintain plant stability even after the failure of any combination of control loops. (This "failure" might be instrument failure, valve failure, or any other circumstance which requires operators to take a control loop out of automatic mode.) Chiu and Arkun then use the Niederlinski Index (or, equivalently, the sign of the RGA values) to rule out controller pairings that do not possess DCLI. The difficulty with this method is the number of tests that must be performed. Chiu and Arkun state that there are $2^k - (k+1)$ tests that must be performed to verify DCLI for a particular pairing in a k-dimensional control problem. In addition, there are k! possible pairings of controlled and manipulated variables. If there are several possible sets of controlled and manipulated variables, the number of numerical tests can be enormous.

2.3.5 Disturbance Gain Analyses

It is useful to analyze the steady-state gains for manipulated variables, but the primary function of most control systems is to reject disturbances. Therefore, it is also useful to have some way to evaluate the disturbance gains as well. Stanley et al. (90) define a measure called the Relative Disturbance Gain (RDG) to help quantify how disturbances affect a plant. Their study shows that different process disturbances have very different dynamic effects on a process. The RDG matrix, unlike the RGA, gives a control engineer information about how different control strategies will perform in response to a specific disturbance. A large RGA value indicates that a plant will be sluggish for some disturbances, but there may be other disturbances for which that plant will respond well. The RDG can help predict which disturbances will cause operational difficulties. The RDG has two basic uses: to determine whether interaction among controllers is favorable or unfavorable in rejecting a disturbance, and to determine what sort of decoupler, if any, should be used in a multi-loop process.

In a 2 x 2 system (2 manipulated variables and 2 controlled variables), the RDG for the first controlled variable (c_1) is defined as

$$\beta_1 = \frac{\left.\frac{\partial m_1}{\partial d}\right|_{c_1,c_2}}{\left.\frac{\partial m_1}{\partial d}\right|_{c_1,m_2}}$$

where

β_1 is the RDG for the first controlled variable;

c_1 and c_2 are the controlled variables;

m_1 and m_2 are the manipulated variables associated with c_1 and c_2;

d is the disturbance.

The RDG is a ratio of the closed-loop gain to an open-loop gain. β_1 measures how much steady-state effect the second controller has on the control of c_1. A small value generally means that the two controllers interact favorably; that is, the disturbance and the corrective action of the second controller push c_1 in opposite directions, cancelling one another and making it easier for c_1 to return to its initial setpoint. A large number generally means that the action of the second controller pushes c_1 in the same direction as the disturbance pushes it, making the first controller work harder to reject the disturbance. The definition of β_2, the RDG for the second controlled variable, is analogous to the definition of β_1.

Note that the RDG depends on the variable pairings. Therefore, it must be recalculated if new pairings are selected.

Stanley et al. (90) present several useful features of the RDG, including the fact that RDG values can be calculated from steady-state gain information. They present an approach for control system design using information from both the RDG and the RGA, including a simple way to choose an appropriate loop decoupler. Although they present examples for 2 x 2 systems only, extension to larger systems is straightforward.

Chang and Yu (9) introduce the Generalized Relative Disturbance Gain (GRDG) to extend the RDG concept to block multivariable controllers. The GRDG is also pairing-dependent, however, so Chang and Yu create the Relative Disturbance Gain Array (RDGA) as a tool for examining the GRDG values for all possible sets of variable pairings. Although the RDGA may be used to select variable pairings for decentralized control, it may also be used to compare all possible multivariable control structures. Thus the RDGA can help a control engineer decide whether a decoupling structure would significantly improve performance, and which structure would perform best. Though the GRDG and RDGA may be easily computed from the steady-state gain matrix, dynamic versions of these measures are also possible when transfer function models are available.

Skogestad and Morari (88) introduce the disturbance condition number to evaluate the ability of control systems to reject disturbances. They use the singular value decomposition (SVD) of the steady-state gain matrix to define an optimum "disturbance direction," the direction of plant movement that requires the least control action to reject. The disturbance condition number measures how much more control action is required for a real disturbance compared to the "best" (most easily rejected) disturbance. Disturbance condition numbers may be used with the RGA to decide whether the control system should use a multivariable controller based on the inverse of the plant gain matrix or use a series of controllers pairing individual controlled and manipulated variables. The disturbance condition number is thus analogous to the ordinary condition number of a matrix: A large value indicates that inverting the matrix may cause stability problems.

Skogestad and Morari define the RDG matrix for a multi-disturbance case and show how their disturbance condition number is related to the RDG. They also briefly discuss the differences between the two measures. Unlike the RDG, the disturbance condition number does not indicate the best way to pair controlled variables with manipulated variables. It can be used, however, to help identify the best sets of these variables.

Williams (100) defines a measure called the Relative Control Gain (RCG) to evaluate the way in which various control strategies propagate variation toward or away from key controlled variables. The RCG is defined as

$$\alpha_i = \frac{\left(\frac{\partial c_i}{\partial d}\right)_{c_{k,k\neq i}}}{\left(\frac{\partial c_i}{\partial d}\right)_{m_k}}$$

where

α_i = RCG value for controlled variable i;

c = a controlled variable;

d = a disturbance;

m = a manipulated variable.

The numerator of the RCG is the sensitivity of a controlled variable to a disturbance when other control loops are closed; the denominator is the sensitivity when the other loops are open. Thus, RCG values with absolute value greater than 1 indicate that the other control loops magnify the variation that a disturbance causes in a particular controlled variable.

2.4 A Plant-Wide Control Design Perspective

Many complex chemical engineering problems, not just control problems, are most easily addressed by breaking them down into smaller pieces which can be "solved" individually and then recombined. Unfortunately, designing the control system for a chemical plant is more complicated than successively finding the best loops for each individual unit. Some other approach must be used.

In an earlier era, when plant units were often separated by storage tanks, disturbances in one process unit were effectively damped out before they affected any other part of the plant. Today, the increased pressure to lower capital costs has eliminated much of this intermediate storage capacity. As a result, unit operations are more sensitive to each other. Also, the increasing cost of energy has encouraged plant designers to make greater use of heat integration and material recycles. These practices are economically efficient, but they create control problems by increasing the interaction between process units.

Luyben (44) asserts that material recycle streams create the most difficult plant-wide control problems. Recycle streams complicate control problems by increasing the degree of interaction between process units. Luyben uses a simple transfer function model to introduce his analysis of recycling systems. His studies show that the recycle loop gain is the key parameter in describing the behavior of these systems. An open-loop recycle system can exhibit integrating, underdamped, overdamped, or unstable oscillatory behavior, depending on the value of the recycle loop gain.

Coupling between process units may invalidate any control study which evaluates only a single unit in a plant. Papadourakis et al. (73) demonstrate that the gain matrix and the RGA for a distillation column, calculated using only the column model, are not an accurate measure of the column's performance as part of a plant. Kapoor et al. (36) demonstrate that a recycle structure in a plant can significantly change the dynamics of distillation, not just the steady-state gains. The time constants of a distillation column in a plant environment may be many times longer than the time constants of the column by itself. Any evaluation of gains or time constants must include the whole plant in the analysis. We may infer that optimum control structures should be designed on a plant-wide basis, as well.

Buckley (8) was among the first to decompose the control problem in terms of plant functions instead of unit operations. He observed that plant control systems perform two functions: maintaining a material balance (unit inventories and production rates) and producing a product with certain quality characteristics. He proposed designing these two classes of controllers with different speeds, so that they would not interact significantly while rejecting disturbances. He recommended the use of ratio controllers for quick approximate correction of material and energy flows, with the ratios adjusted by slower quality measurements. Buckley also suggested that preliminary control design should begin during process design, a theme that has been echoed more recently by several authors.

A plant-wide control perspective is particularly important for proper design of material balance controls. It is obvious that bulk inventories (such as tank levels) must be controlled in a chemical plant so that tanks, reactors, and distillation columns do not overflow or run dry. It is equally important, however, to insure that the plant control structure allows the individual components to enter and leave the plant without accumulating. Furthermore, it is possible to design a control structure in which every unit has a valid material balance control loop although the overall plant material balance is not satisfied. Such problems are often easy to see when they involve bulk inventories, but more difficult when they involve component inventories. Downs (12) recommends the use of component inventory tables to track the path of each component through the plant and to identify the component inventories which do not self-regulate. He also gives examples of the subtle nature of some component inventory problems.

Luyben (49) proposes that every plant has an "eigenstructure," a control structure that is naturally as insensitive as possible to disturbances and that will hold the plant near an optimum operating point for a number of operating conditions. He proposes that, instead of minimizing loop interaction, the control engineer should find control systems that naturally tend to push a plant toward its optimum operating point.

Fonyo (20) elaborates on the concept of a plant's "eigenstructure" and suggests a multi-step procedure to combine various approaches toward finding it. His procedure begins with a plant-wide material balance control structure, followed by steady-state optimizations to determine which manipulated variables should be held at a constraint, then interaction analysis (such as the RGA or Niederlinski Index) to help pair the remaining variables and determine the need for decoupling strategies.

2.5 Approaches to Control System Design

Stephanopoulis (91) describes control system synthesis as an "art" whose most capable artists are those engineers who have a deep understanding of the process they are trying to control. Some insight is required to introduce creative variables (ratios, differences, etc.) and deal with the complicating factors that make mathematical treatment difficult. Nevertheless, there are some systematic methods that will make control system design more straightforward even for the talented artist.

2.5.1 Combining Multiple Mathematical Tools

Marino-Galarraga et al. (51) present a detailed approach to control system design using the RGA and RDG. They provide a thorough interpretation of the combined results of RGA and RDG analyses (e.g., large RGA with small RDG, small RGA with small RDG). The two analyses are used to reduce a large set of possible control strategies to a smaller set of "favored" strategies. Some strategies favored by a disturbance analysis might not perform well with frequent setpoint changes, such as those suggested by many on-line optimization routines, but disturbance rejection is the primary goal of most chemical plant control systems.

Yu and Luyben (107) incorporate a number of different measures into a single procedure for designing a system of SISO controllers for a multivariable process. Assuming that a complete transfer function model is available (relating all controlled variables to all possible manipulated variables and disturbances), they propose the following steps:

- Select a set of controlled variables based on engineering judgment, understanding of the plant objectives, economics, instrumentation capabilities, or other criteria.
- Select the set of manipulated variables that gives the largest Morari Resiliency Index (the minimum singular value of the gain matrix) over a frequency range of interest. Gains are scaled by transmitter spans and valve ranges.
- Eliminate unworkable variable pairings (pairings with a negative RGA value, negative Niederlinski Index, or negative Morari Index of Integral Controllability).
- Tune all remaining pairings. (Yu and Luyben describe and recommend a method called biggest log-modulus, or BLT, tuning).
- Choose the best tuned system with dynamic simulations of load disturbances. This is the best control structure.

The procedures are described in more detail by Luyben (45). This method is straightforward, but it is not a trivial problem to get transfer function models for the tuning procedure, and the number of possible sets of pairings can be prohibitively high.

Papastathopoulou and Luyben (74) calculate the Morari Resiliency Index, condition number, RGA, Niederlinski Index, and Jacobi Eigenvalue Criterion for different control structures on a 2x2 distillation column. They report some difficulties in using ratios (boilup ratio, for example) as manipulated variables, especially in large columns with high reflux ratios. If the ratioed flow measurements are noisy (as they usually are in practice), the calculated ratio may be somewhat unreliable. In addition, the transfer functions relating the ratios to controlled variables may be difficult to derive because the columns are generally very nonlinear. Similar problems might be encountered in other plants where ratios are used to control the process.

Tyreus (95) lists a number of design methods for multivariable systems, and illustrates the use of the Inverse Nyquist Array in designing a control system for a distillation column.

2.5.2 Structural Matrices

Johnston et al. (4, 30, 32, 33) present a non-iterative approach to designing large control systems. Their algorithm is based on the concept of rank deficiency, a condition in which there are too many possible manipulated and/or controlled variables to make full control possible. The procedure begins with a structural matrix, a kind of binary representation of a gain matrix in which each element is represented as either zero or non-zero. To reduce the structural matrix until it has sufficient rank, manipulated and controlled variables are systematically deleted. Then subsystems (small sets of controlled and manipulated variables) are identified within the structural matrix. Within each subsystem, tools like the RGA may be used to pick specific pairings. This procedure may yield several different feasible structures, depending on the engineering judgments about which manipulated variables and control objectives should be deleted.

The algorithm of Johnston et al. (33) is based on analyzing single units. They define "pseudo-variables" to represent interactions between controlled and manipulated variables that are not associated with the same unit, and they illustrate how these pseudo-variables can be used to design a control system for a multi-unit plant. They claim that their procedure will not suggest infeasible control structures, but they assume that every non-zero process gain represents a significant relationship between variables. However, it is not trivial to determine which process gains are sufficiently low to be considered zero, and some mathematical basis needs to be described for making such a decision. In addition, this procedure considers only "generic rank," the maximum possible rank a matrix can have for a given arrangement of zero and non-zero entries. This generic rank is not sufficient to guarantee full rank of the gain matrix. For example, if a plant has a gain matrix of

$$K = \begin{bmatrix} 1 & 1 \\ 2 & 2 \end{bmatrix}$$

then its structural matrix

$$\begin{bmatrix} X & X \\ X & X \end{bmatrix}$$

has all non-zero elements and a generic rank of 2, even though the plant gain matrix has a rank of only 1. A controller designed to try to control two degrees of freedom in such a plant will almost certainly fail badly. Thus, procedures that depend on only the generic rank of a matrix are not guaranteed to produce workable control schemes.

2.5.3 Economic Approaches

Chintapalli and Douglas (10) point out that there are at least two different types of economic control problems. One problem requires that the plant operate as close as possible to some optimum steady-state operating point, despite the presence of disturbances; this problem is addressed by the regulatory control structure. The other problem requires finding a new optimum operating point after a disturbance enters; this problem is addressed by some higher-level process optimization strategy. Their analysis indicates that, in many cases, a well-designed multivariable linear feedback controller is just as effective as most other optimal controllers. They also suggest that there is no economic justification for minimizing control loop interaction, though a lack of interaction might make the process easier to operate.

Perkins et al. (76) discuss a method for choosing a control structure based on economic optimization. They suggest that the best control structure is the one that allows the plant to operate as close to the optimum point as possible, but far enough away that disturbances do not require the plant to violate constraints. Their analysis involves finding the set of controlled variables which permits this "closest-to-optimum" nominal operating point. The mathematical analysis uses a set of perfect controllers, assuming that only the uncontrolled variables deviate from the optimum operating point when a disturbance enters. The effect of such an assumption is to suggest that the best set of manipulated variables must include the constraints as controlled variables. In practice, this assumption is invalid, but their work supports the observation of Fisher et al. (17) that plants may be operated at near-optimum conditions when the set of controlled variables includes those variables which are constrained at the optimum steady-state operating point.

Narraway and Perkins (69) use linear dynamic models and a linear programming algorithm to select a control structure which maximizes plant profit. The analysis includes estimates of instrumentation costs, plus calculation of the operating point which will maximize plant revenue without violating constraints following plant disturbances. For small plants, their linear programming problem can be solved for each possible control structure. For plants with many possible manipulated and controlled variables, Narraway and Perkins suggest some numerical optimization techniques to reduce the number of structures requiring evaluation. Once again, they assume perfect controllers. After control sets have been selected using the economic criterion, some controllability analysis is required to insure that the selected structures will really control the plant.

Norris and Skelton (72) suggest using a quadratic cost function to select the sets of controlled and manipulated variables in a plant. Given a large set of potential variables (including, perhaps, a number of possible different sensor locations), they choose those which factor most heavily in their cost function. They also suggest that, for systems with noise, there is an optimal speed for sensors and actuators, representing a trade-off between speed of control and noise-filtering effects.

Luyben (48) applies economic analysis to the specific problem of distillation column control. He concludes that dual-ended composition control minimizes energy consumption, but that it is often possible to find a single-ended composition control scheme that consumes very little more energy. It is not clear, however, that dual-ended composition control will minimize energy consumption in a plant-wide setting where at least one stream is recycled into the plant.

2.5.4 Heuristic Approaches

Many plants are difficult to describe in the precise mathematical formulations required by some of the advanced mathematical tools for control system design. For complex or poorly defined problems, control engineers often design control systems based mostly on their experience and knowledge about how a process operates. While it may be difficult to quantify this experience and intuition, it is possible to state some guidelines that help create successful control strategies.

Govind and Powers (25) attempt to reproduce the thought process by which control engineers design large control structures without detailed dynamic models or complicated mathematical analysis. Given a process flowsheet, their intuitive control system designs come from evaluating plant objectives and cause-and-effect relationships from either experience or simple transfer function models. They construct a decision tree and illustrate how the choice of control "subgoals" will affect the choice of a control structure. Advances in computing technology have made it possible for control engineers to use more advanced modeling and calculational techniques, but it is still useful to consider how an intuitive approach can produce a workable control scheme.

Morari (66) lists some heuristics commonly employed by control engineers in solving single-variable control problems:

- Choose systems where the manipulated variable has a large effect on the controlled variable (i.e., choose systems with large gains).
- Choose systems where the manipulated variable is "close" to the controlled variable (i.e., choose systems with small time constants and dead times).
- Avoid systems with inverse response characteristics.
- Avoid systems with varying parameters and strong nonlinearities.

Principles such as these often guide the design of more complex problems, too; the heuristic approaches to the Eastman Challenge Problem, discussed later, are based on similar ideas.

Following Buckley's (8) recommendation to separate material balance control from quality control, Price (78, 79) offers a five-stage approach: production rate control, inventory control, product specification control, meeting equipment and operating constraints, and optimizing economic performance. Price tests a large number of control systems for a reactor/column system and states the following heuristics for plant-wide control system design:

- For best rejection of material balance disturbances, inventory controllers should be self-consistent and directed along the primary process path. (Self-consistency is the ability of a material balance structure to reject disturbances without help from composition controllers.)
- Select throughput manipulators internal to the process.

- Using reflux to control accumulator level tends to improve disturbance rejection. (This result may vary depending on the reflux ratio of the process.)
- Tight or mixed level controller tunings tend to perform better than loose tunings for self-consistent inventory control structures.
- Vapor boilup should be used to control the product composition for this CSTR/column system.
- Controlling intermediate compositions will improve plant performance.
- Proportional-only controllers are preferred for controlling intermediate compositions.

Sometimes operating heuristics can help eliminate control structures recommended by steady-state analyses. For example, Papadourakis et al. (73) reject a control structure recommended by the RGA analysis because experience indicates that the dynamic performance will be poor.

2.6 Control Contributions to Plant Design

Control problems are often related to decisions made during a plant's design phase. For example, Skogestad and Morari (88) point out that operating constraints are a design problem, not a control problem. If there is insufficient "capacity" in a manipulated variable to reject a disturbance, perhaps the plant design is insufficient. In a separate publication (87), they show that sensitivity to model uncertainty depends on the design of the plant as well as the design of the control system. Because a plant's design affects the precision required in the model used to design the controls, it is important for design and control engineers to communicate early in the design of a process.

Downs and Ogunnaike (14), from Eastman Chemical and DuPont, respectively, present an industrial view of the control issues involved in plant design. In recent years, product variability has become one of the principal factors in evaluating product quality, and thus one of the distinguishing features of a good plant control scheme is the ability to propagate variability away from the product stream without operator intervention. In turn, the ability of the control scheme to redirect variability often depends on the design of the plant. Unfortunately, design options that increase plant operability and controllability (bigger pipes, more tanks, more instrumentation, etc.) are often rejected because their economic benefit is hard to quantify and because the designers themselves do not operate the plants once they are built. At Eastman, two ideas recur in many design/control evaluations:

- the need to identify and control the inventories of individual components; and
- the benefits of being able to control a plant with simple schemes (such as a m-SISO PID scheme).

Proper plant design can help address both these issues. (For example, Joshi and Douglas [34] identify a procedure for adding exit points to a plant to prevent accumulation of trace components.) Downs and Ogunnaike identify the choice of throughput control as perhaps the most important decision in creating a control scheme.

Fisher et al. (19) illustrate a hierarchical decision procedure in which process controllability is evaluated a number of times during the design procedure. Likewise, they propose a hierarchical decision procedure for designing a control structure as the plant is designed. The proposed decision procedures concentrate on steady-state analysis.

Rijnsdorp and Bekkers (5, 85) describe the development of a process flowsheet software package (EPIC) which evaluates controllability during the process design phase. Integrated software is particularly useful for using design data such as equipment sizes and physical properties to calculate parameters in dynamic models. The EPIC software generates steady-state gains and time constant estimates to help control engineers design multiloop control systems.

Fisher et al. (17) offer a lengthy description of the role of a control engineer in the preliminary stages of plant design. Proper design includes making sure that the plant is controllable and operable. The plant design must include sufficient capacity in the manipulated variables, overdesign in the bottlenecks, and measurement capacity for the controlled variables.

2.7 Test Plants for Control System Design

Mellichamp (59) provides a brief overview of the history of specific industrial problems that have been addressed by control theory. In 1960, Williams and Otto (99) proposed a model of a multi-unit plant for use in evaluating control systems. The plant was made up of a reactor, heat exchanger, decanter, and distillation column. Morari et al. (67) demonstrate a control design procedure for the Williams-Otto plant and give a brief list of other authors who have addressed this problem.

Johnston et al. (33) introduced an expanded version of the Williams-Otto plant in 1985 to illustrate their SISO control system synthesis algorithm. The revised plant included two reactors, three heat exchangers, two decanters, and two distillation columns.

In 1973, Wood and Berry (104) described a simple empirical model of a distillation column, made up of six transfer functions relating product compositions to the flow rates of steam, reflux, and feed. Although the model is small, it has been used repeatedly as a simple test case, as in Lau et al. (42).

Fisher (16) describes a plant for the hydrodealkylation of toluene (HDA process). This plant is a modification of the 1967 AIChE Student Contest Problem. Originally described as a design problem, this HDA process has been used as a case study to investigate process controllability and steady-state optimization (16), the interface between plant design and control (17, 91), the concept of a plant "eigenstructure" (49, 20, 45), and operating heuristics (18).

In 1986, the Shell Process Control Workshop included a model of a heavy oil fractionater to illustrate Shell's application of Quadratic Dynamic Matrix Control. The problem and some proposed solutions are described in Prett et al. (77). The model was a series of first-order-plus-dead-time transfer functions, with stated uncertainties on the gains. The gain uncertainties make this problem more complicated than earlier ones, but the model is still linear and unconstrained.

Papadourakis et al. (73) describe a simple one-reaction plant module consisting of a single CSTR and a distillation column whose distillate is recycled to the reactor. This plant configuration, here called "Module 1," is shown in Figure 2.1.

Williams (100) emphasizes the concept of transformation of variation in her analysis of possible control strategies for the Module 1 plant. She uses component inventory tables, the RDG, and the RCG to describe how the different strategies distribute variation around the plant for disturbances in production rate and feed composition.

Price (78) also analyzes the dynamic performance of a number of control strategies for Module 1. The result of this study is his set of heuristics for plant-wide control design, which have already been cited.

Computers in Chemical Engineering (Volume 17, no. 3) published five relatively new industrial challenge problems in 1993, four of them originally presented at a session of the 1990 AIChE annual meeting. Among those problems is the Eastman Challenge Plant, which will be addressed in this research.

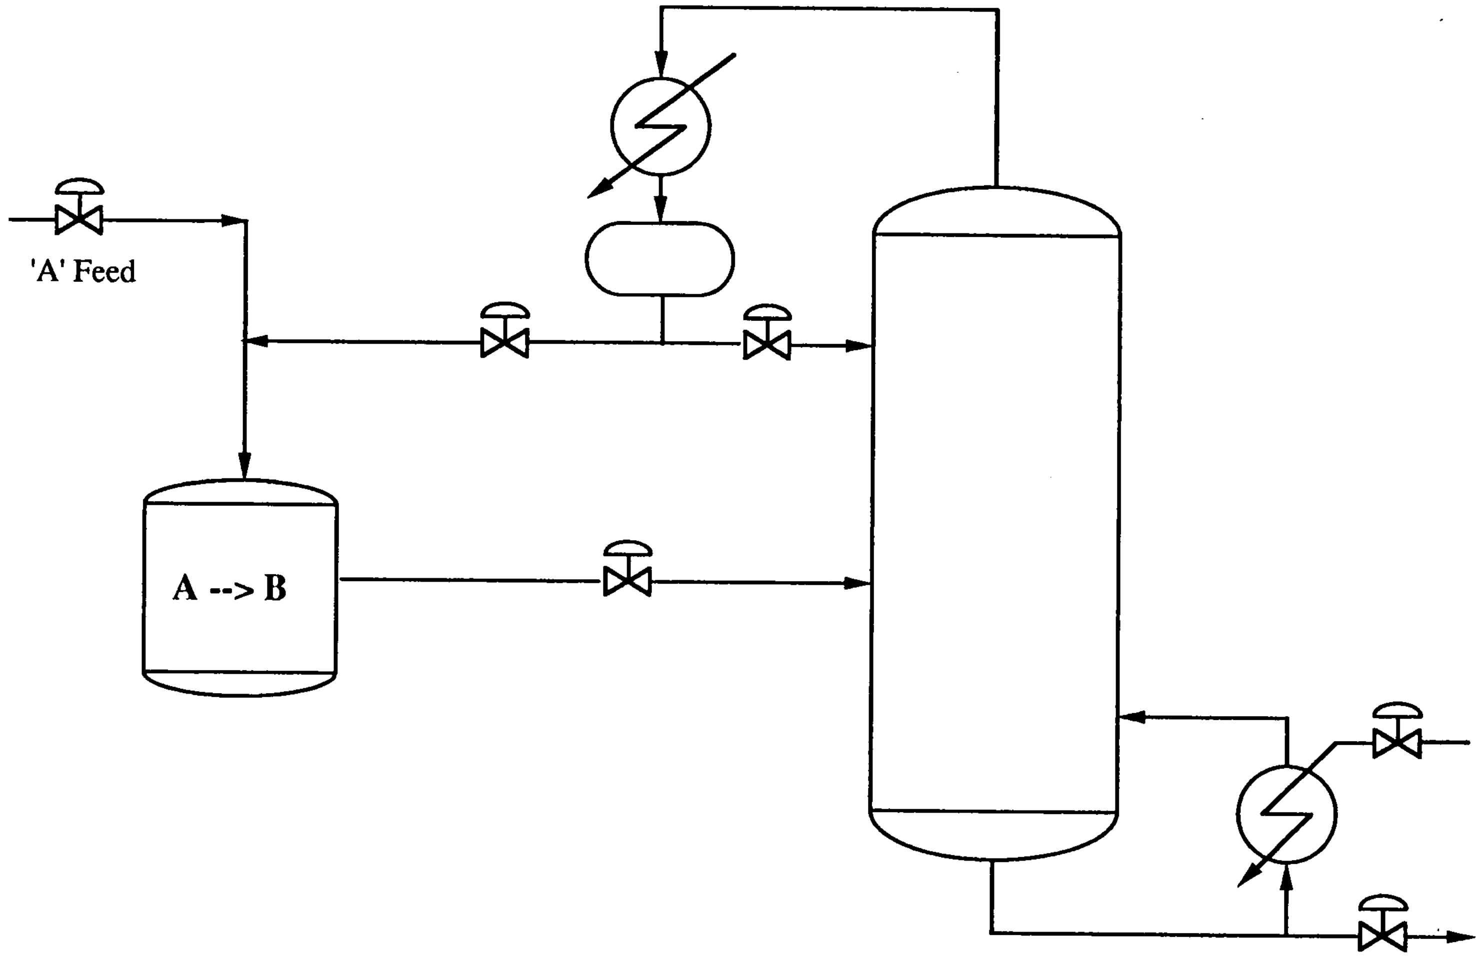

Figure 2.1 Plant Module 1

2.8 The Eastman Challenge Plant

In 1990, engineers at the Eastman Chemical Company proposed a plant model to be used as a test case for control system design procedures. One of the objectives of this research will be to design a m-SISO control strategy for this plant, called the Eastman Challenge Plant (ECP).

2.8.1 Plant Description

Downs and Vogel (15) proposed the plant shown in Figure 2.2 as a test for control system design techniques. Eight components participate in four reactions in this plant. The recycle path is more complicated than in Module 1, and there are more process variables available: 41 process measurements and 12 manipulated variables. Downs and Vogel (15) provide the nonlinear model as a set of intentionally mysterious FORTRAN subroutines for simulation. They also provide data on plant operating modes, economics, control objectives, suggested setpoint changes, and constraints for measurements and valves. Dead time in varying amounts is built into the 19 composition measurements. The FORTRAN code includes 20 different possible disturbances of various kinds, including sticking valves, composition changes, and reaction kinetics.

In addition to the size of the problem, several factors make the Eastman Challenge Plant complicated to analyze:

- Plant models are generally open-loop unstable because they include liquid levels, which are usually integrating variables. In most such cases, simple level controls will stabilize the plant, although the choice of such level controls has a significant effect on the subsequent control performance. The Eastman Challenge Plant, however, is unstable even after the separator and stripper levels have been controlled, so there is another instability to be eliminated before

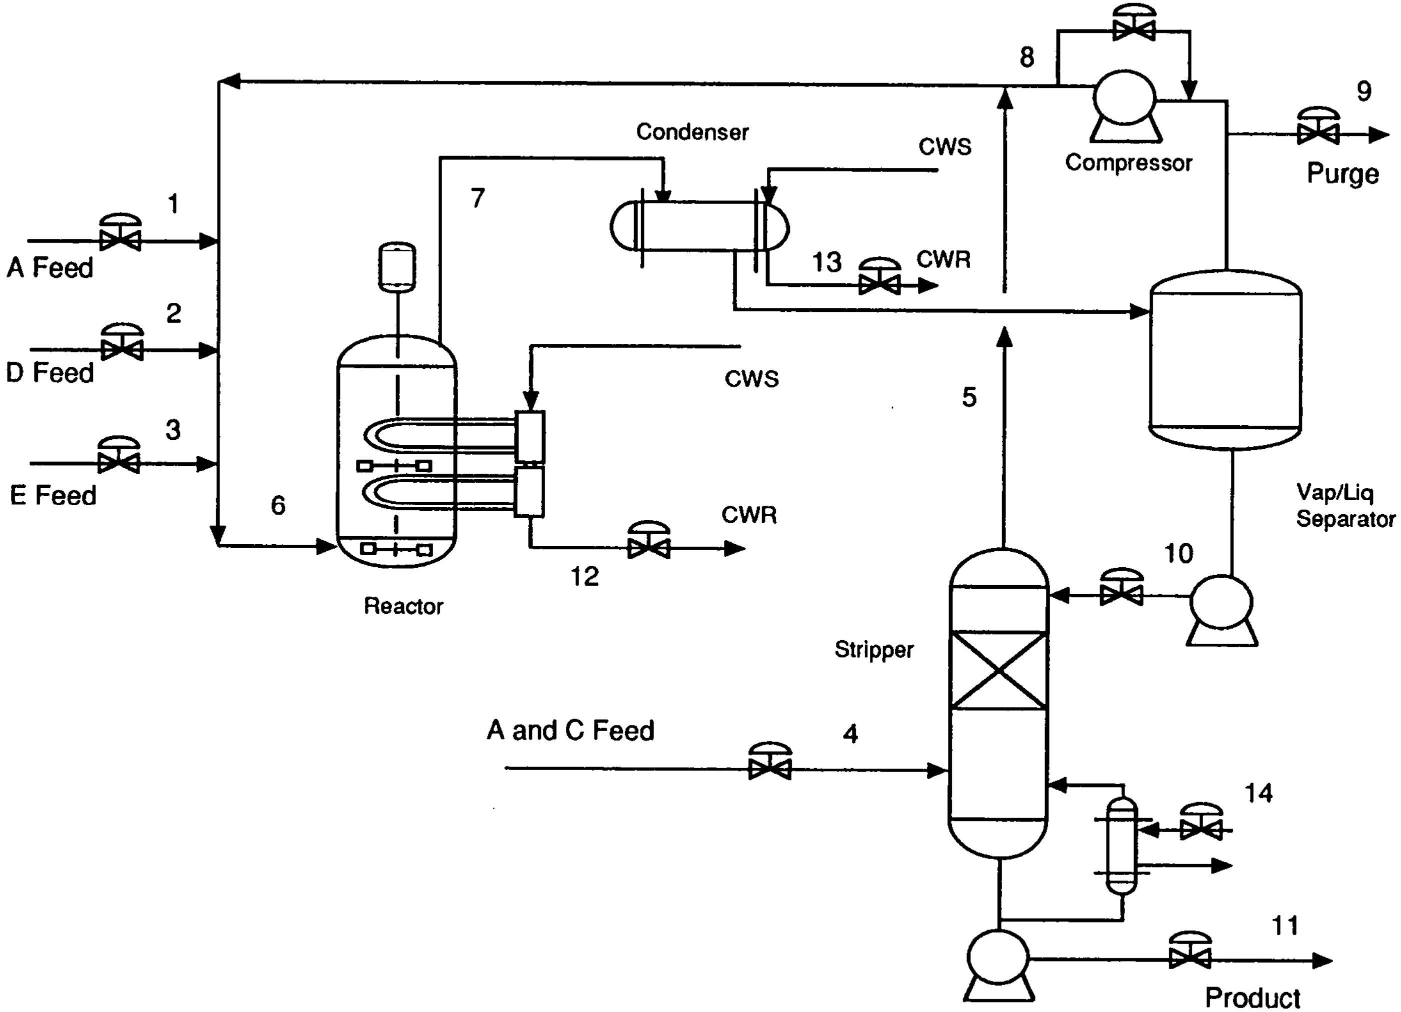

Figure 2.2 The Eastman Challenge Plant

detailed analysis can be attempted. (Curiously, the reactor level is not an integrating variable; it is therefore not necessary to assign a level controller to the reactor to stabilize the plant.)

- The dynamics of the stabilized plant are complex, including some inverse response and other forms of response not easily modeled by simple transfer functions.
- The reactor is two-phase with several components, and the pressure-temperature relationship in the reactor is not intuitive.
- The recycle path is somewhat more complicated than in other, simpler plants.
- Most authors who have addressed the problem seem to agree that the loss of the supplemental 'A' feed is a very difficult disturbance to reject, but different approaches to the problem lead to different opinions about which facet is most difficult. There is no single "most significant control problem" that can be identified at the start of the analysis.

Downs and Vogel (15) propose several different modes of operation for the Eastman Challenge Plant, including changes in the product mix and production rate. Ricker (81) provides optimum steady-state operating conditions for the base case and these different operating modes.

2.8.2 M-SISO Control Systems

McAvoy and Ye (56) in 1993 were the first to suggest a m-SISO control system for the Eastman Challenge Plant. They designed their loops in sets in order of decreasing speed: inner cascade loops first, then the basic operational controls (level, temperature, pressure), followed by the product quality and rate controls. Composition measurements were excluded from the first steps of the design, based on the premise that the control system should perform well even if an analyzer were off-line. (This premise is very similar to Price's [76] suggestion that the best material balance schemes can reject most disturbances without composition data.) Loops were chosen by experience until only about a half-dozen controlled and manipulated variables remained. RGA values were used to determine pairings for the different possible variable sets, and dynamic simulations determined the best control system.

McAvoy et al. (57) published an improved version of the McAvoy and Ye (56) solution, revising some of their previous judgments about which variables in the plant had to be controlled. Most notably, they use reactor temperature to control reactor pressure in a cascade scheme, instead of controlling the two independently. Their revised control strategy gives better control of pressure and throughput rate, with some deterioration in product composition control.

Price et al. (80) follow the plant-wide control guidelines given by Price (78) to design four different control strategies whose main difference is the choice of throughput manipulator (TPM). All four strategies were designed by engineering judgment based on understanding of the process stoichiometry, with little mathematical analysis. The best strategies used reactor coolant flow or reactor 'D' feed as the TPM.

Vinson et al. (97) designed a multi-SISO control structure incorporating extensive variables (calculated from process data). The control loops were designed mostly from heuristic considerations, though the authors do show the RGA for the appropriate set of manipulated and controlled variables. Their control strategy was not able to handle all disturbances; they suggest that "a more fully multivariable control structure with constraint handling should overcome this problem."

Banerjee and Arkun (3) design a control structure by dividing the control loops into two tiers. The first tier includes pressures, levels, and temperatures, which tend to have fast dynamics. The second tier includes the composition loops, whose dynamics are slower. Notice how this division mirrors the approach suggested by Buckley (8): one set of material and energy balance controls, and one set of quality controls. Banerjee and Arkun select sets of manipulated and controlled variables for the first tier using their judgment plus a criterion for robust control based on the singular values of the plant matrix. The controllers for the second tier (using composition measurements) were designed based on the process stoichiometry because the mathematical criteria used for the first tier are not applicable to sampled-data systems. One drawback of the robust stability criterion is that it does not identify the manipulated variables to which the plant is most sensitive. For example, this criterion was unable to recommend using the reactor cooling water flow instead of the agitator speed control to manipulate thermal energy for the reactor; although both manipulated variables have similar effects, the cooling water flow should be preferred because it has a much greater effect on the plant.

Ricker (82) also designs a PID strategy based almost entirely on engineering judgment and physical understanding of the process. Like Price et al. (80), Ricker begins by selecting a throughput manipulator. In addition to the recommendations of Price (78), Ricker suggests that a constrained manipulated variable (such as a feed stream with limited capacity) is an appropriate TPM. The loop design process begins with production rate, followed by bulk inventories (levels and pressure), product composition, component inventories, and temperatures. His attempt to find a self-optimizing strategy (in the design of the pressure controller, for example) echoes Luyben's (49) search for a control "eigenstructure." Ricker's strategy also includes some override logic (primarily for pressure control) to deal with constraints. The control scheme performs adequately for all disturbances, and appears to be well-suited for optimizing the plant operation.

2.8.3 Other Control Strategies

Other authors use more sophisticated control structures for the Eastman Challenge Plant. Kanadibhotla and Riggs (35) design most of their control strategy based on physical process understanding, but they create two nonlinear model-based controllers for the reactor temperature and the stripping factor. Despite adding nonlinear controls, they still require override controls to reject some disturbances. Interestingly, Kanadibhotla and Riggs do not use a controller for the reactor level, though they state that optimizing the operating conditions will probably require reactor level control. They also do not state a mechanism for controlling throughput rate.

Ricker and Lee (84) derive a nonlinear model based on physical principles. Their model has 26 states (compared to 50 for the original plant), 10 manipulated variables, 23 outputs, and 15 parameters which are adjusted online using an Extended Kalman Filter. They use this model (83) to implement nonlinear model predictive control (NMPC) on the challenge plant after using some heuristics to stabilize the plant. The NMPC strategy performed moderately better than the multi-SISO strategy of McAvoy and Ye (56), and significantly better for some specific operating conditions requiring control at some of the plant constraints. Ricker and Lee note that neither the model construction nor the initial choice of manipulated and controlled variables for NMPC was straightforward.

Yan and Ricker (106) propose the use of multi-objective model predictive control (MMPC), implemented with a mathematical programming technique called nonlinear goal programming, which uses a set of prioritized goals to optimize the plant control. Their MMPC strategy performs better than the original MPC solution of Ricker and Lee (83) for certain difficult operating conditions, but it is still unable to handle the complete range of conditions without override controls to reduce the throughput rate when certain constraints were violated.

Surprisingly, Ricker (82) concludes that the MMPC strategy does not perform better than his simpler PID control scheme. Based on his experience with both methods, he offers some observations on their relative merits when applied to this plant:

- Both MPC and m-SISO designs require choosing a set of controlled variables, but the current quantitative methods for choosing a control structure do not adequately address this decision.
- MPC and PID controllers perform comparably in the absence of constraints.
- The conventional MPC formulation is not flexible enough to handle manipulated variable constraints optimally. The Eastman Challenge Plant has many objectives and sets of operating conditions, and the MPC strategy is not able to account for the fact that different objectives and constraints are more significant at different operating points.
- Both MPC and PID solutions are capable of meeting the specifications in the original problem, given proper overrides to handle constraints.

2.8.4 Observations on Solutions to the Eastman Challenge Plant

Several authors point out that there is no substitute for a good physical understanding of the process. Every solution to the Eastman Challenge Plant relies to some degree on engineering judgment, and many solutions are designed primarily by judgment and heuristics rather than rigorous mathematical tools. Price (80) and Ricker (82), for example, use very little math in their approaches. Even Ricker and Lee's (84) nonlinear model predictive control solution required a great deal of engineering judgment in the *a priori* selection of manipulated and controlled variables. This heavy dependence on non-mathematical tools indicates that current process control theory is not adequate to address the wide range of issues that this problem presents.

Unfortunately, none of the authors' judgments quite agree with any of the others. For example, after Kanadibhotla and Riggs (35) claim that a proper control approach is "usually obvious" once the process is physically understood, they cite tight temperature and pressure control as keys to success for this problem, and claim that they used advanced (nonlinear) controls "only where necessary." Other authors, such as McAvoy et al. (57) claim to improve control by relaxing the reactor temperature control requirement; furthermore, they demonstrate that effective control does not require nonlinear controls at all. Ricker (82) cites three pressure control schemes that different authors have advocated based upon their judgment and physical understanding of the process. On the other hand, although these authors' judgments differed, they led to several different successful control structures.

Moore (64, page 145) says of distillation control that "the practice of selecting [sensor] location and type using rule-of-thumb and historical precedent often leads to problems of low sensitivity." A similar problem can be observed in the solutions to the Eastman Challenge Plant. Several authors incorporate the stripper steam flow as one of their manipulated variables even though the sensitivity of the plant to this steam flow is very low. For example, Vinson et al. (97) try to use steam flow to control product composition, and note that the steam valve rapidly saturates trying to maintain that composition after a change in production rate. In fact, the steam flow rate has so little effect on the plant that it that does not really represent a real degree of freedom. Ricker (81) demonstrates that the steam should be turned off for optimum operation for every plant mode listed by Downs and Vogel (15). Some kind of mathematical analysis would be appropriate to identify such manipulated variables; McAvoy and Ye (56) use the RGA as a screening tool, but the RGA does not identify this lack of sensitivity.

Most authors followed the same general pattern in designing the control loops: material balance loops were designed first, then composition and quality loops. This sequence is similar to Buckley's (8) recommendation 30 years ago, although Buckley envisioned the material balance loops as the slower set.

Finally, most authors report that some disturbances require a reduction in the production rate in order to avoid violating plant constraints. This fact may suggest either that the plant design is inadequate (i.e., some equipment has been inadequately sized) or that the control systems do not always direct variation to the right places.

Chapter 3: Research Synopsis

This plant-wide control research will be conducted using models of two different plants. The first plant, designated "Module 2," is similar to Module 1, the plant introduced by Papadourakis et al. (73). The second plant, the Eastman Challenge Plant (ECP), was proposed and described by Downs and Vogel (15).

The purpose of this study is to develop heuristic and mathematical tools to help with the design of control strategies for plants with a relatively large number of controlled and manipulated variables. Among the important issues are these:

- choice of a throughput manipulator (TPM);
- design of a material balance scheme to propagate variation in appropriate directions;
- defining the degrees of freedom (DOF) and feasible operating space for a plant;
- choosing controlled and manipulated variables that will span the feasible operating space; and
- pairing controlled variables with manipulated variables to create a m-SISO control system.

The research will involve developing and applying mathematical and heuristic tools to create control systems for Module 2 and the Eastman Challenge Plant.

The work will begin with a series of simulation studies for the Module 2 plant, designed to define principles for designing good first-level material balance controls. Those principles will then be tested by applying them to the ECP. However, the ECP is sufficiently complex that additional tools will be needed to finish the control strategy design. A mathematical procedure will be described for identifying good candidates for controlled and manipulated variables, and then a series of simulations will be run to test four proposed control strategies for the ECP.

3.1 Module 2

3.1.1 Objectives for Module 2

Although the literature gives a great deal of attention to mathematical tools for using a plant's steady-state gain matrix to define control strategies, considerably less effort has been applied to the initial material balance controls--the throughput manipulator and level controls that must be in place before that gain matrix can be derived. The design of those controls helps determine how the plant behaves, however; Papadourakis et al. (73) show that the inventory control policy has a significant effect on the results suggested by the Relative Gain Array, the most popular of the steady-state tools. The studies with Module 2 are designed to help define a systematic approach to create the best material balance control strategy.

3.1.2 Plant Description

A flowsheet of the plant is shown in Figure 3.1. Module 2 consists of 2 reactors (modeled as perfectly mixed CSTR's) and a distillation column for recycling unreacted raw materials to the first reactor. Four components participate in two reactions:

$$A + B \longrightarrow C \text{ (desired)}$$

$$2B \longrightarrow D \text{ (undesired).}$$

Both reactions are liquid-phase, exothermic, and irreversible, with rate coefficients of Arrhenius form. Each reaction is first-order with respect to the concentration of each of its reactants. There are two feeds to the reactor, one stream of pure A and one of pure B. The product stream is approximately 96.7% C, with small amounts of the other components.

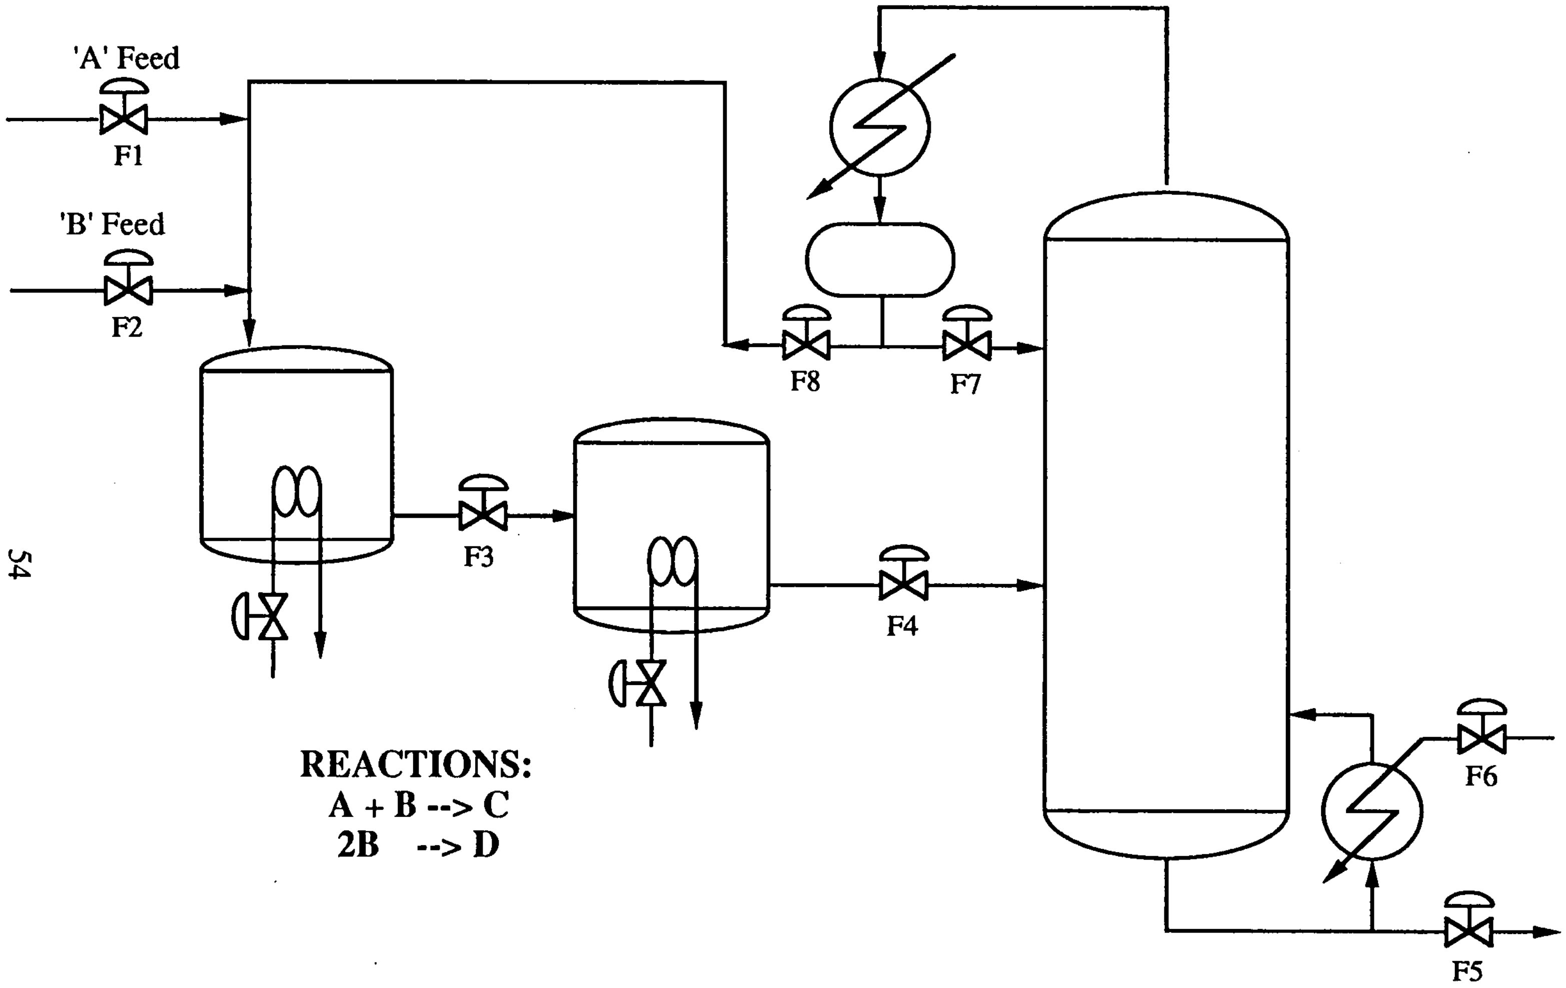

Figure 3.1 Plant Module 2

Specific information about the design parameters for Module 2 are found in Tables 3.1 - 3.5. Gas-phase behavior follows the ideal gas law.

No cost or profit function is defined for Module 2. Studies with this plant will focus on designing control strategies that permit rejection of the largest disturbances while minimizing product variability and energy consumption. The pressure drop through the column is a significant operating constraint; disturbances must be rejected without any danger of column flooding.

3.1.3 Studies with Module 2

The studies with Module 2 will produce a set of heuristics to help guide the design of control schemes for more complex control plants. The studies will investigate the effect of four factors:

Choice of throughput manipulator. The throughput for Module 2 may be set at any of several different locations within the plant. The choice of a TPM affects the subsequent design of the material balance controls and the characteristics of the disturbance rejection. These studies will investigate which TPM is best for this plant and why it is best.

Propagation of variation. When disturbances enter the plant, some manipulated variables must change in order to reject them. The control strategy design determines where this variability occurs, and whether it hurts the plant performance. These studies will address the direction in which control strategies propagate variation.

Control of component inventory. Unless there is some measure of component inventory in the Module 2 plant, one of the reactants will accumulate in the recycle stream. Plants nearly always have controllers on bulk inventories (liquid levels, for example) to prevent accumulation or depletion of material; component inventory controls prevent accumulation or depletion of individual species. Analyzers can measure the "level" of a

Table 3.1: Module 2 Stream Properties

Stream	Flow (lb-mol/hr)	Mole Fractions A	B	C	D	Temp (deg C)
A Feed	95.706	1.000	0.000	0.000	0.000	60.0
B Feed	97.808	0.000	1.000	0.000	0.000	60.0
Reactor 2 Feed	188.00	0.396	0.186	0.415	0.003	60.1
Column Feed	165.30	0.314	0.073	0.608	0.005	60.1
Distillate	67.318	0.757	0.157	0.086	0.000	60.7
Reflux	266.3	0.757	0.157	0.086	0.000	60.7
Product	97.982	0.010	0.015	0.967	0.008	73.5

Table 3.2: Module 2 Reactor Specifications

	Total Volume (cu. ft.)	Holdup (lb-mol)	Reactor Level
Reactor 1	1000	500	50%
Reactor 2	1000	500	50%

Table 3.3: Module 2 Distillation Column Specifications

Number of Stages	Feed Stage	Top Pressure	Bottom Pressure
24 (ideal)	15	760.0 torr	770.3 torr

Stages are numbered from the reboiler (1) to the top tray (24).

Table 3.4: Module 2 Reaction Parameters

	Reaction 1	Reaction 2
Stoichiometry	A + B --> C	2B --> D
Rate Law	$r_C = k_1[A][B]$	$r_D = k_2[B]$
Rate Constant Pre-Exponential Factor	$1.3723x10^{10}$ $(ft^3)^2$/(hr*lb-mol)	$2.9459x10^{17}$ ft^3/hr
Activation Energy	15,000 cal/(mol*K)	30,000 cal/(mol*K)
Heat of Reaction	-15,000 BTU/(mole A)	-15,000 BTU/(mole D)

Rate constants are of the form $k_i = k_{i,0}exp(-E_0/(1.987*T))$, with T in Kelvin. Rates are in lb-mol/hr. Negative heats of reaction indicate exothermic reactions.

Table 3.5: Module 2 Physical Properties

	A	B	C	D
Molecular Wt. (lb/lbmol)	50.0	40.0	90.0	80.0
Liquid Molar Volume (cc/gmol)	62.5	62.5	62.5	62.5
Liquid Density (lb/cu.ft.)	49.92	39.94	89.86	79.87
Antoine Constants for Vapor Pressure (torr)				
'A'	21.73	21.56	21.04	18.73
'B'	-5000	-5000	-5000	-5000
'C'	273.2	273.2	273.2	273.2
Heat of Vaporization (BTU/lb)	360.1	360.1	360.1	360.1
Liquid Heat Capacity (BTU/lb-degF)	0.5	0.5	0.5	0.5

The Antoine equation for vapor pressure is of the form P = exp(A + B/(T+C)), with T in degrees Celsius.

particular component, but they are sometimes expensive or slow. These studies will address what kind of component inventory measure is best, which component should be measured, and where the measurement should be taken.

Use of reaction volume. Luyben (44) recommends using reaction volume to control throughput. These studies will investigate the proper use of reaction volume as a manipulated variable.

Sixteen different control strategies will be created for Module 2, each strategy representing a different approach to the issues listed above. Dynamic studies will determine the effectiveness of these strategies for rejecting disturbances. The schemes will be evaluated according to the following criteria:

- variability in the compositions of the product stream;
- energy consumption during disturbance rejection;
- sensitivity of the throughput manipulator; and
- ability to reject disturbances without violating operating constraints.

Once the control strategies have been evaluated and ranked, patterns within the rankings will be used to infer heuristics for designing material balance controls.

3.2 The Eastman Challenge Plant

3.2.1 Objectives for the ECP

Control of the Eastman Challenge Plant involves many issues more complex than material balance control. It is, however, a reasonable plant on which to test some of the conclusions drawn from the study of Module 2. The heuristics drawn from the Module 2 simulations will be used to design the initial material balance controls for the ECP. In order to finish the control strategy design, however, some method is required for selecting controlled and manipulated variable sets from the many available inputs and measurements.

A mathematical procedure will be developed and demonstrated for variable set selection. The material balance heuristics, the variable set reduction method, and existing steady-state analysis tools will be used to create several possible control strategies, which will be evaluated using the set of disturbances built into the ECP model. The ECP will be a test for the proposed approach to control strategy design, an approach that incorporates both heuristic and mathematical tools in a predictable fashion.

3.2.2 Plant Description

The ECP and some proposed control systems were described briefly in the literature review in Chapter 2. The plant model offered by Downs and Vogel (15) is a set of FORTRAN subroutines. The plant, shown again in Figure 3.2, consists of a 2-phase reactor, an overhead product condenser, a vapor/liquid separator whose vapor product is recycled to the reactor, a compressor for the recycle stream, and a stripper in which the liquid product is stripped using one of the feeds. There are eight components participating in four reactions:

$$A(g) + C(g) + D(g) \rightarrow G(liq)$$
$$A(g) + C(g) + E(g) \rightarrow H(liq)$$
$$A(g) + E(g) \rightarrow F(liq)$$
$$3D(g) \rightarrow F(liq)$$

Each reaction is irreversible and exothermic. Each is approximately first-order with respect to the concentrations of its reactants, though the kinetics may vary. Products G and H are desired, F is undesired, and component B is an inert present in some of the feed streams. Downs and Vogel describe the base-case operating conditions plus five other modes in which the plant might be expected to operate. The six modes vary in their production rate and in the ratio of the desired products. The FORTRAN model includes measurement

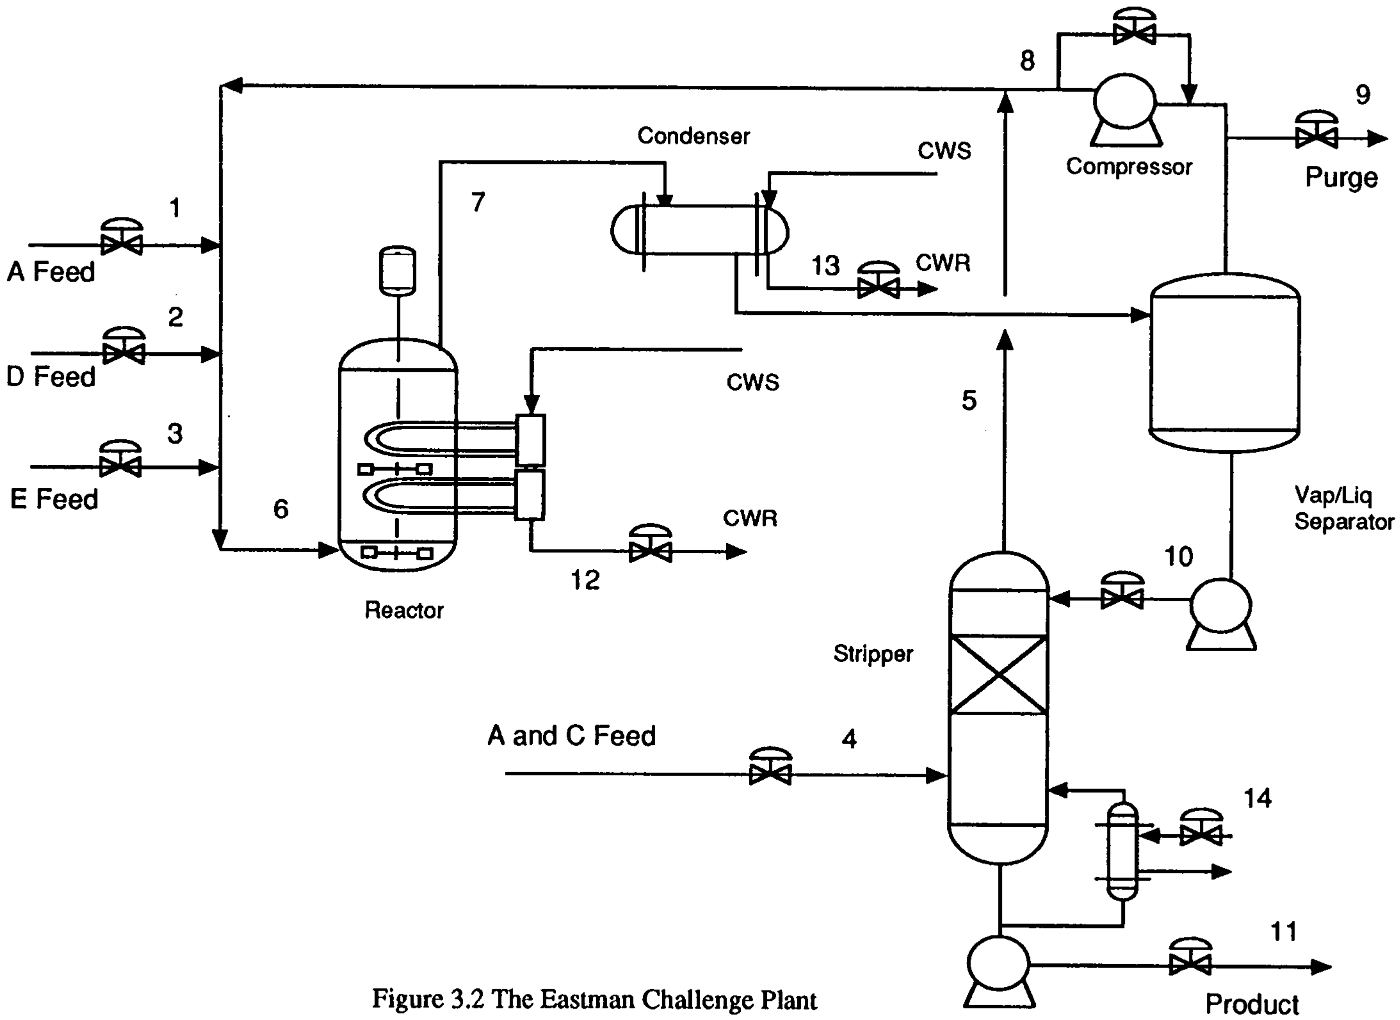

Figure 3.2 The Eastman Challenge Plant

noise; analyzer sampling times and dead times; process constraints on levels, reactor temperature, and reactor pressure; and 20 different possible disturbances. The inert leaves the process primarily in the purge stream; the undesired byproduct leaves in the purge and the product stream. More specific information may be found in Downs and Vogel (15).

The studies with the Eastman Challenge Plant will focus on disturbance rejection. The problem statement lists these objectives:

- Maintain process variables at setpoints.
- Keep process operating conditions within equipment constraints.
- Minimize variability in product rate and quality during disturbances.
- Minimize variability in flow rates that affect upstream processes.
- Recover quickly and smoothly from disturbances, throughput changes, or product mix changes.

The problem statement also includes data on the cost of operation. Economic optimization is beyond the scope of this project, but the control scheme design should consider economic factors.

3.2.3 Studies with the Eastman Challenge Plant

The control structure design for the ECP will begin with material balance controls. Heuristics from the literature and from the study of Module 2 will be used to select a throughput manipulator and necessary level controls. Once the plant has been stabilized, controllers will be designed for a set of self-regulating measurements.

The ECP is sufficiently complex that it is not immediately obvious how many degrees of freedom are available in the plant. There are many more potential controlled variables than manipulated variables, so the plant gain matrix is not square. A suitable subset of controlled and manipulated variables must be chosen before pairing can begin. A mathematical procedure will be developed to eliminate variables (either inputs or measurements) that do not help define the feasible operating region for the plant. This mathematical procedure will be used in conjunction with judgment to choose several feasible sets of controlled and manipulated variables from which control strategies may be developed.

For each set of manipulated and controlled variables, a control strategy will be designed for the ECP using heuristics and steady-state analysis tools. These control strategies will be tested and compared using the disturbances provided by Downs and Vogel (15).

Chapter 4: Studies with Module 2

Table 4.1 summarizes the control loop configuration for each of 16 different control strategies created for Module 2. The studies with this plant were designed to distinguish different approaches to first-level material balance controls. The control strategies were designed to highlight four different choices that must be made in the design of the initial controls.

4.1 Alternative Control Strategies

4.1.1 Choice of Throughput Manipulator (TPM)

The throughput manipulator is the variable (or combination of variables) that must be set to control the production rate of a plant. Downs and Ogunnaike (14) identify the choice of throughput control as perhaps the most important decision in creating a control strategy. Because the TPM is often a stream flow rate, the choice of a TPM dictates much of the rest of the bulk inventory control structure. For example, in Module 2, if the production rate is set using the feed rate to the second reactor (stream F3 in Figure 4.1), then the level in the second reactor must be controlled by the reactor effluent (stream F4), and the level in the first reactor must be set by some flow rate upstream from the reactor. In a similar fashion, the bulk inventory controls for any plant radiate generally outward from the TPM.

Four stream flow rates were chosen as potential throughput manipulators: the column feed (stream F4), the feed to the second reactor (stream F3), the 'A' feed (stream F1), and the product (stream F5). The fifth potential TPM was the reactor level setpoint;

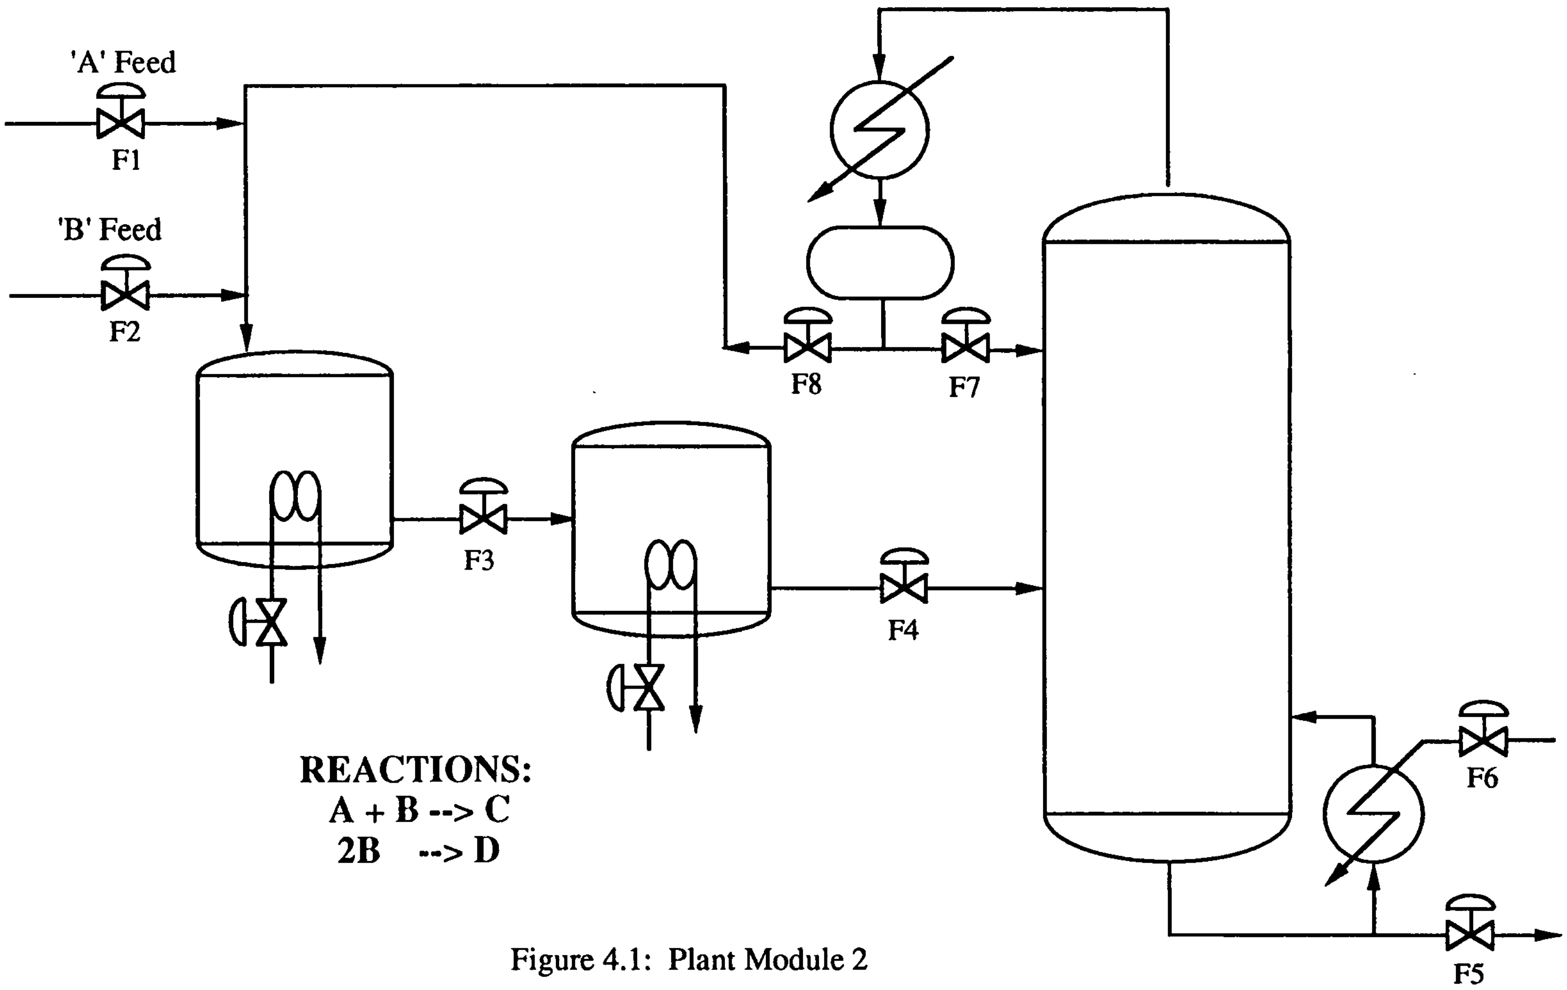

Figure 4.1: Plant Module 2

Table 4.1: Module 2 Control Schemes

Scheme	TPM	LR1	LR2	LB	CI	Other
1	F4	F1	F3	F5	R1B	Constant ratio F7/F4
2	F3	F1	F4	F5	R1B	Constant ratio F7/F4
3	F1	F3	F4	F5	R1B	Constant ratio F7/F4
4	F5	F1	F3	F4	R1B	Constant ratio F7/F4
5	LR1	F1	F3	F5	R1B	Constant ratio F7/F4
6	F4	F1	F3	F5	R1A	Constant ratio F7/F4
7	F4	F1	F3	F5	R1D	Constant ratio F7/F4
8	F4	F1	F3	F5	DB	Constant ratio F7/F4
9	F4	F1	F3	F5	RSB	Constant ratio F7/F4
10	F4	F1	F3	F5	RDF	Constant ratio F7/F4
11	F4	F1	F3	F5	R1B	Constant ratio F7/F8
12	F4	F1	F3	F5	R1B	Ratio F7/F4 controls distillate 'C' composition
1b	As Scheme 1; reactor level setpoint ratioed to TPM					
2b	As Scheme 2; reactor level setpoint ratioed to TPM					
6b	As Scheme 6; reactor level setpoint ratioed to TPM					
12b	As Scheme 12; reactor level setpoint ratioed to TPM					

Key:

TPM	Throughput manipulator	CI	Component inventory measurement
Fn	Flow rate of stream 'n' (see Figure 4.1)	R1x	Mole fraction of component 'x' in reactor 1
		DB	Mole fraction of component B in distillate
LRn	Level in reactor 'n'	RSB	Ratio of steam flow to bottoms product flow
LB	Level in column bottom	RDF	Ratio of distillate flow to column feed

All control schemes use the following pairings:

1. Reboiler heat duty controls the total A&B mole fraction in the product stream (F5).
2. Cooling duties control reactor temperatures.
3. Feed ratio F2/F1 controls component inventory, measured as indicated in the table.
4. Distillate flow controls level in column accumulator.

increasing the reaction volume increases the total reaction rate. Control strategies 1-4 are identical except for the bulk inventory controls dictated by the choice of TPM.

Stream flow rates are the most common throughput manipulators, but Luyben (44) notes that increasing production rate at constant reaction volume for recycling systems can lead to very high recycle rates. He recommends fixing one flow rate inside a recycle loop and using reaction volume to control the production rate. Strategy 5 incorporates this suggestion--the reactor level setpoints control the production rate, and the column feed rate is held constant. Note that Strategy 1 and Strategy 5 are therefore identical for disturbances which do not involve changes in the production rate.

Figures 4.2 and 4.3 illustrate how the bulk inventory controls depend on the choice of TPM. Figure 4.2 shows the inventory controls when the column feed is the TPM; Figure 4.3 shows the inventory controls when the 'A' feed is the TPM.

4.1.2 Choice of Component Inventory Measurement

Assuming that the distillation column does its job properly and separates the products from the unreacted raw materials, there will be very little 'A' or 'B' in the product stream. There is the potential, then, for component accumulation in the recycle path. The base case operating point for Module 2 requires a precise ratio of the molar flows of components 'A' and 'B'; if the actual feed ratio is slightly off, one of the raw materials will be fed at a rate greater than that at which it can be reacted away. With no place to exit the plant, that component will accumulate in the plant until some process constraint is exceeded and the plant ceases to operate. The solution to this problem is to include some measure of component inventory to monitor the level of one of the reactants and adjust the feed ratio if that inventory deviates from the desired level. (In fact, very small imbalances in the feed ratio might be regulated by the rate of the second reation, which produces byproduct 'D'

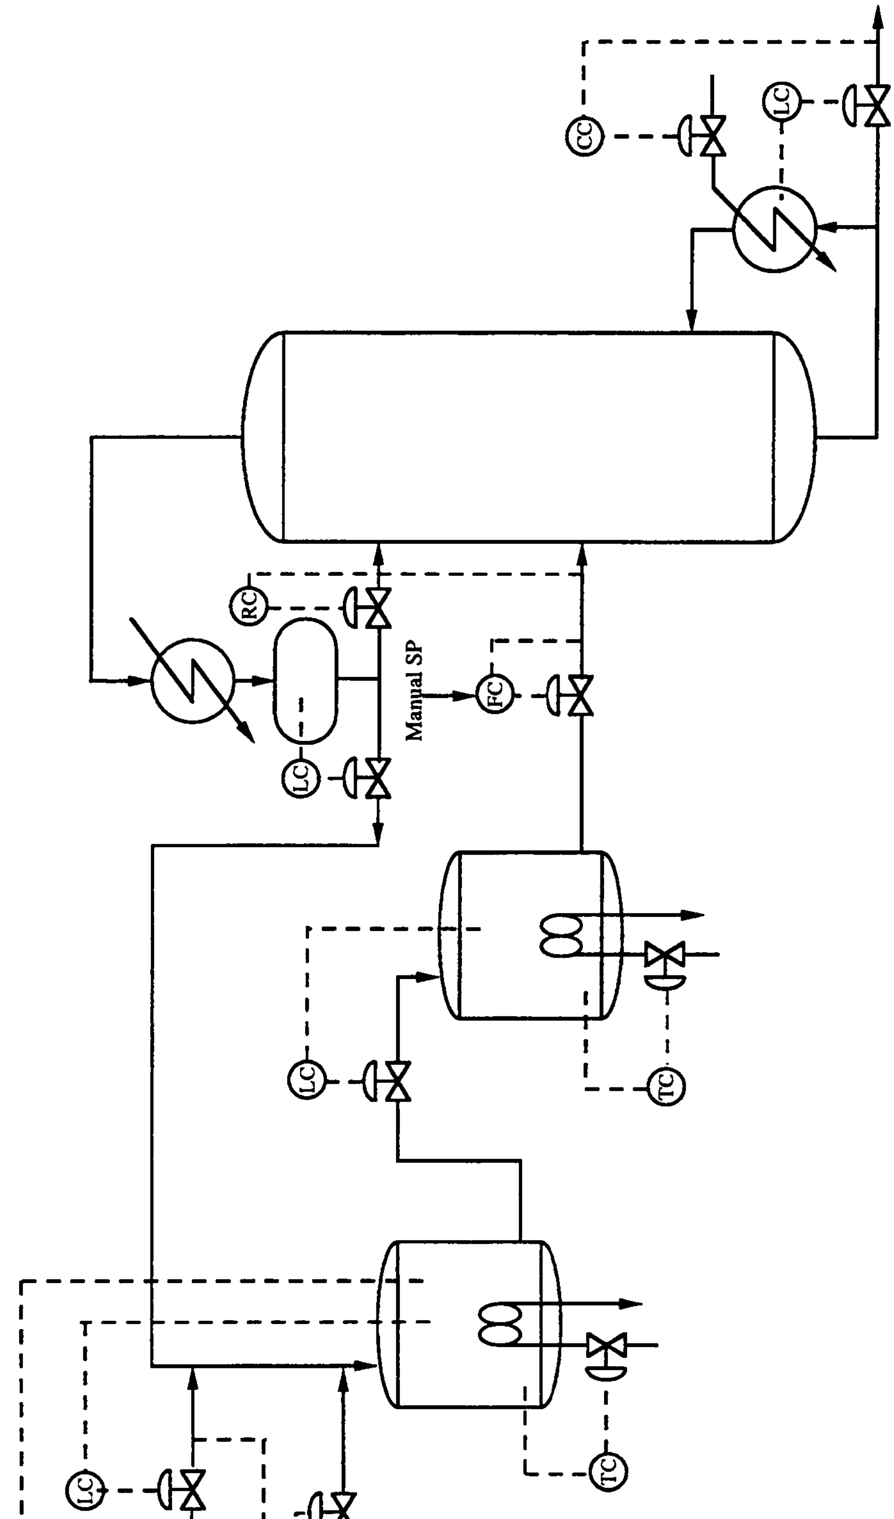

Figure 4.2 Throughput Set at Column Feed

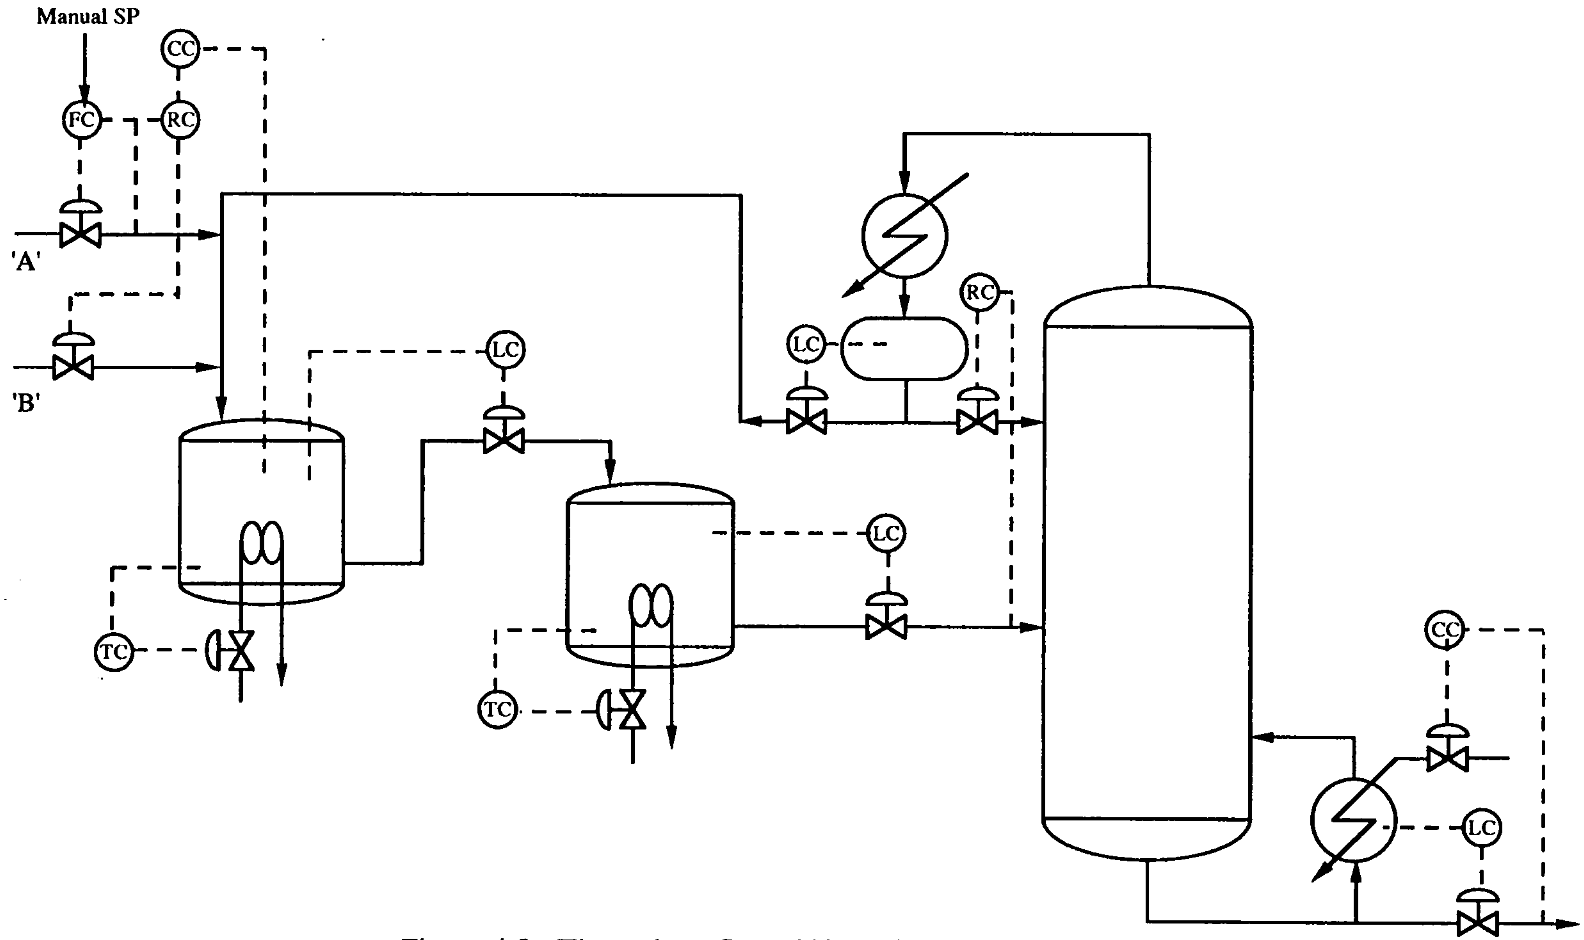

Figure 4.3: Throughput Set at 'A' Feed

from raw material 'B,' but the rate of this reaction is comparatively small. Realistically, the plant requires some measure of component inventory.)

The most easily interpreted component inventory measures are composition measurements. Unfortunately, composition analyzers usually involve some amount of measurement dead time, and analyzers may be less reliable than other process measurements. This study examined several different kinds of component inventory measurements to determine not only which component inventory should be measured, but also what kind of measurement will work best.

Four compositions were tested in an effort to determine which was the best indicator of a component buildup: the mole fraction of 'A' in Reactor 1, the mole fraction of 'B' in reactor 1, the mole fraction of 'D' in reactor 1, and the mole fraction of 'B' in the distillate. Others could have been chosen, but these were sufficient to demonstrate differences between components and locations of inventory measurement. The composition measurements were taken only every two hours to imitate the sampling characteristics of a real analyzer.

In addition, two indirect measurements were tested: the ratio of steam flow to product flow in the reboiler, and the ratio of distillate to column feed. When the inventory of one component begins to accumulate in the process because of a feed ratio imbalance, the reboiler must supply more steam to boil up the excess, and a greater fraction of the column feed will be sent overhead in the distillate. These two ratio measurements, then, might provide an indication of a feed imbalance. The advantage of these indirect measurements is that flow rates can be measured almost instantaneously.

4.1.3 Use of Reflux

Because the distillate flow is not a product stream, it is not necessary to control the composition of the distillate. It is possible to set the reflux flow based on composition or some other process objective. Three different reflux schemes, shown in Figure 4.4, were evaluated in this study. Strategies 1, 11, and 12 differ only in their use of the reflux stream. Strategy 1 ratios the reflux to the column feed. Strategy 11 ratios the reflux to the distillate (the classic "reflux ratio" scheme). Strategy 12 ratios the reflux to the column feed, but adjusts that ratio to control the distillate 'C' composition. Note that strategies 11 and 12 will vary the reflux flow rate in response to transient column conditions; in that sense, they reflect variation back into the column. Strategy 1, on the other hand, does not change the reflux to react to variation in vapor rate or distillate composition; material balance variations at the top of the column are propagated back to the reactor.

4.1.4 Control of Residence Time

Strategies with a 'b' designation (1b, 2b, 6b, and 12b) are "derived" strategies--they are identical to previous strategies except that the reaction volume is manipulated in an attempt to maintain a constant reactor residence time. Strategy 1b, for example, is identical to Strategy 1, but the reactor level setpoint is maintained in a constant ratio with the column feed rate (the throughput manipulator). The "snowball effect" described by Luyben (44, 47) results from a kinetic limitation--as single-pass conversion increases in a reactor, small changes in the desired production rate require very large changes in the recycle rate. Because the recycle rate is closely related to the distillation column load, large changes in the recycle may cause the column to run dry or flood. This snowball effect can be averted by increasing the reaction volume as the desired production rate increases. This study

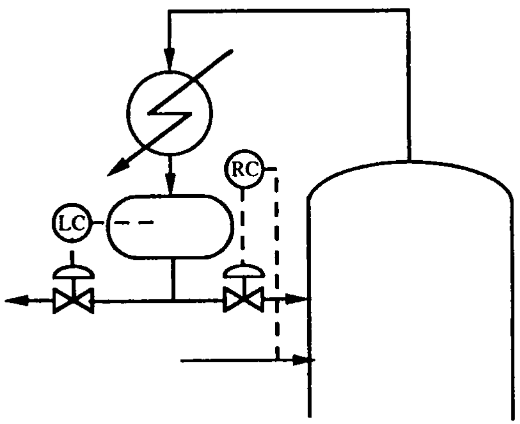

(1) Maintain Ratio of Reflux to Column Feed

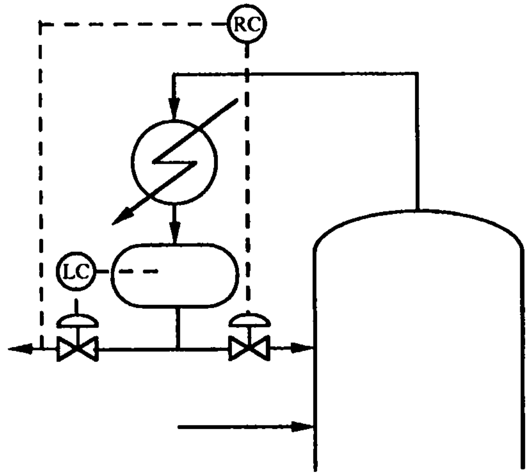

(2) Maintain Ratio of Reflux to Distillate

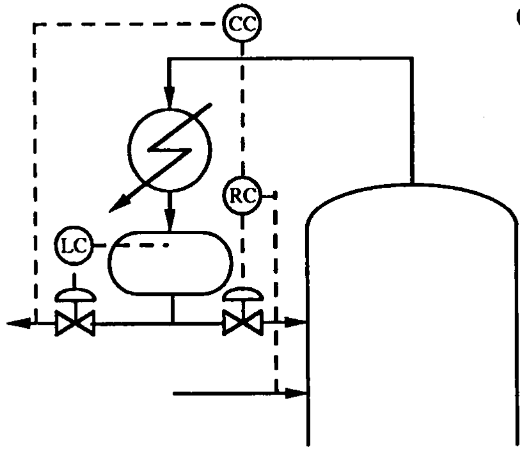

(3) Maintain Distillate 'C' Composition by Adjusting Reflux/Feed Ratio

Figure 4.4: Options for Reflux Control

investigates the effect that such a level-control policy has on the dynamic behavior of the plant.

4.1.5 Loops Common to All Strategies

Not every possible control configuration for Module 2 was tested. Some control loops are common to every strategy in an attempt to maintain a uniform basis of comparison. All strategies use reactor cooling duties to control reactor temperatures, distillate flow to control accumulator level, reboiler duty to control the total mole fraction of raw materials in the product (X_A+X_B), and feed ratio (F2/F1) to control the component inventory measurement.

4.2 Model Disturbances

Two different kinds of disturbances were used to evaluate the control strategies: throughput changes and feed composition changes. The throughput change was a step change in the TPM sufficient to increase the production rate by 5%. Table 4.2 lists the four different composition changes introduced into the feed streams.

Table 4.2: Module 2 Feed Composition Changes

	Mole Frac A in Stream F1	Mole Frac B in Stream F1	Mole Frac A in Stream F2	Mole Frac B in Stream F2
Base Case	1.00	0.00	0.00	1.00
Step Test 1	1.00	0.00	0.05	0.95
Step Test 2	0.95	0.05	0.00	1.00
Step Test 3	1.00	0.00	0.10	0.90
Step Test 4	0.90	0.10	0.00	1.00

4.3 Strategy Evaluation

The dynamic responses of these control systems to different disturbances are shown in Appendix B. Fourteen different criteria were used to evaluate the steady-state and dynamic behavior of the control strategies. These criteria are explained in detail in Table 4.3. In summary, these criteria measure energy efficiency, the ease of throughput control, changes in product quality, and ability to reject disturbances without flooding the column. Tables 4.4 and 4.5 show the numerical values for each of these criteria for each control strategy.

Criterion 2 requires more explanation than Table 4.3 provides. Each of the feed composition changes can be rejected perfectly by changing the feed ratio. During the transient response, however, there may be a temporary buildup in whichever of the components is in excess because of the disturbance. This component excess causes an increase in the recycle rate, including the load on the distillation column. The larger the feed disturbance, the larger the transient change in column pressure drop will be. Assuming that there is some maximum pressure drop beyond which the column may flood, there is some maximum feed composition disturbance that can be rejected without risking column flooding. One measure of the control system's robustness, then, is the maximum feed composition change that can be rejected without exceeding the upper limit on the column pressure drop. No more than 40% of the wrong reactant was introduced into any feed stream; the strategies that successfully rejected a change of this size were ranked by the size of the maximum change in column pressure drop during the transient disturbance response. Note that this pressure drop criterion is a purely dynamic evaluation--it is impossible to detect this difference in control strategies using steady-state gains.

Table 4.3: Criteria for Evaluating Control Strategies

Criterion Number	Description
1	Sensitivity to throughput manipulator; expressed as (% change in throughput)/(% change in TPM); measured for a 5% change in throughput.
2	Either the maximum feed composition step change that can be rejected without exceeding maximum column pressure drop, or the smallest change in pressure drop during the transient response to the largest feed composition change.
3	(Deleted).
4	IAE (integral of absolute error) of the total composition of raw materials (total mole fractions of 'A' and 'B') in the product following a feed composition change of .05 'A' in the 'B' feed.
5	IAE of the 'D' composition in product stream following a feed composition change of .05 'A' in the 'B' feed.
6	IAE of the total composition of raw materials in the product following a feed composition change of .05 'B' in the 'A' feed.
7	IAE of the 'D' composition in product stream following a feed composition change of .05 'B' in the 'A' feed.
8	IAE of the total composition of raw materials in the product following a feed composition change of .10 'A' in the 'B' feed.
9	IAE of the 'D' composition in product stream following a feed composition change of .10'A' in the 'B' feed.
10	IAE of the total composition of raw materials in the product following a feed composition change of .10'B' in the 'A' feed.
11	IAE of the 'D' composition in product stream following a feed composition change of .10'B' in the 'A' feed.
12	IAE of the total composition of raw materials in the product following a throughput rate change of +5%.
13	Steady-state change in 'D' composition in product stream following a throughput change of +5%.
14	Steady-state change in ratio of steam flow to product flow following a throughput change of +5%.
15	Integral of steam flow following feed composition changes (total for all four feed composition changes).

Table 4.4: Simulation Results for Criteria 1-8

Strategy	Evaluation Criterion 1	2	4	5	6	7	8
1	0.58	0.40A/11.7	.0040	.252	.00431	.215	.00792
1b	0.99	0.40A/11.7	.0040	.252	.00431	.215	.00792
2	0.59	0.39A/13.0	.0091	.266	.00766	.227	.0215
2b	1.00	0.39A/13.0	.0091	.266	.00765	.227	.0215
3	0.99	0.06A/13.0	.1725	2.37	.168	2.33	.386
4	1.00	0.20A/13.0	.0530	.317	.0392	.272	.133
5	0.37	0.40A/11.7	.0040	.252	.0043	.215	.00792
6	0.80	0.40A/11.3	.0036	.138	.00389	.123	.00705
6b	1.00	0.40A/11.3	.0036	.138	.00389	.123	.00705
7	0.93	0.40A/11.3	.0046	.255	.00657	.230	.00782
8	0.63	0.40A/12.5	.0056	.214	.00857	.190	.00946
9	1.43	0.40A/11.3 0.08B/Marginal	.0085	.540	.0208	.942	.0142
10	1.00	0.40A/11.4 0.19B/Marginal	.0058	.426	.0103	.449	.00917
11	0.67	0.20A/13.0	.0175	.274	.0117	.235	.0473
12	0.54	0.40B/11.7	.0301	.217	.0267	.184	.0677
12b	1.00	0.40B/11.7	.0301	.217	.0267	.184	.0677

(Note: Schemes 9 and 10 have two results listed for Criterion 2; these schemes were robust for one kind of feed composition change but not for the other.)

Table 4.5: Simulation Results for Criteria 9-15

	Evaluation Criterion						
Strategy	9	10	11	12	13	14	15
1	.595	.0094	.432	.0100	3.51E-4	4.21	1070
1b	.595	.0094	.432	.0083	6.30E-7	1.30	1070
2	.629	.0145	.454	.0084	3.47E-4	4.17	1210
2b	.630	.0145	.454	.256	2.10E-7	0.73	1180
3	4.98	.34	4.93	.164	3.46E-4	4.15	34750
4	.744	.0721	.543	.0649	3.47E-4	4.16	2220
5	.595	.0094	.432	.025	5.90E-4	-3.95	1070
6	.315	.0083	.252	.0087	2.47E-5	2.15	480
6b	.315	.0083	.252	.0135	1.90E-7	0.73	480
7	.567	.0155	.465	.0118	1.93E-4	1.43	860
8	.525	.021	.383	.0106	2.55E-4	2.15	1380
9	1.08	Unstable	Unstable	.04	7.68E-4	0.00	(Estimated)
10	.881	.027	1.04	.0137	2.87E-4	1.13	1410
11	.646	.0209	.47	.0603	3.60E-4	4.58	1660
12	.515	.0537	.369	.0349	3.37E-4	3.99	700
12b	.515	.0537	.369	.0293	0.00	0.73	700

Several of these criteria measure the dynamic response for different sizes of a single kind of disturbance. In order to avoid skewing the final rankings in favor of strategies that responded well to that particular criterion, criteria 5, 7, 9, and 11 were grouped into a single composite measurement describing the product 'D' composition IAE following a feed composition change. Similarly, criteria 4, 6, 8, and 10 were grouped into a single composite measurement describing the product raw material composition IAE following a feed composition change. Thus, the final criteria used for ranking were 1, 2, 12, 13, 14, 15, and two composite measures.

4.4 Strategy Rankings

Table 4.6 shows the ordinal rankings for each control strategy. Equal rankings were given to strategies whose performance measurements were sufficiently close to be considered equal. In cases for which two strategies were structurally equivalent for a particular disturbance, the pair were given equal rankings and were considered only a single strategy for purposes of "counting down" the list. (For example, strategies 1 and 5 are equivalent for feed composition changes, and the strategies with 'b' designations are exactly equivalent to their "parent" strategies for feed composition changes.) For this reason, not all rankings go from 1 to 16.

The control strategies were ranked from best to worst based on the sum of the rankings from each of the individual criteria, as shown in Table 4.6. The rankings for individual areas of comparison are shown in Tables 4.7-4.9. The patterns in these rankings suggest some guidelines for designing good control strategies.

Table 4.6: Ordinal Rankings of Module 2 Strategies

	Evaluation Criterion								Rank	Final
Strategy	1	2	C1	C2	12	13	14	15	Total	Rank
1	7	3	2	4	4	13	12	4	49	9
1b	2	3	2	4	1	4	7	4	27	3
2	7	6	5	6	2	10	12	5	53	10
2b	2	6	5	6	16	3	4	5	47	7
3	2	11	11	11	15	10	12	11	83	16
4	2	7	10	8	14	10	12	10	73	14
5	9	3	2	4	9	15	1	4	47	7
6	4	1	1	1	3	5	9	1	25	2
6b	2	1	1	1	7	2	4	1	19	1
7	3	1	3	5	6	6	8	3	35	5
8	6	5	4	3	5	7	9	6	45	6
9	1	10	8	10	12	16	2	9	68	13
10	2	9	6	9	7	8	6	7	54	11
11	5	7	7	7	13	14	16	8	77	15
12	8	3	9	2	11	9	11	2	55	11
12b	2	3	9	2	10	1	3	2	32	4

Criterion C1 is a composite of criteria 4, 6, 8, and 10 (the IAE of raw material composition in the product stream following four feed composition step disturbances).

Criterion C2 is a composite of criteria 5, 7, 9, and 11 (the IAE of 'D' composition in the product stream following the same four feed composition step changes).

Table 4.7: Throughput Manipulator Rankings (From Best to Worst)

Strategy	TPM
5	Reactor Level Setpoint
1	Distillation Column Feed
2	Feed to Reactor 2
4	Product Flow
3	'A' Feed

Table 4.8: Component Inventory Measurement Rankings (From Best to Worst)

Strategy	Component Inventory Measurement
6	'A' composition in Reactor 1
7	'D' composition in Reactor 1
8	'B' composition in distillate
1	'B' composition in Reactor 1
10	Ratio of recycled distillate to column feed
9	Ratio of reboiler steam flow to product flow

Table 4.9: Reflux Control Rankings (From Best to Worst)

Strategy	Reflux Scheme
1	Reflux ratioed to column feed (the TPM)
12	Reflux ratioed to column feed, with ratio reset by distillate 'C' composition
11	Reflux ratioed to distillate flow (constant reflux ratio)

4.4.1 Ranking by Throughput Manipulator

4.4.1.1 Ranking All Strategies

Strategies 1-5 differ only in their bulk inventory control schemes, with the level controls dictated by the choice of the TPM. Table 4.7 shows the relative rankings for these strategies. In each of these strategies, there is one stream whose flow rate is not manipulated for inventory control, and the scores from Table 4.7 show a definite pattern: the farther this "fixed stream" is from the column feed, the worse the ranking, especially for feed composition changes. The best performers are Strategies 1 and 5, which fix the flow rate of the column feed manually. Strategies 2 and 4 fix, respectively, the flow rates of the Reactor 2 feed and the product stream (each one unit operation removed from the column feed), and they do not perform as well. Strategy 3 fixes the flow rate of a plant feed (two unit operations removed from the column feed), and it performs by far the worst.

The effect of the TPM on performance is best explained by examining the direction in which each control strategy propagates variation. Buckley (8) described material balance controls in terms of whether they were set up in the direction of flow (with manipulated variables downstream of the levels they control) or in the direction opposite flow (with manipulated variables upstream of the levels they control). The same directional description can be applied to variation, also: the direction of the material balance controls is also the direction in which those controls propagate variation. If the material balance controls propagate variation toward critical locations (bottlenecks or product streams), measured performance will be worse.

Figure 4.5 demonstrates the effect that the control strategy (including the choice of TPM) can have on everyday operation. Assuming that there is some upper limit to the pressure drop in the distillation column, there is some maximum feed composition step change that can be rejected without violating this limit. This maximum feed composition

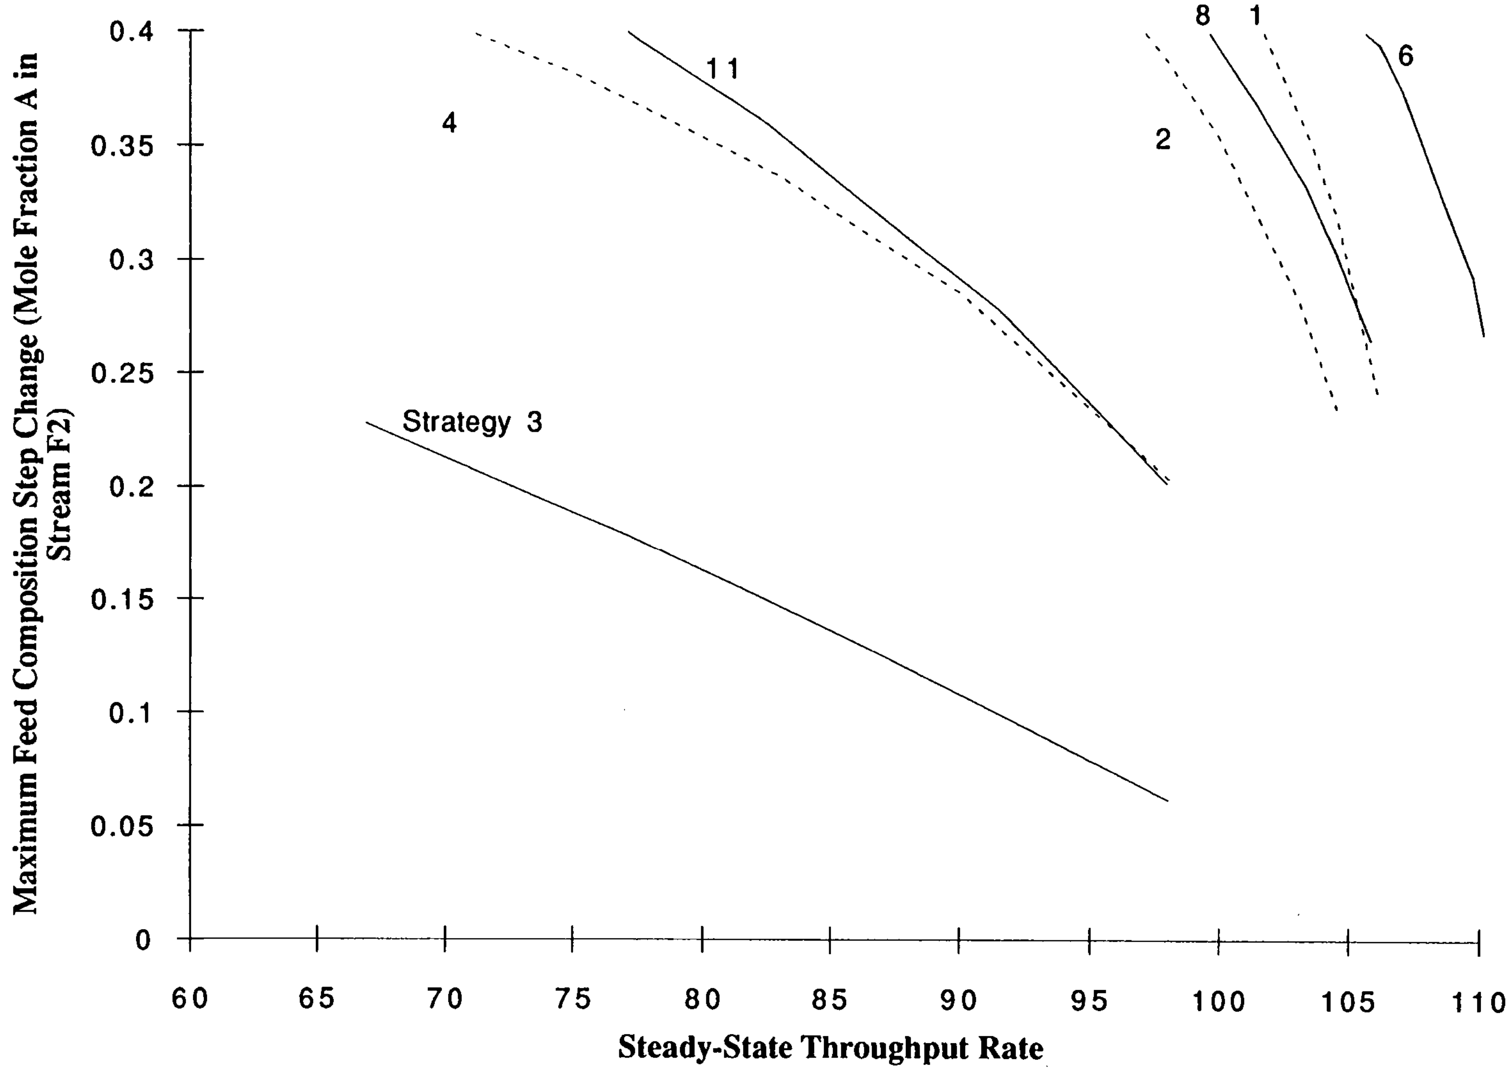

Figure 4.5: Maximum Allowable Feed Composition Change for Various Production Rates

change is also a function of the throughput rate: lower throughput rates (and lower base-case column loads) permit larger feed composition changes. Figure 4.5 shows the maximum throughput rate that can be maintained for a given range of feed compositions. Control strategies with operating lines in the upper right hand corner can reject large feed composition changes even at high throughput rates; strategies with operating lines in the lower left corner require low throughput rates to reject even small feed composition changes. It is apparent from this figure that two plants with identical equipment but different control strategies may have very different throughput ratings. Furthermore, the strategies with the best operating lines on Figure 4.5 set the throughput at the column feed, and by far the worst line is for Strategy 3, whose TPM is the 'A' feed. Other reasons for the poor performance of Strategy 3 are explained later.

Note from Table 4.6 that there was only a small difference in the overall scores for Strategies 1 and 5. The two main differences between them are that that Strategy 5 required much less steam than Strategy 1 after a production rate increase (because the increase in residence time increased the overall conversion), although Strategy 5 forced more raw material into the product stream during the transient period following the throughput change.

4.4.1.2 Poor Performance for Strategy 3

It is apparent from Tables 4.4 and 4.5 that Strategy 3 often performed much worse than any other strategy, even worse than can be explained by propagation of variation. There is another factor that makes Strategy 3 so bad, a reason which suggests that a plant feed stream is often a poor choice for the TPM. At the base case, the 'A' feed stream provides 95.7 lbmol/hr of 'A' to the plant, almost exactly the same as the number of moles of 'C' which are produced by reaction. Suppose now that the feed composition changes. If some amount of 'B' is present in the 'A' feed, the feed rate of component 'A' drops and so does the production rate. If some amount of 'A' is present in the 'B' feed, the component inventory controller calls for more 'B' feed to restore the correct feed ratio, but this control action will also raise the total feed rate of 'A'! In fact, if 'A' enters with the 'B' feed, a severe change in the production rate may result as the proper feed ratio is restored, even though the TPM has not changed.

In practice, streams that use an internal flow as the TPM are able to restore the original steady-state operating conditions following a feed composition disturbance, because they are able to simultaneously raise one feed rate and lower the other. Strategy 3 is unable to do so, and therefore is unable to regulate the production rate in the face of feed composition disturbances. Figure 4.6 illustrates how three different strategies respond to the same disturbance. Strategies 1 and 2 experience a fairly small change in column pressure drop, which decays over time as the original molar feed rates are restored. Strategy 3, on the other hand, cannot restore the original steady state molar feed rates; the dramatic permanent rise in column pressure drop is the result of the fact that the production rate rises dramatically even for a modest feed disturbance. This same observation can be made of any control strategy which uses a raw material feed rate as TPM: that stream will be ineffective as a TPM if it is possible for the molar flow of that raw material into the plant to change.

4.4.1.3 Manual vs. Automatic Production Rate Control

In the control strategies used for Module 2, there was no actual control loop that monitored the product flow rate and adjusted the TPM. Instead, the TPM was set manually at whatever value was necessary to maintain the proper production rate. Many of the difficulties with Strategy 3 are related to the fact that feed composition changes produce throughput changes as well. These difficulties might be partly relieved by an automatic

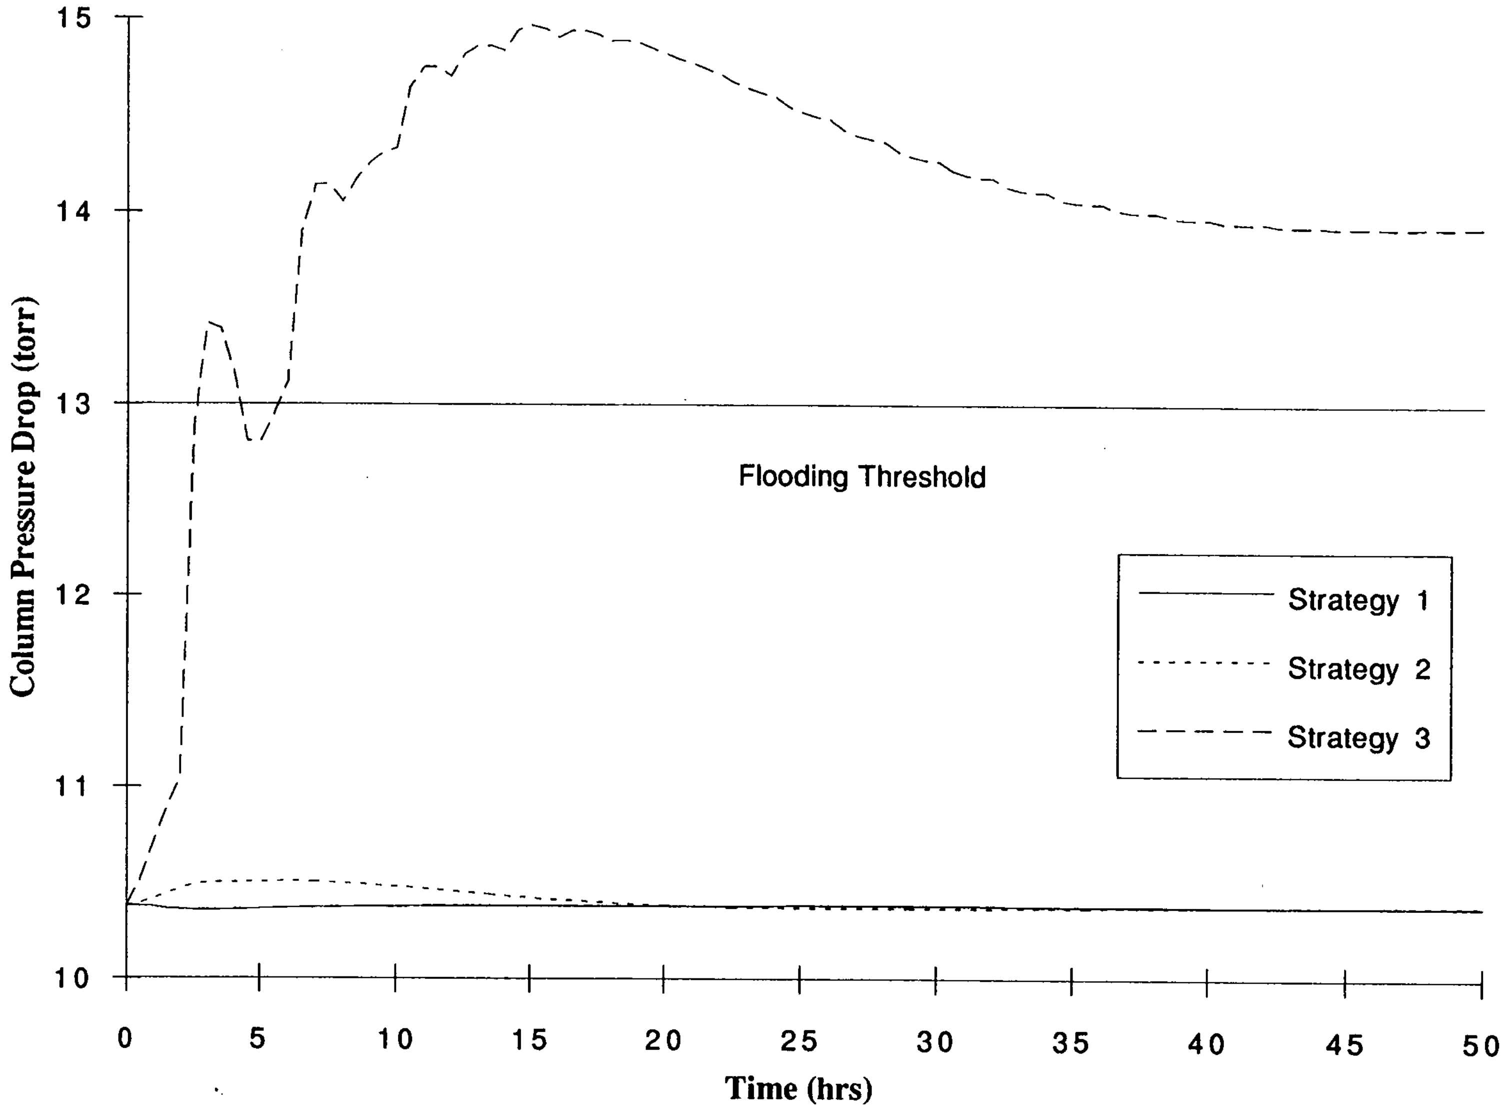

Figure 4.6: Response of Module 2 Control Strategies to Feed Composition Change (10% A in B Feed)

throughput rate controller. However, the dynamics of such a controller would be necessarily very slow, so that most of the dynamic problems associated with Strategy 3 would still manifest even with a production rate controller in place. The steady-state problems would be eliminated only if the frequency of the feed composition changes were very low.

4.4.2 Ranking by Component Inventory Measurement

Strategies 1 and 6-10 differ only in their component inventory measurement. According to Tables 4.4 and 4.5, all of these strategies performed comparably in response to a change in production rate. Feed composition changes, however, distinguished good from bad strategies; Table 4.8 ranks these strategies from best to worst. Strategy 6 was clearly the best performer in the group, while Strategies 9 and 10, which used indirect measurements of component inventory, were inferior.

The biggest problem with the indirect component inventory measures in Strategies 9 and 10 is a combination of nonlinearity and slow controller response. This plant exhibits a kind of multiple steady state behavior related to the compositions of A and B in the reactors. At the base case, there is an excess of A in the plant; the B is almost completely consumed. Under these conditions, an increase in the amount of recycled material, as measured by the column distillate/feed ratio or steam/product ratio, indicates that there is too much A in the plant, and that the B/A feed ratio should be increased (either by increasing the flow rate of B or decreasing the flow rate of A, depending on which way the level and ratio controllers are set up). The plant can operate, though, with an excess of B instead of A; under those conditions, an increase in either of the two measured ratios indicates that the B/A ratio should be decreased because there is too much B in the plant.

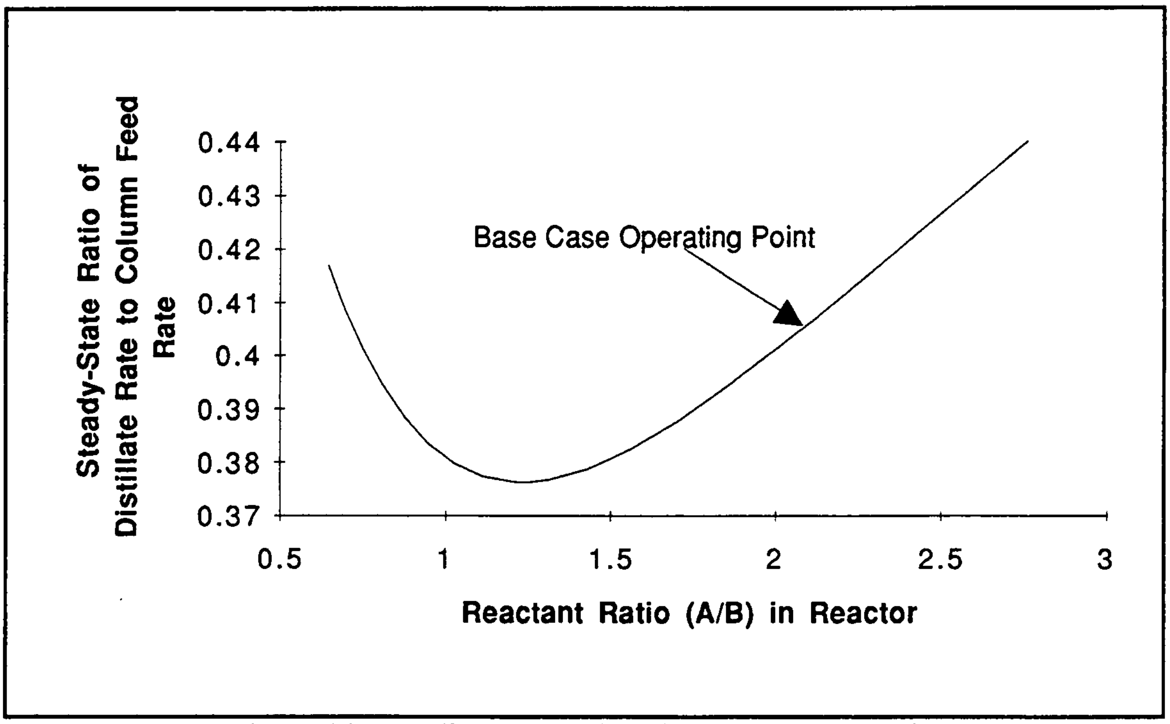

Figure 4.7: Nonlinear Behavior of Distillate/Feed Ratio

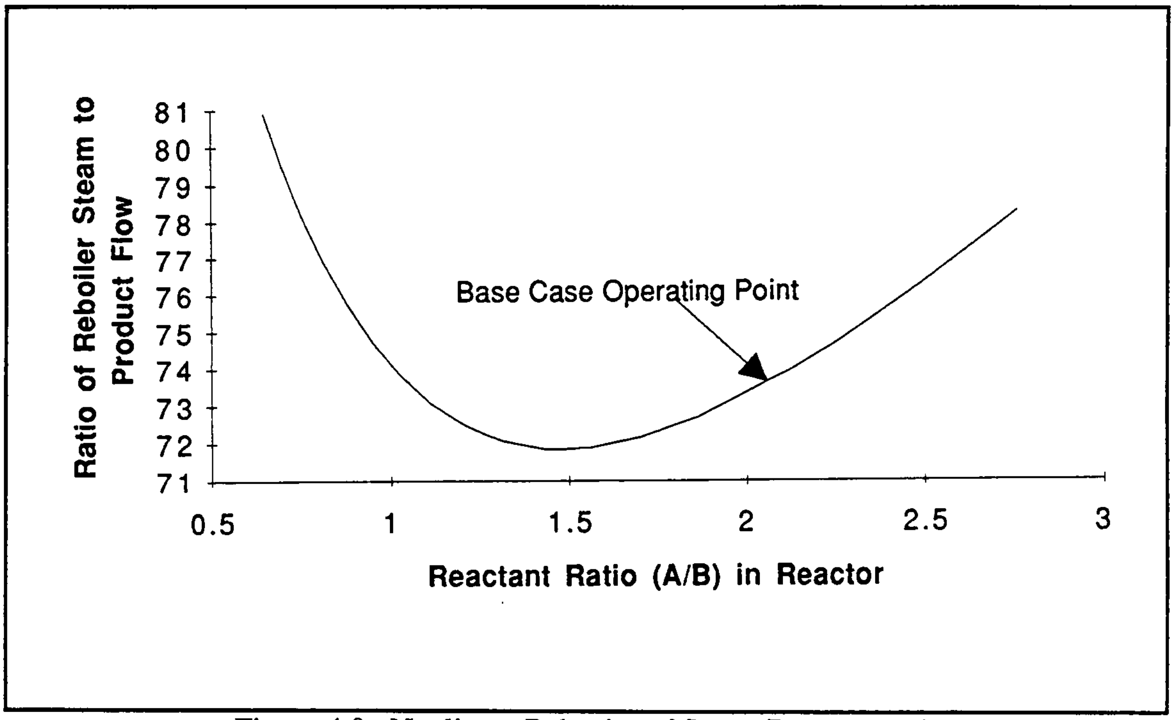

Figure 4.8: Nonlinear Behavior of Steam/Product Ratio

In short, the relationship between the measured ratio and the B/A ratio is very nonlinear, as shown in Figures 4.7 and 4.8. On the left half of either curve in these figures, increasing the B/A feed ratio will decrease the measured ratio. On the right half, increasing the B/A feed ratio will increase the measured ratio. The base case operating point is marked on both graphs.

This nonlinearity becomes important when the feed composition changes so that the B/A feed ratio increases (for example, if some B is introduced into the A feed). If the composition change is small, the plant stays in the A-rich operating region, and the component inventory controller adjusts the feed stream flow ratio correctly to restore the original steady-state. If the composition change is large, however, the plant might move well into the B-rich operating region. (The response time of this control loop for Strategies 9 and 10 is very long, so the plant moves well away from the base case before the controller can respond adequately.) If the plant moves far enough into the B-rich region, the controller will take the wrong action to try to restore the balance, and the plant will be unstable.

This very nonlinear behavior is the reason that Strategy 9 is listed as "Unstable" for some feed composition disturbances in Table 4.5. Raising the controller gain will increase the speed of the response and perhaps prevent the plant from moving into the wrong operating region, but there is an upper limit to the allowable controller gain. There is no controller gain that will stabilize the plant for every size of feed composition disturbance.

The control strategies that use composition as a component inventory measure do not suffer from this problem. An increase in the B/A feed ratio will move composition measurements in the same direction regardless of whether A or B is in excess. Therefore, the inventory control strategies based on composition measurements can always return the plant to the original operating condition after a feed composition disturbance.

Of course, there is at least one other drawback to indirect component inventory measurements: they may be affected by other things besides feed ratio imbalances. For example, if the component inventory is controlled by monitoring the steam/product ratio, a change in the steam enthalpy will change this ratio and lead to a false "correction" of the feed ratio.

4.4.3 Ranking by Use of Reflux

Strategies 1, 11, and 12 differ only in their use of the reflux flow. Strategy 1 maintains the reflux in a constant ratio with the column feed. Using this strategy, the reflux is constant except immediately following a throughput rate change. Strategy 11 maintains a constant reflux/distillate ratio. Strategy 12 controls the distillate 'C' composition by adjusting the ratio of reflux to column feed, thus establishing composition control at both ends of the column. Table 4.9 ranks these strategies from best to worst.

When the column profile changes in response to a disturbance, Strategy 1 does not vary the reflux to propagate any of this variation back into the column. Strategies 11 and 12 do vary the reflux in response to changes in the column profile; as a result, there is increased variation in the product quality. Strategy 12 ranks higher than Strategy 11 largely because the dual-ended composition control strategy uses less energy than the other strategies, in agreement with Luyben's (48) statement that dual-ended composition control is the most energy-efficient way to control a column.

4.4.4 Ranking by Control of Residence Time

Luyben (44, 47) describes a problem he calls the "snowball effect"--a nonlinear relationship between recycle rate and throughput rate. In plants in which the reaction volume is fixed, a throughput increase requires an increase in the conversion, which requires an increase in recycle rate. As the desired conversion increases, a large increase in recycle flow may be required for only a small increase in throughput rate. Luyben recommends using the reactor level setpoint as a possible throughput manipulator to help eliminate this problem. One difficulty with this suggestion is that large level setpoint changes may be required for small throughput changes.

These approaches (increasing flow rates or increasing level setpoints) address two different objectives: increasing the plant throughput and keeping the recycle rate low. In an attempt to find a compromise between these objectives, four control strategies were created in which the reactor level setpoints were increased in proportion to increases in the TPM, so that the residence times remained nearly constant. For these strategies, the reactor levels and recycle rate both increased moderately in order to increase throughput. In every case, the derived strategy (with constant steady-state residence time) performed better than the "parent" strategy.

4.5 Module 2 Summary

- When the column feed is the throughput manipulator, the bulk inventory control strategy directs variation away from the plant bottleneck and rejects the largest disturbances at high throughput rates. As the TPM moves farther from this point, plant performance deteriorates. The best option for throughput control is to adjust the column feed and reactor levels proportionately.
- The reactant feed was an especially poor choice for the TPM. First, it is as far as possible from the plant bottleneck. Second, the plant responds very poorly to feed composition disturbances because they cause net throughput changes as well. Strategies which use feed streams to control inventories performed much better.

- There is no need for dual-ended composition control on the distillation column. Plant performance is best when reflux is ratioed to the column feed. Product quality is noticeably less variable when the reflux does not respond unnecessarily to material balance changes in the column.
- Composition measurements are the best choice for controlling component inventory measurements. Other indirect measurements of component inventory are less effective because of poor steady-state and dynamic behavior. The best measurement for controlling the component inventory is the composition of A, the reactant in excess.
- Maintaining a constant reactor residence time improves the dynamic and steady-state performance of the control systems for which it was tried.

4.6 Extrapolating Module 2 Results

There is always some level of uncertainty associated with generalizing the results from the study of a specific plant. However, experience with Module 2 suggests that the following policies may be useful in creating stable, robust control systems for a variety of chemical plants.

- At the beginning of the control strategy design process, determine which plant constraints are likely to limit the plant performance <u>both at steady state and during the dynamic response to disturbances</u>. Choose a throughput manipulator and bulk inventory control strategy that will direct material balance variation away from those points and from significant quality measurements.
- A good throughput manipulator should not be subject to any disturbance that would significantly change the plant throughput at steady state.

- Component inventories are best controlled using composition measurements. Less direct measures should only be used with caution. Components present in excess might be preferred as inventory indicators.
- Unlike most bulk inventories, reaction volume has a significant effect on plant operation. Insightful manipulation of reactor level setpoints can improve performance.

Chapter 5: Bulk Material Balance Controls for the Eastman Challenge Plant

The purpose of the studies with Module 2 in Chapter 4 was to develop some heuristics for the design of first-level material balance controls (throughput manipulator, level controls, etc.). The Eastman Challenge Plant (ECP), introduced in Chapter 2, provides a worthy test for these principles. The flowsheet for this process is shown again in Figure 5.1.

5.1 Applying Flow Controls

The first step is to apply flow control to streams for which flow sensors and valves have been provided:

- Stream 1 (A feed)
- Stream 2 (D feed)
- Stream 3 (E feed)
- Stream 4 (A&C feed)
- Stream 9 (Purge flow)
- Stream 10 (Separator underflow)
- Stream 11 (Product flow)
- Stream 14 (Steam flow)

These flow loops were tuned to give slight overshoot to a setpoint change.

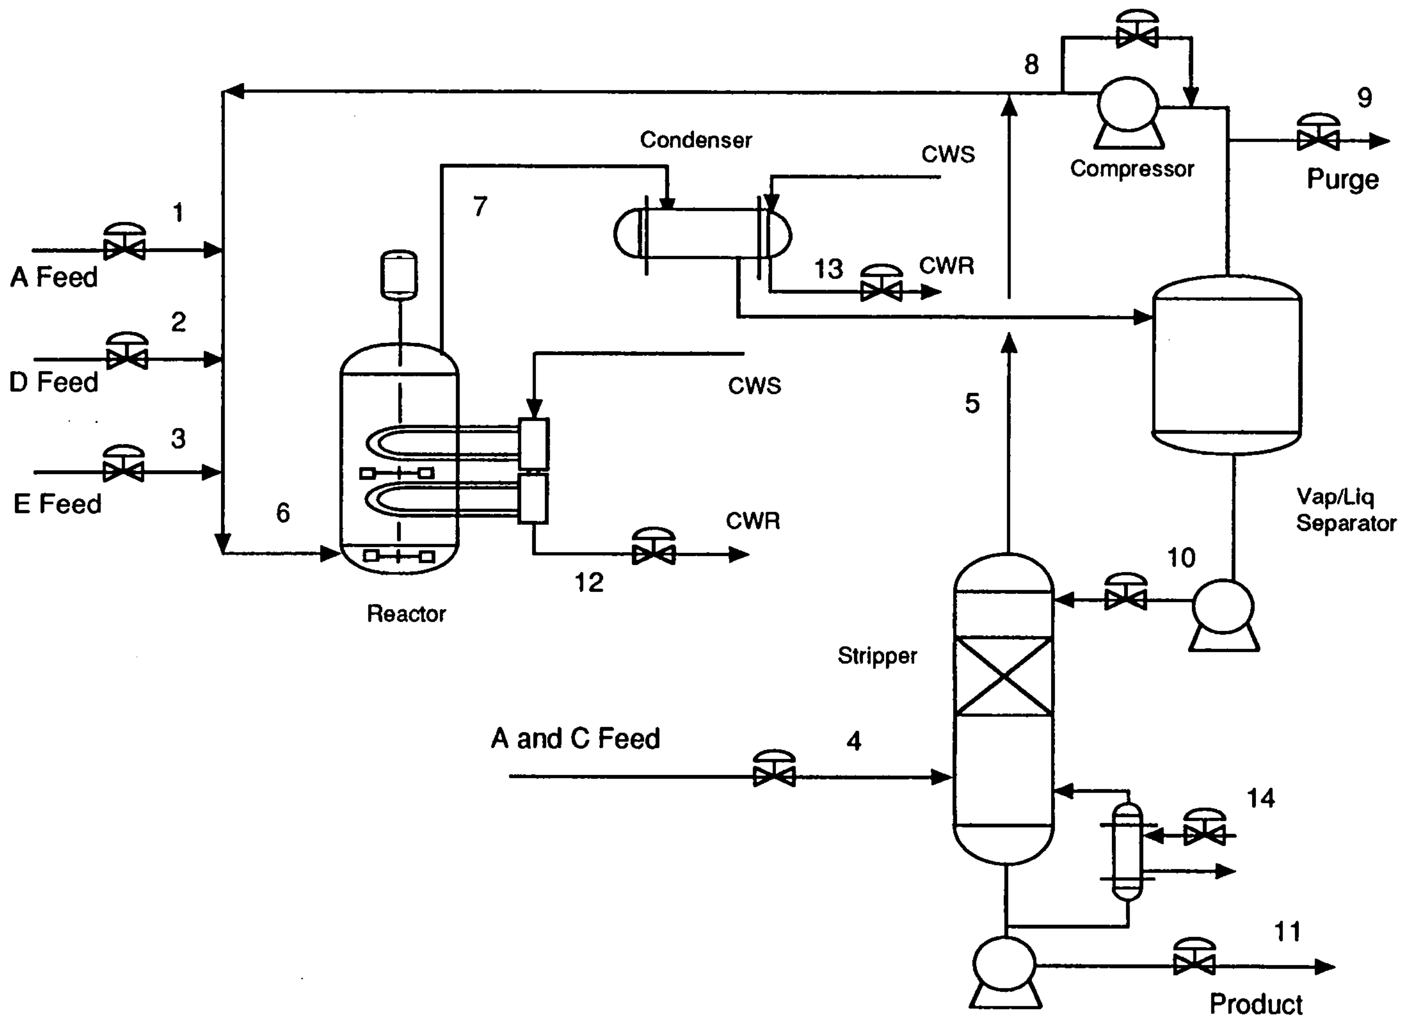

Figure 5.1: The Eastman Challenge Plant

Even this first step is not completely routine. Some of the flow rates depend on the process pressure. Putting the feed streams on flow control helps to protect the process from disturbances in the feed supply pressures. However, putting the purge stream on flow control eliminates some of the self-regulating nature of the process--small variations in process pressure might be compensated by slightly higher purge flows without the flow control. The purge stream is small enough that the benefits of flow control probably outweigh the disadvantages.

5.2 Selecting a Throughput Manipulator

The real control system design decisions begin with the selection of one variable as the TPM and the subsequent design of bulk material balance controls. Some principles from the Module 2 studies and from Price (78) can help guide these choices:

- Set the production rate in front of a process bottleneck to minimize variation propagated to that unit operation.
- If all else is equal, give preference to a TPM internal to the process.
- Do not choose a TPM subject to disturbances that would significantly alter the throughput.
- For inventory control, give preference to stream flows that can be adjusted at the inlet or outlet of the holdup.
- Inventory control loops should propagate variation away from "sensitive" locations (bottleneck, product purification, etc.).

Downs and Vogel (15) do not directly identify a process bottleneck, but they do list some operating constraints, including the reactor pressure. The nominal operating point is fairly close to the maximum allowable pressure, and brief experience running the simulation suggests that the reactor pressure is sensitive to even small disturbances. These observations suggest that the reactor pressure is probably a bottleneck.

Immediately we are faced with a dilemma: the bottleneck and the product stream are on opposite ends of the plant. Any inventory control strategy that propagates variation away from one of these locations must propagate it toward the other, so we must choose which location is more critical and move variation away from that point. Given the unusual nature of the reactor and the fact that pressure problems will cause plant shutdown, we will choose reactor pressure as the more critical variable and seek to direct variation away from the reactor.

What, then, is the TPM that will protect the reactor pressure from excessive variation? There are several candidates available--the A feed, the D feed, the E feed, the A&C feed, the condenser coolant flow, and the recycle flow are obvious possibilities. Following the example of Module 2, we may wish to consider manipulating reaction inventory (either by the level setpoint or the pressure setpoint) along with some flow rate, but that opportunity will come later. The A feed may be quickly eliminated as a prospect because it is so small and because it may sometimes be unavailable. The A&C feed would be a poor choice because it is subject to composition changes. Similarly, the condenser coolant flow might also be a poor choice because the condenser coolant inlet temperature is subject to random variation. The recycle flow is questionable; at this point in the design process, it is not certain that increasing throughput requires increasing the recycle flow. (In fact, Ricker's (81) analysis of optimum operating points suggests that the optimum recycle flow does not increase significantly with increased production rate.)

Only the D and E feeds remain as possibilities. High-frequency variation is undesirable in the D feed, but not necessarily in the E feed. Setting the D feed rate manually will prevent high-frequency variation, so we prefer the D feed as a TPM.

There is another significant reason to prefer the D feed as TPM. Ricker (82) proposes that the D feed is a good candidate for TPM because it is constrained at maximum production. If the D feed were used to control some other variable, it would not be effective at maximum production because of the constraint.

5.3 Stabilizing the Plant

In order to derive a steady-state gain matrix, it is necessary to have sufficient control of the plant that it will be stable following small changes in the plant inputs. In general, this stability requires at least level controls. Once the reactor has been identified as the most critical unit, the level controls for the separator and stripper are straightforward--they will be designed to direct variation downstream and away from the reactor. The separator level will be controlled with the separator underflow setpoint, and the stripper level will be controlled with the product flow setpoint.

Curiously, the reactor level is not an integrating variable like the other levels. For this reason it is not necessary to define a level controller for the reactor before proceeding. Reactor level gains will be calculated just like the gains for any other process variable.

On the other hand, there is a degree of instability in the plant that does not have to do with unit levels. Even with the level controls in place, without any disturbances, the process is still unstable at the nominal operating point. The reactor pressure, shown in Figure 5.2, shuts down the process within 3 hours. It is not immediately obvious what causes the instability, but it can be remedied by adding a loop that controls the reactor coolant outlet temperature with the coolant flow. For the sake of consistency, we add a similar loop on the condenser coolant. The resulting plant is stable.

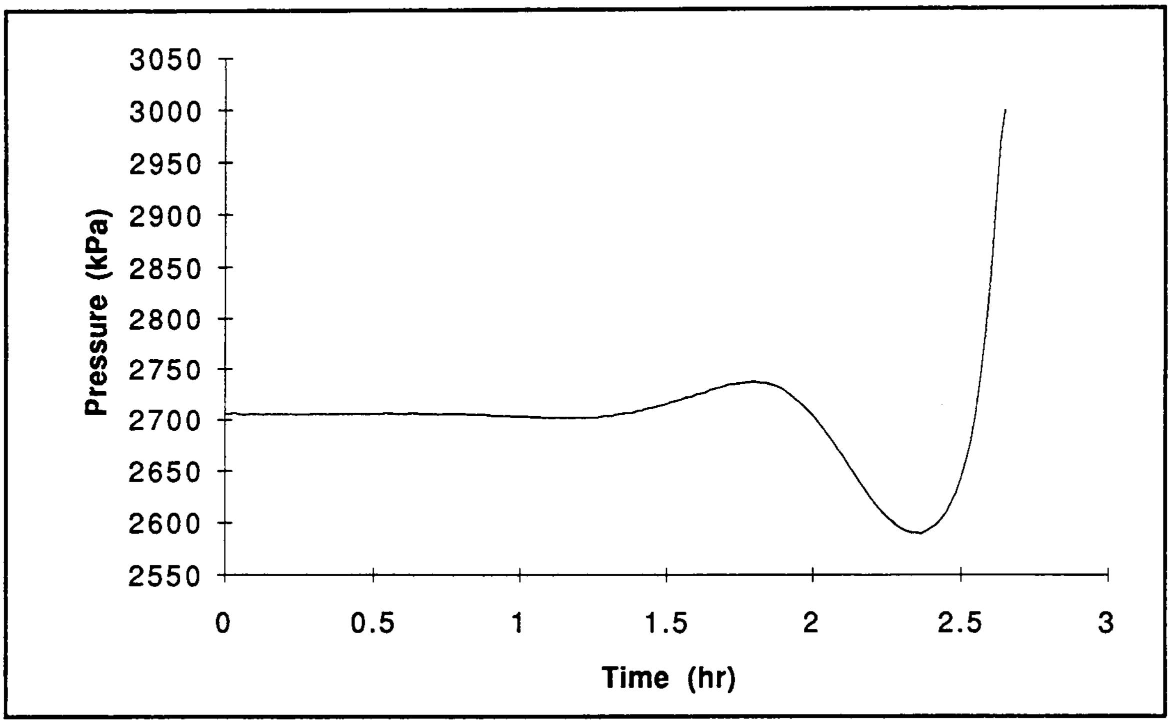

Figure 5.2: ECP Reactor Pressure Instability

5.4 An Impasse: Selecting Control Sets

After applying these basic flow and temperature controls, the plant still has 29 measurements and 9 inputs (shown in Table 5.1) available for control. Standard steady-state tools such as the Relative Gain Array require square gain matrices, but it is not at all obvious which subset of these variables should be chosen for control. None of the authors who have addressed this problem have provided any systematic approach to variable selection. Banerjee and Arkun (3) use a mathematical criterion that is able to account for only a few loops; other authors rely solely on their own judgment. Ricker (82) makes the following observations about variable selection:

- "Both MPC [model-predictive control] and decentralized methods require one to make critical decisions without quantitative justification. Foremost among these is the selection of the controlled variable set."
- "It is difficult to decide how many internal concentrations to control, which chemicals, which locations, etc. Existing quantitative methods for structure selection are inadequate for this purpose."

Instead of relying solely on judgment to choose sets of controlled and manipulated variables, it would be better to have some systematic criterion for deciding which variables might be appropriate to use for control. Such an approach will be presented in the following chapter.

Table 5.1: Variables Available for Control

Potential Manipulated Variables (9 Inputs)	Potential Controlled Variables (29 Measurements)
A feed setpoint	Reactor feed flow
E feed setpoint	Reactor level
A&C feed setpoint	Reactor pressure
Reactor coolant temperature setpoint	Reactor temperature
Separator coolant temperature setpoint	Separator temperature
Agitator power	Separator pressure
Purge flow setpoint	Compressor work
Steam flow setpoint	Stripper pressure
Recycle valve position	Stripper temperature
	Recycle flow
	Reactor feed %A
	Reactor feed %B
	Reactor feed %C
	Reactor feed %D
	Reactor feed %E
	Reactor feed %F
	Purge %A
	Purge %B
	Purge %C
	Purge %D
	Purge %E
	Purge %F
	Purge %G
	Purge %H
	Product %D
	Product %E
	Product %F
	Product %G
	Product %H

Chapter 6: Variable Set Selection

Many control strategy design tools such as the Relative Gain Array and Neiderlinski Index are calculated using the plant gain matrix. As Papadourakis et al. (73) point out, some basic material balance strategy must be in place before such a matrix can be derived, and the results obtained from these steady-state tools depend heavily on how the material balance controls are constructed. The studies with Module 2 primarily addressed how to choose the best such basic material balance control strategy. In many plants, however, there is an additional problem: which variables to use for control. If there is a physical plant already in place, then the variables may be limited to the valves and measurements already in place, whether or not they are sufficient to the task. If there is a plant model available, though, there are many more possible choices. A detailed model may allow the user to monitor every composition in every stream, every tray temperature in a distillation column, and many other kinds of variables. If the plant control strategy is designed as the plant is designed, then the plant design can be modified on paper to permit such a control strategy to be implemented during construction. Thus, proper selection of controlled and manipulated variables in the design phase can lead to a better plant design.

This chapter describes some concepts that can be used to reduce the variable set and give useful sets of controlled and manipulated variables. The interested reader may refer to Appendix C for more information on the mathematical details of this procedure.

6.1 Defining Classes of Variables

The amount of control that can be applied to a plant is sometimes measured in "degrees of freedom." It may be helpful to think of degrees of freedom as a measure of the number of independent directions in which the plant measurement vector can move. Each control loop (or set of cascaded loops) applied to the plant represents an attempt to control one of these directions--the existence of the loop implies a desire to be able to move one measurement or set of measurements without moving the others.

Given a plant with *n* measurements, *m* inputs, and *r* degrees of freedom, *r* can be no larger than the smaller of *n* and *m*. The goal of the variable set reduction method is to take the sets of measurements and inputs and identify a subset of *r* variables to control and a subset of *r* variables to manipulate. In linear algebra terms, if a plant has *r* degrees of freedom, then its measurement vector exists in an *r*-dimensional space, with *r* different independent directions of movement available. The challenge is to find measurements that best correspond to those directions, and to find inputs that move the plant in those directions.

In order to describe the problem of variable set reduction, we introduce two classes of plant variables: nonessential and equivalent.

- A <u>nonessential</u> variable is not useful for plant control. It does not correspond to any direction of significant plant movement.
- A group of <u>equivalent</u> variables is useful for control, but each of them represents essentially the same direction of plant movement.

Thus, there are really two levels of simplification associated with variable set selection: nonessential variables need to be eliminated from consideration, and equivalent variables need to be grouped.

These categories are neither exhaustive nor mutually exclusive. A member of an equivalent group might also be nonessential because it has such small sensitivity, or a variable might be essential and not equivalent to any other available measurement. These terms are introduced simply because they are useful for describing the procedure to follow.

6.1.1 Nonessential Variables

The operating state of the plant is observed, of course, by measuring the values of certain process variables. For control purposes, though, some measurements are not useful--if there are some measurements that are not affected by any manipulated variable, then those measurements are poor choices for controlled variables, because they cannot really be controlled. Likewise, it is possible that some process inputs do not have a significant effect on plant operation; these inputs are not appropriate choices for manipulated variables. Nonessential variables are those that should not be chosen in the control set.

As an extreme example of a nonessential measurement, consider the weather. Several weather-related factors (ambient temperature, rainfall, etc.) may have an effect on the operation of the plant, but (unfortunately) no plant input can change the weather. Weather measurements may exist, but they are "nonessential" in the sense that they cannot be controlled at all and therefore should not be chosen as controlled variables.

6.1.2 Equivalent Variables

In any plant with multiple inputs and outputs, it is possible that some variables represent essentially the same information. Several measurements might always move together in response to changes in plant inputs. Such variables cannot be independently controlled; moving one in some direction will automatically move the other in a predictable fashion. Similarly, there may be multiple plant inputs that have the same effect on the set of plant measurements. These inputs cannot be used to control different measurements. These equivalent variables should be grouped together; they cannot be used independently.

As an example of equivalent measurements, consider the temperatures on adjacent trays of a distillation column. Assuming that the trays are sufficiently removed from any feed or product stream, temperatures on adjacent trays are usually strongly coupled, and it is nearly impossible to move one temperature without changing the temperature on the adjacent tray. Adjacent tray temperatures, therefore, should not be chosen as separate controlled variables because they cannot be independently controlled.

It is important to remember that the plant gain matrix **K** will give information only about which measurements are equivalent with respect to the inputs. Within a group of "equivalent" measurements, it is possible that some subset might not always move together in response to the plant disturbances. When choosing equivalent measurements to use in a cascade configuration, then, one must be careful to careful to insure that disturbances do not unduly disturb the implied relationship between the measurements.

6.1.3 Using Nonessential and Equivalent Variables

The label "nonessential" does not imply that a variable is useless. Likewise, equivalence between two variables does not imply that either is superfluous. Some disturbances that enter the process may be reflected in the nonessential measurements (weather, for example). Even if those disturbances cannot be rejected, it may still be useful for the operating personnel to know about them so that they can better understand the behavior of the process. If two measurements are equivalent, they may provide a level of redundancy in case of sensor failure, or perhaps they can be combined into a single measurement to reduce the effects of noise in any single one of them. Alternatively, if two measurements have equivalent steady-state effects but differing dynamic effects, and the most important measurement is the slower one, then it might be appropriate to arrange them in a cascade configuration. Equivalent inputs might be manipulated together to provide more effective control action than either could provide individually. These possibilities will be demonstrated in the analysis of the Eastman Challenge Plant later in this work.

6.2 A Mathematical Approach to Variable Set Selection

6.2.1 Exact Solution Using the Inverse

As a background for a mathematical approach to variable set selection, it is helpful to review the simple case for which the plant gain matrix is invertible. Such problems are found frequently in the literature dealing with steady-state tools for control strategy design. Assuming some linearized model for the plant, the control problem may be written as a matrix equation:

$$\mathbf{Kx} = \mathbf{c}$$

where

$\mathbf{K}$ = set of steady-state gains; a square matrix of dimension n

$\mathbf{x}$ = manipulated variable (input) values; a vector of length n

$\mathbf{c}$ = controlled variable (output) values; a vector of length n.

With the $\mathbf{c}$ values specified (as setpoints), the function of the control system is simply to find the values for $\mathbf{x}$ that will satisfy the control equation. Whatever control strategy may be in place, the steady state will be reached at $\mathbf{x} = \mathbf{K}^{-1}\mathbf{c}$. As long as the gain matrix is square, invertible, and accurate, any appropriate control system will find its way to this steady-state point.

But what if $\mathbf{K}^{-1}$ does not exist? For a general plant gain matrix of dimension nxm, with rank r, there are at least 3 situations for which no inverse will exist:

- $n > m$, $r=m$. This case represents a plant with more measurements than inputs. There is no general solution to $\mathbf{Kx} = \mathbf{c}$; not all the variables can be controlled independently.
- $n < m$, $r=n$. This case represents a plant with more inputs than measurements. There is no unique solution to $\mathbf{Kx} = \mathbf{c}$; all the variables can be controlled, but more than one input vector $\mathbf{x}$ will move the process to the proper setpoint.
- $r < \min(n,m)$. Both the columns of $\mathbf{K}$ (inputs) and the rows of $\mathbf{K}$ (measurements) are dependent. In most cases, a solution does not exist. In those special cases for which a solution exists, it is not necessarily unique.

6.2.2 Optimum Least-Squares Solutions Using the Pseudoinverse

There is a mathematical approach to solving a system of equations such as $\mathbf{Kx} = \mathbf{c}$ when $\mathbf{K}^{-1}$ does not exist. The best solution is defined by the pseudoinverse $\mathbf{K}^{+}$. [See Strang (92) for a description of the pseudoinverse of a matrix.] The best solution is given by $\mathbf{x} = \mathbf{K}^{+}\mathbf{c}$. This solution is "best" in two senses:

- It is a least squares solution; that is, it minimizes the Euclidian norm of the error vector $\mathbf{e} = \mathbf{Kx} - \mathbf{c}$. (In control terms, this solution moves the set of controlled variables as close as possible to the set of setpoints.)
- If there is a set of vectors x that all give the optimum error vector, $\mathbf{K}^{+}\mathbf{c}$ is the shortest such vector. (In control terms, if there is more than one way to get to the target point, $\mathbf{K}^{+}\mathbf{c}$ is the way which requires the smallest change in the set of manipulated variables.)

6.2.3 The Pseudoinverse and Multi-Loop Control

If the process uses a multivariable controller, the pseudoinverse is an appropriate way to robustly control an overdetermined or underdetermined process. However, the information in the gain matrix and its pseudoinverse are also useful for identifying nonessential and equivalent variables. For this work, we propose that the diagonal elements of the matrix $\mathbf{KK}^+$ indicate which measurements are nonessential, and that the diagonal elements of the matrix $\mathbf{K}^+\mathbf{K}$ indicate which inputs are nonessential. For example, if the third element along the diagonal of $\mathbf{KK}^+$ is sufficiently small, then the measurement represented by the third row of $\mathbf{K}$ is nonessential--it does not change significantly when the inputs are manipulated, and so it cannot be effectively controlled. See Appendix C or more details and some justification for this method of identifying nonessential variables.

As demonstrated in Appendix C, there are different ways of identifying groups of equivalent variables. The important fact, though, is that all the variables within an equivalent group behave the same way. Within a group of equivalent measurements, moving one of them effectively moves them all in a predictable direction; within a group of equivalent inputs, manipulating one ha moves the plant in the same direction as manipulating any other one.

Once nonessential variables have been eliminated from consideration and groups of equivalent variables have been identified, the final manipulated and controlled variable sets may be chosen using the control engineer's best judgment.

- Equivalent inputs may be combined into a single manipulated variable in order to provide a greater range of control than any one of them could provide.
- From a set of equivalent measurements, the most appropriate choices for controlled variables are measurements that exhibit fast dynamic response or are closely related to plant constraints or quality measurements.

- Within a group of equivalent measurements, if one is related to an important quality objective but another exhibits much faster response, then a cascade scheme is appropriate, with the quality measurement in the "master" loop and the fast measurement in the "slave" loop.

This method of variable set selection will be demonstrated with the Eastman Challenge Plant. The method for constructing a multi-loop strategy is summarized in the following steps:

1. Once the basic material balance controls are in place, obtain **K**, the steady-state gain matrix for all the available inputs and measurements.
2. Scale the gain matrix using the physical spans of the inputs and sensors.
3. Using the singular value decomposition, determine *r*, the number of degrees of freedom available for control in the plant.
4. Construct $\mathbf{K}^+$, the rank-*r* pseudoinverse of **K**.
5. Using the diagonals of $\mathbf{KK}^+$ and $\mathbf{K}^+\mathbf{K}$, identify nonessential variables and eliminate them from consideration.
6. Using any appropriate method, segregate the remaining variables into equivalent groups.
7. Using engineering judgment, choose *r* controlled variables (one from each of *r* groups of measurements). If appropriate, use multiple variables within a set for cascade control.
8. Choose *r* groups of inputs as manipulated variables.
9. Construct a new "reduced" *rxr* gain matrix using the scaled gains for the controlled and manipulated variables chosen in Steps 7-8.
10. Apply existing steady-state analysis techniques to the reduced gain matrix to create the final multivariable or multi-loop control strategy.

11. Test the proposed control strategy by dynamic simulation.

If more than one variable set seems feasible from Steps 7-8, construct separate reduced gain matrices for each set. Each matrix then can be subjected to the analyses in Step 10.

6.2.4 Variable Set Selection and Multivariable Control

The control strategies designed in this study are multi-loop PID strategies. For some plants it may be appropriate to apply multivariable control based in some fashion on the inverse of the gain matrix. The techniques described in this chapter will still give valuable information in the design of a multivariable controller.

Skogestad and Morari (89) show that multivariable controllers based on the inverted gain matrix will give poor control for plants with large RGA values. Large RGA values indicate that the gain matrix is nearly singular; that is, some of the inputs or measurements are either nonessential or equivalent. For these plants, multivariable controllers will behave poorly if they can stabilize the plant at all. When the gain matrix is not invertible, the RGA cannot be calculated, but the same principle applies--a multivariable controller will be unable to control the plant properly if it presumes more degrees of control freedom than actually exist in the plant. In these cases, eliminating nonessential variables and grouping equivalent variables will improve the quality of multivariable control.

6.2.5 Demonstrating the Method

Appendix C contains some short numerical examples of the techniques described in this chapter for identifying nonessential and equivalent variable. On a larger scale, this procedure will be demonstrated in the next chapter on the Eastman Challenge Plant. Following the preliminary work on this plant in Chapter 6, 29 measurements and 9 inputs remain to be addressed. In the next chapter, these variable sets will be reduced to yield four different 7x7 control strategies.

Chapter 7: Alternative Control Strategies for the Eastman Challenge Plant

7.1 The Gain Matrix

7.1.1 Deriving and Scaling Gains

At the end of Chapter 5, the basic temperature and material balance controls had reduced the Eastman Challenge Plant (ECP) from a 41x12 problem, as it was originally stated, to a 29x9 problem. Chapter 6 and Appendix C present a rationale for eliminating and combining some of these variables to reduce the problem size even further. In order to apply this procedure, we must derive the 29x9 gain matrix for the plant. (These will be gains for the inputs, not the disturbances. Although there are mathematical procedures for using disturbance gains, the Eastman Challenge Plant is not presented in a format that permits calculating disturbance gains.)

Small step tests in both the positive and negative directions were performed for each of the 9 inputs. To insure sufficient linearity, the tests were redone with a smaller step size if the step tests in the two directions gave gains that differed by more than about 10%.

Physical scaling of a gain usually involves multiplying the gain by some significant range of the input (usually the total span of the input) and dividing by some significant range of the output (usually the sensor span or the base-case value of the output). This procedure gives gain dimensions of (%output)/(%input). The scaling factors used for the ECP gains are shown in Tables 7.1 and 7.2, and the scaled gains are shown in Table 7.3.

Table 7.1: Scaling Factors for Inputs

Input	Significant Range
A feed setpoint	1.017
E feed setpoint	8354
A&C feed setpoint	15.25
Agitator power	100
Recycle valve	100
Steam flow setpoint	(Special)
Purge flow setpoint	(Special)
Reactor coolant temp setpoint	(Special)
Separator coolant temp setpoint	(Special)

Note: Inputs labeled "Special" above are setpoints for which it is difficult to estimate a range between upper and lower limits. For these cases, the gain was calculated relative to the valve position necessary to achieve the given setpoint.
This substitution is equivalent to assuming that the quantity

(HiSetpoint-LowSetpoint)/SetpointRange

is approximately equal to

±(HiValvePosition - LowValvePosition)/ValveRange

with the sign dependent on whether the loop is direct or reverse acting. This substitution is exact for linear systems, and approximate for other systems. The valve ranges were 0-100% for all valves.

Table 7.2: Scaling Factors for Measurements

Measurement	Significant Range
Reactor feed flow	84.68
Reactor level	100.00
Reactor pressure	2705.00
Reactor temperature	100.00
Separator temperature	100.00
Separator pressure	2705.00
Compressor work	682.86
Stripper pressure	2705.00
Stripper temperature	100.00
Recycle flow	53.80
Reactor feed %A	100.00
Reactor feed %B	100.00
Reactor feed %C	100.00
Reactor feed %D	100.00
Reactor feed %E	100.00
Reactor feed %F	100.00
Purge %A	100.00
Purge %B	100.00
Purge %C	100.00
Purge %D	100.00
Purge %E	100.00
Purge %F	100.00
Purge %G	100.00
Purge %H	100.00
Product %D	100.00
Product %E	100.00
Product %F	100.00
Product %G	100.00
Product %H	100.00

Table 7.3: Scaled Gains for the Eastman Challenge Plant

	A feed Setpoint	E feed Setpoint	A&C feed Setpoint	Agitator Power	Rx Coolant Temperature	Sep Coolant Temperature	Purge Flow	Steam Flow	Recycle Valve
Rx feed flow	-0.556	2.244	-3.467	0.125	1.192	-1.472	0.011	0.015	0.139
Rx level	-9.429	46.302	-83.456	-0.580	-5.551	-15.662	1.828	0.233	3.763
Rx pressure	-3.410	11.041	-18.323	0.089	0.839	-4.897	0.131	0.063	1.187
Rx temp	0.214	-0.627	1.525	-0.220	-1.991	0.260	-0.021	-0.008	-0.050
Sep temp	1.971	-6.297	10.218	-0.357	-3.412	1.403	-0.070	-0.040	-0.707
Sep press	-3.354	10.848	-18.041	0.089	0.822	-4.801	0.130	0.063	1.198
Comp work	0.113	0.382	-0.391	0.123	1.180	-1.103	-0.019	0.008	0.466
Str press	-3.741	12.122	-20.155	0.096	0.907	-5.329	0.145	0.070	1.113
Str temp	1.902	-7.038	11.704	-0.273	-2.611	1.325	-0.107	0.106	-0.610
Recyc flow	-0.813	2.950	-4.889	0.197	1.884	-2.293	0.007	0.020	0.206
Rx feed XA	0.687	-0.209	0.305	0.089	0.852	-0.494	0.015	0.002	0.167
Rx feed XB	-0.042	0.117	-0.092	0.013	0.119	-0.135	-0.221	0.001	0.011
Rx feed XC	-1.071	-0.585	1.754	0.019	0.176	-0.273	0.141	-0.001	0.123
Rx feed XD	0.033	-0.047	-0.053	-0.004	-0.036	0.170	0.008	-0.001	-0.018
Rx feed XE	-0.275	2.924	-5.643	-0.012	-0.106	0.169	0.123	0.008	-0.070
Rx feed XF	0.177	-0.475	0.885	-0.052	-0.498	0.143	-0.043	0.000	-0.048
Purge XA	1.016	0.000	-0.153	0.137	1.306	-0.753	0.031	0.005	0.266
Purge XB	0.000	0.000	0.229	0.000	0.000	0.000	-0.347	0.000	0.000
Purge XC	-1.758	-0.470	1.906	0.038	0.361	-0.545	0.228	0.003	0.210
Purge XD	-0.043	0.316	-0.671	0.015	0.146	0.008	0.015	0.001	-0.003
Purge XE	-0.219	3.446	-6.863	-0.031	-0.301	0.390	0.164	0.008	-0.150
Purge XF	0.270	-0.740	1.365	-0.075	-0.720	0.240	-0.059	-0.001	-0.074
Purge XG	0.486	-1.710	2.772	-0.055	-0.527	0.458	-0.020	-0.010	-0.166
Purge XH	0.244	-0.825	1.414	-0.028	-0.265	0.225	-0.012	-0.005	-0.082
Product XD	-0.003	0.013	-0.024	0.000	0.004	-0.002	0.000	0.000	0.001
Product XE	-0.122	0.573	-1.011	0.011	0.101	-0.080	0.013	0.000	0.029
Product XF	-0.001	0.017	-0.024	-0.002	-0.018	-0.001	-0.002	0.000	0.001
Product XG	0.000	-0.365	-0.229	-0.004	-0.037	0.026	0.021	0.000	-0.008
Product XH	0.126	-0.230	1.258	-0.004	-0.046	0.052	-0.033	0.000	-0.024

7.1.2 Identifying Degrees of Freedom

After scaling the gain matrix, the next step is to decide how many degrees of freedom should be retained for control. The number of degrees of freedom (DOF) defines the number of independent loops that can be applied to the plant. Mathematically, the number of DOF is defined by the number of sufficiently large singular values of the gain matrix. An SVD of the gain matrix in Table 7.3 gives the singular values shown in Table 7.4. Because the gain matrix is scaled to units of (%output)/(%input), the singular values are scaled as well. A singular value around 1.0 indicates a direction in which a change of some percentage in the value of the inputs causes a change of equal percentage in the value of the outputs. Large values indicate directions in which the outputs are very sensitive to the inputs; small values indicate directions in which the outputs are insensitive to the inputs. Two things are immediately obvious from Table 7.4: the first singular value is very large, and the last (ninth) singular value is very small.

Table 7.4: Singular Values for the ECP Gain Matrix

Position	Singular Value
1	107.48
2	7.79
3	4.16
4	2.71
5	0.82
6	0.64
7	0.45
8	0.11
9	0.01

Table 7.5 shows the first left and right singular vectors of the gain matrix for the ECP. (These are the vectors associated with the large first singular value.) The u_1 vector is dominated by the 2nd measurement, the reactor level. The v_1 vector is composed primarily of the 3rd input (the A&C feed rate), with some contribution from the 2nd input (the E feed rate). The large singular value shows that the reactor level is extremely sensitive to these two flows; we might expect closed-loop control difficulties if we use some other input to control the reactor level. (In a real plant this large singular value might also predict difficulties with valve resolution. For example, if we use the A&C feed to control the reactor level, changing the reactor level by 1% would require a steady-state change of only about .01% in the A&C flow rate, and it is doubtful that any normal control valve can control with such precision.)

Some judgment is required to determine how small a singular value must be to be considered zero and eliminated from consideration. In general, an input direction is insensitive if its singular value is substantially less than 1, although the sizes of the other singular values may affect the decision as well. For example, if the singular values are [2, 1.5, 1.3, 0.9, 0.7, 0.1], then the last singular value is probably too small to retain. However, if the singular values are [0.3, 0.2, 0.15, 0.1], then perhaps a singular value of 0.1 is not too small.

For the ECP gain matrix, the ninth singular value is sufficiently small that it may be safely considered to be zero, but the eighth singular value is questionable. Because it is considerably smaller (by a factor of 4) than the next highest singular value, and because it is small compared to 1, it is probably too small to retain. Thus we conclude that the ECP gain matrix has 7 effective degrees of freedom.

Table 7.5: First Singular Vectors for the ECP Gain Matrix

U1	V1
-0.0400	0.1054
-0.9061	-0.4797
-0.2059	0.8533
0.0149	0.0035
0.1125	0.0333
-0.2026	0.1663
-0.0062	-0.0166
-0.2262	-0.0023
0.1277	-0.0402
-0.0558	
0.0035	
-0.0014	
0.0150	
0.0001	
-0.0579	
0.0094	
-0.0011	
0.0019	
0.0147	
-0.0067	
-0.0695	
0.0146	
0.0307	
0.0154	
-0.0002	
-0.0108	
-0.0003	
-0.0002	
0.0112	

7.2 Reducing the Variable List

7.2.1 Eliminating Nonessential Variables

The variable selection procedure from Chapter 6 may be applied now that the degrees of freedom have been determined. The diagonals of $\mathbf{KK}^+$ and $\mathbf{K}^+\mathbf{K}$, where $\mathbf{K}^+$ is the pseudoinverse which retains 7 degrees of freedom, are shown in Table 7.6.

The input list is easy to evaluate. There are seven inputs with diagonal values very close to 1, and two with diagonal values very close to 0. The steam flow and agitator power, with small diagonal values, are nonessential and may be eliminated from the list of potential manipulated variables. (Further analysis reveals that the agitator power is equivalent in effect to the reactor coolant exit temperature. The agitator power is nonessential, though, because its effect is so much smaller than that of the reactor coolant.)

In practice, Buckley (8) recommends establishing ratio controls to move several streams together, and then using the ratio setpoints as the actual manipulated variables. To that end, the E feed setpoint will be ratioed to the D feed setpoint. In addition, in imitation of the reaction stoichiometry, the A feed and A&C feed setpoints will be ratioed to the total of the D and E feed setpoints. (We ratio to the feed setpoints instead of the flow measurements to minimize noise effects in the flow measurements.) These ratios will be the actual manipulated variables in the control strategies.

The measurement list is only slightly more difficult to evaluate than the input list. It is obvious from Table 7.6 that the seven degrees of freedom are spread out among many different measurements. Only the reactor level has a diagonal value close to 1. There is a clear distinction, though, between the essential and nonessential measurements--there are 9 measurements with diagonal values less than 0.02; the next lowest diagonal value is 0.137. These nine measurements, shown in Table 7.7, may be eliminated from the list of potential controlled variables.

Table 7.6: Diagonals of $\mathbf{KK}^+$ and $\mathbf{K}^+\mathbf{K}$ for the ECP

Plant Variable	Corresponding Diagonal Value
Inputs:	
A feed SP	1.000
E feed SP	0.999
A&C feed SP	1.000
Agitator power	0.011
Reactor coolant temp SP	0.989
Separator coolant temp SP	0.999
Purge flow SP	1.000
Steam flow SP	0.002
Recycle valve	1.000
Outputs:	
Reactor feed flow	0.235
Reactor level	0.963
Reactor pressure	0.231
Reactor temp	0.322
Separator temp	0.472
Separator press	0.233
Compressor work	0.307
Stripper press	0.333
Stripper temp	0.398
Recycle flow	0.507
Reactor feed %A	0.173
Reactor feed %B	0.137
Reactor feed %C	0.225
Reactor feed %D	0.006
Reactor feed %E	0.157
Reactor feed %F	0.018
Purge %A	0.392
Purge %B	0.335
Purge %C	0.527
Purge %D	0.006
Purge %E	0.283
Purge %F	0.037
Purge %G	0.017
Purge %H	0.004
Product %D	0.000
Product %E	0.002
Product %F	0.000
Product %G	0.352
Product %H	0.329

Table 7.7: Nonessential Variables in the ECP

Inputs	Outputs
Agitator power	Reactor feed %D
Steam flow setpoint	Reactor feed %F
	Purge %D
	Purge %F
	Purge %G
	Purge %H
	Product %D
	Product %E
	Product %F

Identifying nonessential variables can provide some insight into the process capabilities. For example, we observe that the only product compositions that can be affected by the inputs are the mole fractions of G and H, the desired products. The inputs will not significantly affect the composition of F, the undesired product. If it becomes necessary to reduce byproduct formation, or to reduce the fraction of raw materials in the product, the current plant design provides no way to do so near the base case operating conditions. Also, the stripper steam flow and the agitator speed have no significant steady-state effect on the plant operation. It is possible that such knowledge could be useful during plant design; perhaps the steam supply or the variable-speed drive for the agitator could be eliminated, assuming they were included in the design only for control purposes.

7.2.2 Grouping Equivalent Variables

After deleting nonessential variables, the ECP gain matrix is now a 20x7 matrix. The inputs do not need to be grouped; with 7 degrees of freedom and 7 inputs available, there should be no redundancy in the inputs. (If there were, the seventh singular value would have been smaller.) The measurements, though, can be grouped by either of the two methods outlined in Appendix B. Using the method of inflated gains suggests the variable groupings shown in Table 7.8.

Note that there are more groups in Table 7.8 than there are degrees of freedom. This phenomenon is not hard to explain; consider the vectors shown in Figure 7.1. There are only two dimensions on the graph, but there are three obvious groupings of vectors. Although there are only two independent directions, there are three groups of variables, any two of which could be used for control. (Of course, all the variables in all three groups could be used in a multivariable controller, but the multivariable controller could still control in no more than two directions.)

7.2.3 Selecting Variable Sets

By now we have reduced the Eastman Challenge Plant to a 7x9 problem with 7 control loops desired. There are at most 72 possible control strategies that could be formed from these variables, as opposed to the 5.5 trillion possibilities from a 29x9 system. Although some exhaustive search technique is now feasible, we will continue by other methods.

Seven controlled variables must be chosen to pair with the manipulated variables. Of the nine groups represented in Table 7.8, only seven may be chosen, and no more than one variable should be chosen from each group. This step in the selection process depends largely upon the judgment of the control engineer. Some selections are simple, based on plant objectives:

- The reactor level, in group 2 by itself, is an obvious choice. (Remember that the reactor level dominates the first singular vector of the plant gain matrix.)
- The reactor pressure should be controlled because of the constraint on it.

Table 7.8: Equivalent Measurements

Group	Variables
1	Reactor feed flow rate, recycle flow rate
2	Reactor level
3	Reactor pressure, separator pressure, stripper pressure
4	Reactor temperature, product %G, product %H
5	Separator temperature, stripper temperature
6	Reactor feed %A, purge %A
7	Reactor feed %B, purge %B
8	Reactor feed %C, purge %C
9	Reactor feed %E, purge %E

Note: Compressor power was weakly correlated with several of these groups, but not strongly correlated with any of them. It might be included as a separate group if it were an important variable to control.

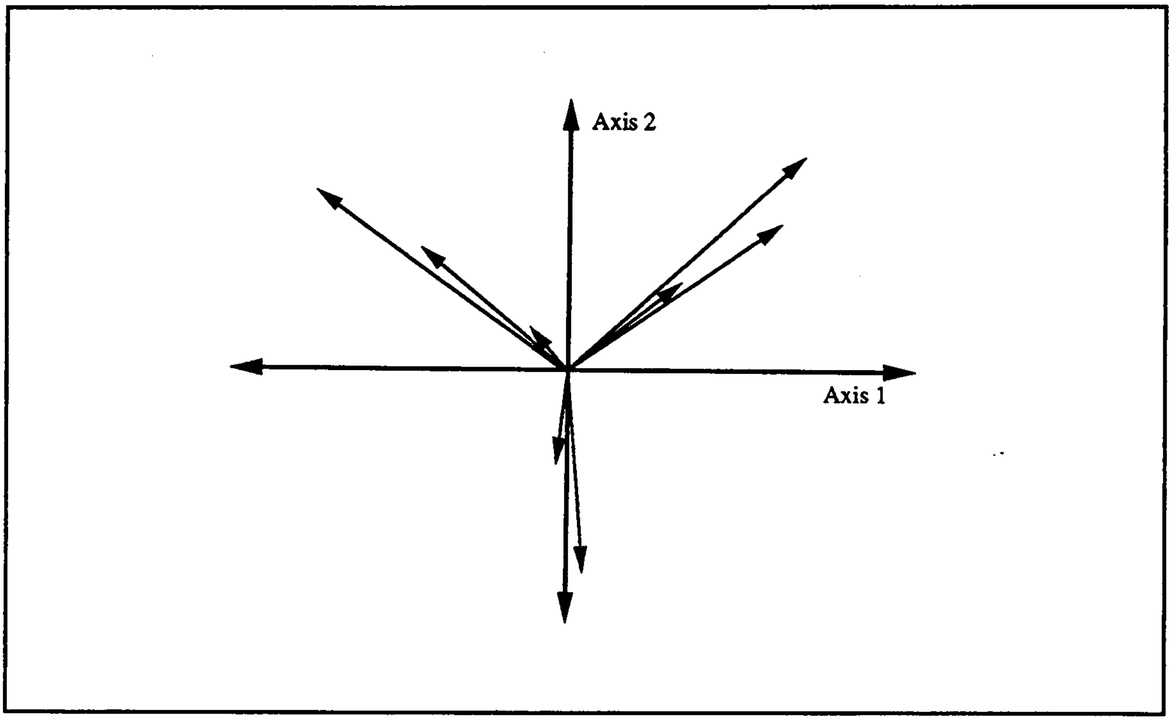

Figure 7.1: Three Variable Groups in Two Dimensions

- The product mix (the split between G and H) must be controlled to move between the various required operating modes given by Downs and Vogel (15); for our initial control strategies, we choose to control the %G in the product stream, with the understanding that we may later use the %G and %H as a single variable and control (%G)/(%H).
- The stripper temperature and separator temperatures (which are equivalent variables) affect the product separation. We choose to control the stripper temperature because it is closer to the product.
- Downs and Vogel (15) require in the original problem statement that the mole fraction of 'B' in the purge stream must be controllable; one of the tests that they propose for control strategy evaluation is a change in that composition. The %B in the purge is therefore a necessary choice for a controlled variable.

After specifying these five controlled variables, two remain to be chosen. Several criteria might be considered as a basis for the final choices: component inventory control requirements, best singular values, quickest dynamic response, and propagation of variation.

7.2.3.1 Component Inventory Analysis

Because the Eastman Challenge Plant contains a recycle loop and the purge stream is relatively small, there is the potential for any of the light components (A, B, C, D, E, or F) to accumulate. Based on the results of the variable selection process, only the inventories of A, B, C, and E are adequately controllable. The inventory of B has been chosen as a controlled variable; other inventories may be chosen for control based on the process stoichiometry.

The process stoichiometry is based on the following reactions:

$$A(g) + C(g) + D(g) \rightarrow G(liq)$$

$$A(g) + C(g) + E(g) \rightarrow H(liq)$$

$$A(g) + E(g) \rightarrow F(liq)$$

$$3D(g) \rightarrow F(liq)$$

The third and fourth reactions are byproduct reactions which do not consume nearly as many moles as the reactions that produce G and H. Therefore, the stoichiometric considerations will be based only on the first two reactions.

Reactants A and C are consumed together in the product reactions. Neither reactant exits the process in large amounts; the only exit is via reaction. It will be important to insure that A and C are fed in the proper ratio. The difficulty in the Eastman Challenge Plant is that A and C both enter the plant primarily in the same feed stream. There is a small A feed as well, but it is not always available. (Disturbance #6 is the loss of the A feed.) Unfortunately, manipulating this A feed is the only way to affect the relative feed rates of A and C.

Reactant D is consumed in the reaction that makes product G; reactant E is consumed in the reaction that makes product H. In order to maintain the proper product mix, it is necessary to maintain the D and E feeds in the proper ratio. The product %G (and, by implication, the G/H mix) has already been selected as a controlled variable; monitoring the E inventory might smooth the control, but it should not be essential, provided that the product composition controller is directly connected to the E feed. Under most circumstances, an imbalance in the G/H ratio in the product should indicate a problem with the D/E ratio in the process. The exception might be some disturbance that affects the reaction temperature, which has <u>not</u> been chosen as a controlled variable. The reactions have significantly different activation energies, so reaction temperature changes might disturb the predicted relationship between E inventory and product composition.

The Module 2 analysis suggested that the components in excess might be preferred as component inventory measures. The compositions of the recycle and purge streams shows that component A exists in the greatest excess in the plant.

The inventories of A, C, and E may be measured by analyzers on the purge stream or the reactor feed stream. Because the purge composition measurements for each component are equivalent to the reactant feed composition measurements, dynamic or instrumentation considerations will dictate where these compositions should be measured if they are chosen as controlled variables. The available analyzers operate at the same frequency, but measuring compositions in the purge stream might be preferred for economic reasons--assuming that one analyzer can provide measurements for several components, measuring all the compositions from the purge stream might keep instrumentation costs lower. On the other hand, the analyzer in the reactor feed stream is immediately downstream of the A and E feeds that would be used for inventory control; that analyzer might provide better dynamic response. For this work, purge compositions will be chosen arbitrarily instead of reactor feed compositions.

Based only on stoichiometric analysis, we conclude:

- Control the A inventory or the C inventory to insure that A and C are fed in the proper ratio. The A feed stream is the only one available for changing this feed ratio; when it is not available, upsets in the composition of the A&C feed stream could very possibly shut the plant down. The A inventory might be preferred for control because the A feed is the makeup stream and because A exists in the greatest excess in the plant.
- If the product composition controller manipulates the E feed, there should be no need under most circumstances to control the E inventory. Problems might arise with the E inventory if disturbances significantly change the reaction temperature.

- Component inventory measurements can be made in the purge stream or the reactor feed. There is no steady-state reason to prefer one over the other.

7.2.3.2 Singular Value Analysis

The four remaining candidates for controlled variables are recycle flow (or its equivalent, the reactor feed flow), purge %A, purge %C, and purge %E. There are six possible ways to choose two variables from this set of four, and there is a 7x7 gain matrix corresponding to each of these six choices. From one perspective, the best choice of controlled variables is the set that gives the largest minimum singular value. Table 7.9 shows the singular values for the six possible gain matrices. The first six singular values are approximately the same regardless of which two variables are controlled. The last singular value (shown in bold in Table 7.9) is significantly different from one variable set to the next.

Table 7.9: Singular Values for Reduced (7x7) Variable Sets

Final Two Controlled Variables	Singular Values 1	2	3	4	5	6	7
Recycle flow, purge %E	101.3	4.73	2.68	0.89	0.40	0.32	**0.19**
Recycle flow, purge %C	101.0	4.83	2.52	1.49	0.55	0.32	**0.11**
Recycle flow, purge %A	101.0	4.85	2.07	1.60	0.49	0.31	**0.26**
Purge %E, purge %C	101.1	4.30	2.98	1.17	0.44	0.32	**0.14**
Purge %E, purge %A	101.1	4.27	2.57	1.54	0.43	0.32	**0.14**
Purge %C, purge %A	100.8	4.36	2.62	1.67	0.54	0.36	**0.09**

7.2.3.3 Dynamic Analysis

Of the four measurements remaining, from which we must choose our final two controlled variables, three are compositions and one is a flow rate. In general, compositions respond more slowly than flow rates to disturbances and inputs. In addition, the compositions are not measured continuously and therefore might not respond as quickly as the recycle flow. Therefore, dynamic considerations favor using the recycle flow as one of the controlled variables.

7.2.3.4 Propagation of Variation Analysis

We have previously identified the reactor (and specifically the reactor pressure) as the critical point in the plant, the point away from which we should direct variation. Controlling the recycle flow will probably help stabilize the reactor pressure, depending on which input is used to control the recycle. However, compositions in the plant will likely have a significant effect on the reactor pressure because of the complicated vapor-liquid equilibrium in the two-phase reaction mixture. Therefore it is difficult to choose intuitively among the possible controlled variables based on the propagation of variation.

7.2.3.5 Choosing the Final Two Controlled Variables

Based on the above considerations, the choices for the final two controlled variables seem fairly obvious: recycle flow and purge %A. This combination addresses the most pressing component inventory problem (the A-to-C ratio), yields the largest minimum singular value, and uses the fast-responding recycle flow. Depending on the dynamic simulations which will be run later, other possibilities might be acceptable:

- recycle flow and purge %E (if the E inventory is problematic)
- purge %A (or %C) and purge %E, depending on which inventories are problematic.

Thus we now have a preferred set of controlled variables and some alternatives for comparison. The same seven manipulated variables will be used for all control strategies: A feed ratio setpoint, E feed ratio setpoint, A&C feed ratio setpoint, reactor coolant temperature setpoint, separator coolant temperature setpoint, purge flow setpoint, and recycle valve position. There are four possibilities for the controlled variable set; these sets are listed in Table 7.10. Strategy 1 is recommended by our design procedure; the other strategies are alternatives in case particular problems develop during dynamic testing. For example, if Strategy 1 has problems with the E inventory after certain disturbances, perhaps Strategy 2 or 3 will perform better.

7.2.3.6 Conclusions of Other Authors

Some of our conclusions from the variable selection process can be compared to decisions made by other authors who have addressed the Eastman Challenge Plant. McAvoy and Ye (56) were the first to note that the agitator power and the reactor coolant flow were equivalent in effect. [Banerjee and Arkun (3) noted the same fact, but their analysis was unable to provide justification for choosing one of these variables over the

Table 7.10: Possible Controlled Variable Sets

Strategy 1	Strategy 2	Strategy 3	Strategy 4
Reactor level	Reactor level	Reactor level	Reactor level
Reactor pressure	Reactor pressure	Reactor pressure	Reactor pressure
Product %G	Product %G	Product %G	Product %G
Stripper temp	Stripper temp	Stripper temp	Stripper temp
Purge %B	Purge %B	Purge %B	Purge %B
Recycle flow	Recycle flow	Purge %A	Purge %C
Purge %A	Purge %E	Purge %E	Purge %E

other one. Our variable selection procedure identifies the coolant flow as the preferred choice.] Ricker (81) observed that the steam flow had no significant effect on the plant operation and would be best turned off. Vinson et al. (97) attempted to use steam flow to control the %E in the product and reported that control was poor. Our variable selection procedure verifies all of these observations--the agitator power and reactor coolant are equivalent; the steam flow, agitator power, and product %E are nonessential.

In other ways, though, the variable selection procedure differs from decisions by other authors. Several authors attempt to control reactor temperature, reactor pressure, and product composition simultaneously. McAvoy et al. (57) differ slightly when they cascade reactor pressure control onto the temperature controller. Our variable selection procedure suggests that there are only two degrees of freedom among those three variables and that reactor temperature, if controlled, should be tied to the product composition. McAvoy et al. (57) report that product composition control worsened somewhat when they applied their cascade controller, as we might expect if the reactor temperature and product composition are closely related.

7.3 Creating Controllers

7.3.1 Defining the Alternatives

Twelve control loops were created to stabilize the plant before deriving the gain matrix. These control loops, described in Table 7.11, were implemented for all the tested control strategies. The subsequent loops were designed using steady-state tools.

The Relative Gain Array has gained wide acceptance as a tool to help pair controlled with manipulated variables once the sets have been chosen. While the RGA cannot indicate definitively which pairing configuration is the best one, it can indicate (by very small or

Table 7.11: Control Loops Common to All Strategies

Controlled Variable	Manipulated Variable
A feed rate	A feed valve
D feed rate	D feed valve
E feed rate	E feed valve
A&C feed rate	A&C feed valve
Purge flow rate	Purge valve
Separator underflow rate	Separator underflow valve
Product flow rate	Product flow valve
Stripper steam flow rate	Stripper steam valve
Reactor coolant exit temp	Reactor coolant valve
Separator coolant exit temp	Separator coolant valve
Stripper level	Product flow rate setpoint
Separator level	Separator underflow setpoint

negative values) which pairings might be undesirable. So far we have identified four different possible variable sets; the RGA matrices for these sets are shown in Table 7.12.

Some pairings are obvious from the RGA or stoichiometric considerations. The purge %B can be effectively controlled only by the purge flow rate. The reaction stoichiometry requires a change in the D/E feed ratio to effect a change in the product G/H ratio, so the E feed ratio setpoint is the logical choice to control the product composition. Likewise, the purge %A (or perhaps %C, depending on which was chosen to indicate the relative inventories of A and C) must be ultimately controlled by the A feed ratio setpoint, so it makes sense to assign that variable pairing for those strategies that control the A or C inventory. The recommended variable pairings for each control strategy are indicated by italics in the RGA matrices in Table 7.12. A diagram of Strategy 1 is shown in Figure 7.2.

The RGA for Strategy 2 deserves special mention, because the variable pairings recommended in Table 7.12 do not seem straightforward. Given only the RGA matrix, we would almost certainly choose to control the reactor pressure with the A feed ratio. However, Downs and Vogel (15) indicate that the A feed is sometimes not available for

Table 7.12: RGA Matrices for Alternative Control Strategies

(a) Strategy 1

	A feed ratio	E feed ratio	A&C feed ratio	Reactor coolant temperature	Separator coolant temperature	Purge flow setpoint	Recycle valve position
Reactor level	-0.80	0.86	<u>1.63</u>	0.13	0.07	-0.02	-0.88
Reactor pressure	1.34	-0.98	-1.67	-0.06	0.38	0.01	<u>1.98</u>
Product %G	0.00	<u>0.78</u>	0.25	-0.01	0.00	-0.01	-0.01
Stripper temp	0.04	0.24	0.63	<u>0.52</u>	-0.44	0.00	0.01
Purge %B	0.00	0.00	-0.03	0.00	0.00	<u>1.03</u>	0.00
Recycle flow	-0.08	0.09	0.19	0.34	<u>0.97</u>	0.00	-0.51
Purge %A	<u>0.50</u>	0.00	0.00	0.08	0.02	0.00	0.40

(b) Strategy 2

	A feed ratio	E feed ratio	A&C feed ratio	Reactor coolant temperature	Separator coolant temperature	Purge flow setpoint	Recycle valve position
Reactor level	-3.18	0.02	<u>1.98</u>	-0.04	0.24	-0.03	2.01
Reactor pressure	5.65	0.03	-2.05	-0.19	<u>0.12</u>	0.01	-2.58
Product %G	0.00	<u>0.78</u>	0.25	0.00	0.00	-0.01	-0.01
Stripper temp	-1.18	-0.08	0.75	0.72	-0.40	0.00	<u>1.20</u>
Purge %B	0.00	0.00	-0.04	0.00	0.00	<u>1.04</u>	0.00
Recycle flow	-0.54	-0.04	0.24	<u>0.47</u>	1.03	0.00	-0.15
Purge %E	<u>0.25</u>	0.29	-0.13	0.04	0.02	0.00	0.53

(continued on next page)

Table 7.12 (continued)

(c) Strategy 3

	A feed ratio	E feed ratio	A&C feed ratio	Reactor coolant temperature	Separator coolant temperature	Purge flow setpoint	Recycle valve position
Reactor level	-0.40	0.28	0.19	<u>0.57</u>	-2.85	0.00	3.22
Reactor pressure	0.61	-0.27	-0.08	0.27	<u>4.96</u>	0.00	-4.49
Product %G	0.00	<u>0.78</u>	0.25	-0.01	0.00	-0.01	0.00
Stripper temp	0.25	0.02	0.11	-0.01	-1.07	0.00	<u>1.70</u>
Purge %B	0.00	0.00	-0.03	0.00	0.00	<u>1.03</u>	0.00
Purge %A	<u>0.59</u>	0.00	0.01	0.28	0.29	0.00	-0.17
Purge %E	-0.04	0.20	<u>0.54</u>	-0.11	-0.33	-0.01	0.75

(d) Strategy 4

	A feed ratio	E feed ratio	A&C feed ratio	Reactor coolant temperature	Separator coolant temperature	Purge flow setpoint	Recycle valve position
Reactor level	-0.18	0.30	0.04	<u>0.62</u>	-3.10	0.00	3.32
Reactor pressure	0.17	-0.30	0.09	0.31	<u>5.38</u>	0.00	-4.66
Product %G	0.00	<u>0.78</u>	0.25	-0.01	0.00	-0.02	0.00
Stripper temp	-0.16	-0.01	0.30	0.20	-0.88	0.00	<u>1.55</u>
Purge %B	0.00	0.00	-0.02	0.00	0.00	<u>1.02</u>	0.00
Purge%C	<u>1.08</u>	-0.01	0.08	-0.08	-0.22	0.01	0.14
Purge %E	0.09	0.24	<u>0.24</u>	-0.04	-0.18	0.00	0.65

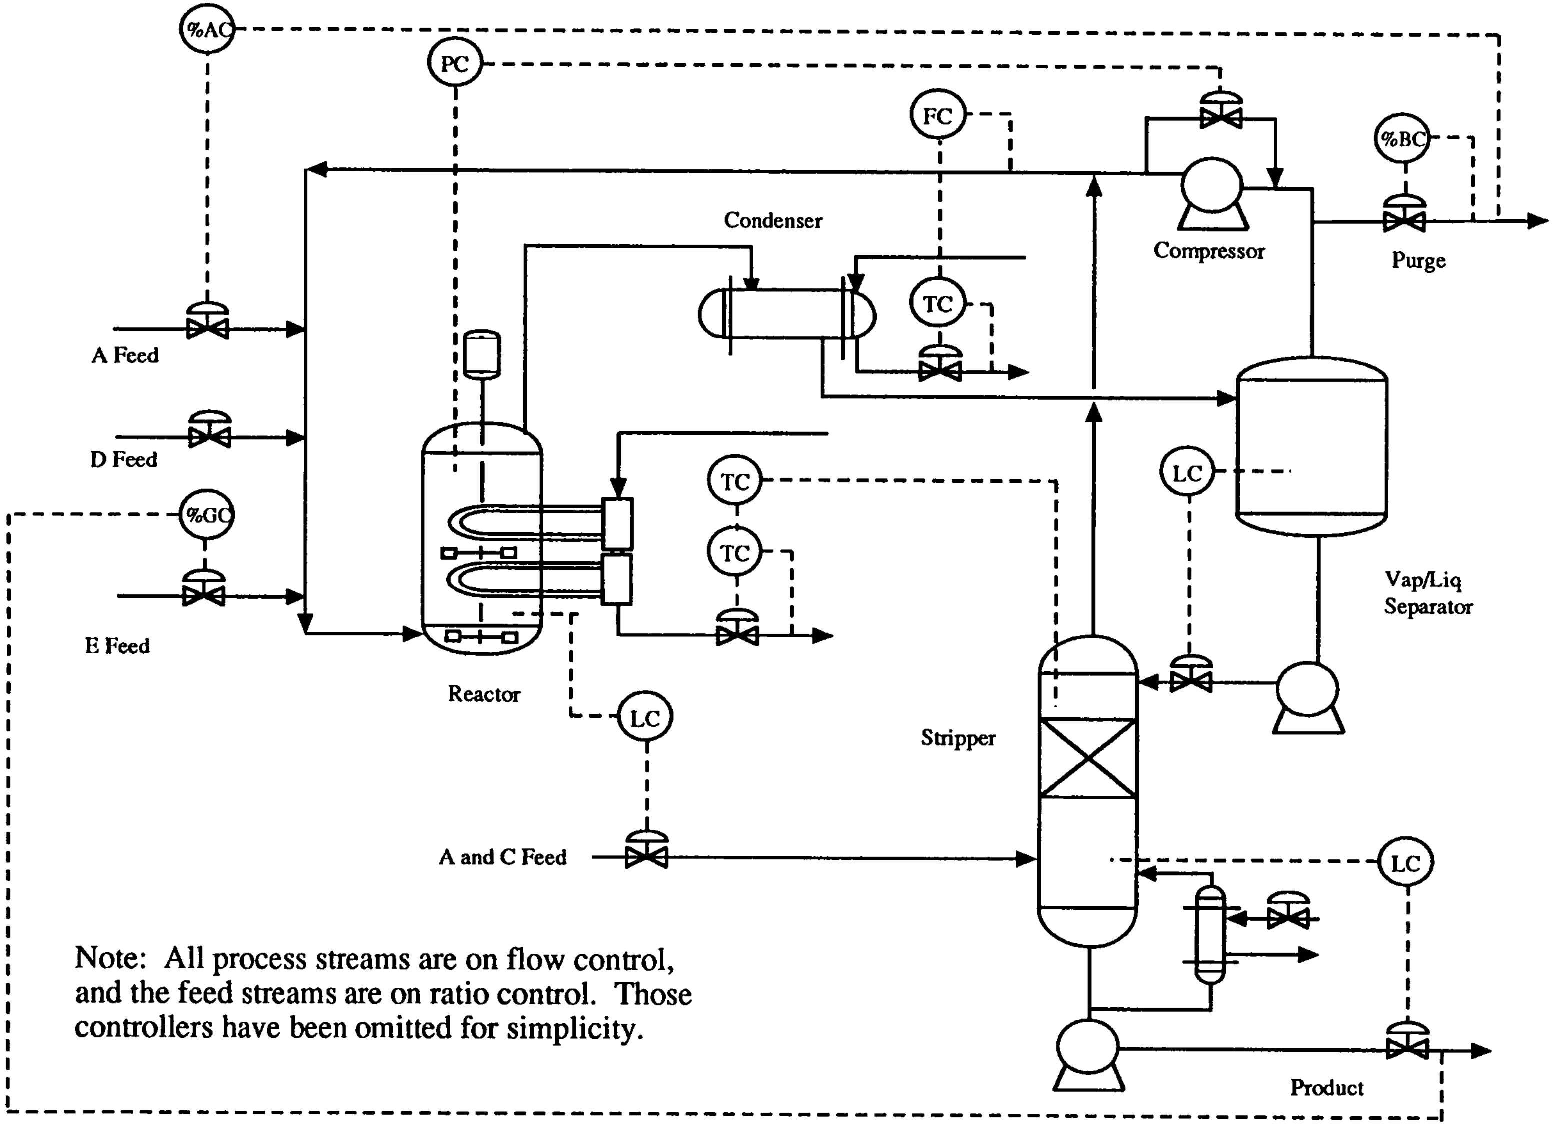

Note: All process streams are on flow control, and the feed streams are on ratio control. Those controllers have been omitted for simplicity.

Figure 7.2 ECP Control Strategy 1

control. We should not use the A feed, then, to control a variable as critical as the reactor pressure. The only real remaining choice is to control the purge %E with the A feed and the reactor pressure with the separator coolant temperature. The poor RGA values for these pairings suggest that control may be poor for this strategy, but there is little else to do with these variables if the A feed is not always available.

Table 7.13 shows the RGA and recommended variable pairing for a strategy designated 2b. Strategy 2b uses the variable pairings that we would choose if the A feed were always available for control. This strategy has much more favorable RGA values than the original Strategy 2, but we expect that it will fail quickly whevenver the A feed becomes unavailable.

Given the four different options for controlled variable sets, the RGA might help select from among them. Skogestad and Morari (89) show that large elements in an RGA matrix are a predictor of stiff control, so we would prefer strategies whose maximum RGA value is close to 1. The maximum RGA values for each of our 6 strategies are listed in Table 7.14. It is immediately obvious that Strategy 1 fares much better in this analysis than the others. Table 7.14 also shows the Niederlinski Index (NI) for each strategy. A negative NI indicates that the chosen pairings are unstable regardless of tuning. None of the proposed control strategies has a negative Niederlinski Index.

At this point, based on our preliminary evaluations, we might expect Strategy 1 to perform best for the following reasons:

- This strategy was the result of our attempt to find the set of variables that best fit all the criteria of Section 7.2.3 above.
- This strategy has a lower maximum RGA values than the others, perhaps indicating more robust control.

Because it is reasonably straightforward to do dynamic testing, however, we will evaluate Strategies 1, 2b, 3, and 4.

Table 7.13: Strategy 2b--An Alternative Pairing for Strategy 2

	A feed ratio	E feed ratio	A&C feed ratio	Reactor coolant temperature	Separator coolant temperature	Purge flow setpoint	Recycle valve position
Reactor level	-3.18	0.02	1.98	-0.04	0.24	-0.03	2.01
Reactor pressure	5.65	0.03	-2.05	-0.19	0.12	0.01	-2.58
Product %G	0.00	0.78	0.25	0.00	0.00	-0.01	-0.01
Stripper temp	-1.18	-0.08	0.75	0.72	-0.40	0.00	1.20
Purge %B	0.00	0.00	-0.04	0.00	0.00	1.04	0.00
Recycle flow	-0.54	-0.04	0.24	0.47	1.03	0.00	-0.15
Purge %E	0.25	0.29	-0.13	0.04	0.02	0.00	0.53

Table 7.14: Steady-State Measures for Proposed Control Strategies

Strategy	Max RGA	NI
1	1.98	0.85
2	5.65	2.13
2b	5.65	0.86
3	4.96	2.18
4	5.38	1.19

7.3.2 Tuning the Control Strategies

Tuning the control loops for each strategy was a formidable task because of the complex dynamics of the plant. There was no significant dead time associated with any of the measurements; the dead time built into the model for composition measurements was small compared to the response time following any change in the inputs. Many loops exhibited either inverse response or underdamped behavior, making standard PI tuning procedures ineffective. Some loops were tuned with integral-only controllers to smooth the response.

Compared to the other control strategies, the dynamics encountered in tuning Stragegy 2 were much slower and more complex. The tuning was so unwieldy that this strategy was abondoned during tuning. Strategy 2b, using the same measurements but different pairings, was much easier to tune.

The loops were tuned sequentially, usually starting with the reactor level loop and the reactor pressure loop. The tuning constants were chosen to give a small overshoot in response to a setpoint change, but several of the tuning constants had to be modified during disturbance testing (see below) to reject the set of disturbances.

7.4 Simulating the Control Strategies

All four of the final candidate control strategies maintain the plant base case conditions indefinitely in the absence of disturbances. Downs and Vogel (15) provide a set of 20 test disturbances with the code for the plant. Each of the control strategies was subjected to each of these 20 disturbances. As expected, the reactor pressure is the bellwether measurement--its behavior is a good predictor of the behavior of the other plant measurements as well. The reactor pressure was the only measurement that was ever in danger of violating its constraint.

In the initial simulations, several disturbances caused the plant to violate its pressure constraint for each control strategy. Retuning some of the control loops (not always the pressure loop!) eliminated some of these problems; disturbances 1 and 18 in particular cound not be rejected with initial loop tuning. Other disturbances could not be rejected at the base case operating conditions given by Downs and Vogel, but they could be rejected when the operating conditions were modified.

Some disturbances were, of course, rejected more easily than others. Table 7.15 shows how each control strategy responded to each disturbance. (The disturbances are listed according to the numbers assigned by Downs and Vogel; Table 7.16 explains each disturbance.) Table 7.15 separates the disturbances into four classes:

- rejected with virtually no effect on process operation;
- rejected easily, with only a minor effect on process operation (reactor pressure remaining within 100kPa of setpoint)
- rejected with difficulty, with major upsets in process operation (reactor pressure deviating more than 100 kPa from setpoint)
- not rejected, with reactor pressure finally exceeding 3000 kPa.

Table 7.15: Disturbance Rejection Properties

	Rejected with no difficulty	Rejected with minor upset	Rejected with major upset	Not rejected (overpressure)
Strategy 1	2,3,4,5,7,9, 11,12,14,15 16,12/15[5]	10,19	1,6[1,2,4],8,13, 17[3],18,20	
Strategy 2b	3,4,5,7,9,10, 11,12,14,15, 16,12/15[5]	19	2,8,13,17[3], 18,20	1,6
Strategy 3	2,3,4,5,7,9, 11,12,14,15, 16	10,19,12/15[5]	1,8,13,17[3], 18,20	6[6]
Strategy 4	3,4,5,7,9,11, 12,14,15,16, 12/15[5]	10,19	1[3],2,8,13, 18,20	6[6],17

[1] With reactor level at 60%.

[2] Requires full purge.

[3] Requires full agitation. In practice, the agitator would probably run at full speed all the time, but these disturbances were originally tested under the base case conditions in the problem statement. At the original base case, the agitator operates at half speed.

[4] Under the given conditions, the reactor pressure will not return to setpoint at steady state due to valve saturation, but the reactor will not overpressurize. Lowering the throughput rate will lower the steady-state pressure, but lower throughput is not necessary to maintain successful plant operation.

[5] Downs and Vogel recommend using disturbances 12 and 15 together.

[6] Reactor overpressurized after 13 hours of operation. Disturbance could be rejected if it were of shorter duration.

Table 7.16: Disturbances for the Eastman Challenge Plant

Disturbance number	Description
1	Stream 4 step change in A/C ratio; B composition constant
2	Stream 4 step change in B composition, A/C ratio constant
3	Stream 2 step change in feed temperature
4	Stream 12 step change in cooling water inlet temperature
5	Stream 13 step change in cooling water inlet temperature
6	Stream 1 loss of feed
7	Stream 4 step decrease in header pressure
8	Stream 4 random variations in A, B, C compositions
9	Stream 2 random variations in feed temperature
10	Stream 4 random variations in feed temperature
11	Stream 12 random variation in cooling water inlet temperature
12	Stream 13 random variation in cooling water inlet temperature
13	Drift in reaction kinetics
14	Sticking valve on reactor cooling water
15	Sticking valve on condensor cooling water
16	Unknown
17	Unknown
18	Unknown
19	Unknown
20	Unknown

It is possible, but not certain, that some of the disturbances in this fourth class could be rejected with different loop tuning.

Table 7.15 also notes disturbances that could be rejected by changing process operating conditions. Ricker (81) notes that the base-case operating conditions are not optimal and that a lower reactor level gives better steady-state performance. Such a change in operating conditions improves dynamic performance for some disturbances, too.

7.5 Evaluating Control Strategy Performance

7.5.1 Comparing the Four Candidates

Downs and Vogel do not provide dynamic performance criteria for the Eastman Challenge Plant. Evaluation of various strategies by error integral measurements, such as those measured for Module 2, are probably not completely reliable, either--the plant performance is so sensitive to tuning that differences in such criteria may reflect tuning differences rather than true differences between strategies. This difficulty will be especially apparent when comparing strategies proposed by different authors.

Downs and Vogel recommend that disturbances 1, 4, 8, and 12/15 (12 and 15 in combination) be used as a common basis for comparison. The graphs in Appendix D show the dynamic response of the reactor pressure to these disturbances. Based on these graphs, experience running the plant, and the data in Table 7.15, we draw the following conclusions:

- Strategies 1 and 2b were easiest to tune.
- All four of the proposed control strategies were able to reject nearly all of the disturbances, usually with little difficulty.
- Only Strategy 1 was able to reject all disturbances without causing the plant to shut down.

- Strategy 1 rejected disturbance 1 much more easily than did the other control strategies. Strategies 3 and 4 were also able to reject this disturbance, but much more slowly.
- All four strategies rejected disturbances 4 and 12/15 equally well. These disturbances provide no real basis for distinguishing between the strategies.
- Strategies 1, 3, and 4 were equally good at rejecting disturbance 8. Strategy 2b was somewhat worse; the difference could not be completely attributed to tuning.
- Overall, Strategy 1 appears to be the best of the four candidates.

Disturbance 6, the loss of the A feed was especially difficult to reject. Strategy 1 was able to handle this disturbance, but only by changing the reactor operating conditions, as Table 7.15 indicates. Lowering the reactor level to 60% is a simple change, and Ricker (81) suggests that such a move has economic benefits. The purge requirement, however, is much more significant. When the A feed is lost, there is not enough of component A to react away all the C, which can be removed only by purging it. In order to purge the C, however, we must sacrifice control of the purge %B. Even with these changes to the plant operating conditions, disturbance 6 is still difficult to reject, but it can be done without reducing the throughput rate.

7.5.2 Comparing to Other M-SISO Solutions

At least five other m-SISO solutions to the Eastman Challenge Plant have been published. Solutions by McAvoy et al. (56, 57), Price et al. (80), Vinson et al. (97), Banerjee and Arkun (3), and Ricker (82) were covered briefly in the literature review in Chapter 2. It is somewhat difficult to compare quantitative results among authors because of differences in tuning procedures and differences in the choice of "important" variables.

However, it is possible to make some qualitative comparisons based on disturbance rejection properties.

The two related solutions by McAvoy et al. (56, 57) successfully reject disturbances 1, 4, 8, and 12/15. (Limited data are given for disturbances 2 and 6, and no data are given for other disturbances.) However, these solutions use the A feed to control reactor pressure. As a result, the control strategy must employ a selector to control pressure with the purge stream when the A feed is lost (as in disturbance 6). In practice, this is not so very different from automatically using full purge when the A feed is lost, as we have done with Strategy 1. The McAvoy strategies also must lower the production rate for disturbance 6. Because these strategies use the E feed for reactor level control, they cannot achieve maximum production without losing level control.

Price (80) presents four different possible control strategies. One of these strategies is able to reject all 20 disturbances, though disturbances 6 and 14-20 were simulated with a 15% reduction in plant throughput. This strategy, however, uses the condenser exit temperature as the throughput manipulator; in an industrial setting, this strategy could be effectively implemented only if the coolant supply temperature were very stable. Price does present another strategy using the D feed as the TPM. This strategy failed only for disturbance 6, even with a 15% reduction in throughput. No data are given to indicate whether these strategies would be stable for disturbances 6 and 14-20 if the throughput were not reduced.

Vinson et al. (97) tested only disturbances 1, 6, and 8. Their strategy was unable to reject disturbance 6, and they report difficulties using steam to control the %E in the product composition. (In the analysis earlier in this chapter, the stripper steam flow and product %E were both discovered to be nonessential, so they might be expected to behave poorly as a control loop pairing.)

Banerjee and Arkun (3) present a control system capable of rejecting all 20 disturbances, though their control strategy also requires a decrease of about 17% in throughput rate to reject disturbance 6. They chose not to control the %B in the purge.

Ricker (82) provides the most in-depth view of his heuristic approach to control system design. Ricker's strategy is PID-based, although he includes several override algorithms to handle constraints and push the plant toward optimum operation. His enhanced PID strategy rejected all disturbances effectively. Ricker argues that controlling the %B in the purge is not economical, and he provides no means for doing so.

Strategy 1, the best-performing strategy from this research, compares favorably with the strategies enumerated above:

- Unlike most of the listed solutions, Strategy 1 does not require a reduction in throughput to reject disturbance 6 or any other disturbance. Strategy 1 does require modification of the base operating conditions to reject disturbance 6, but not a reduction in throughput.
- Strategy 1 permits operating at maximum throughput.
- The control strategies presented in this research provide control of the %B in the purge stream, in agreement with the original problem statement. It might be possible to achieve better control without this requirement, as Ricker has chosen to do, since the purge flow would be free for other uses. Strategy 1, however, is able to deal with all disturbances while controlling the purge composition.
- The control strategies presented in this research do not include override controls to deal with possible constraint violations or valve saturation. Control might be improved for some cases with such controls, but the focus of this research is on design of the PID system.

- Some other control strategies seem to maintain lower variability in reactor pressure than does Strategy 1. It is not clear whether this difference is due to tuning or control strategy design.

7.6 Other Observations on the ECP Solution

The majority of control system design research focuses on steady-state aspects of control, but the dynamics of this plant are critical. Often the critical factor in disturbance rejection is not whether the control strategy can reach the correct steady-state endpoint but whether it can do so fast enough to keep the reactor from overpressurizing. This observation suggests that steady-state tools must be used cautiously in designing loops.

Section 7.3 suggested that large values in the RGA might indicate slow or stiff control. The results of this study bear out that conclusion--Strategy 1 had by far the smallest maximum RGA value (1.98 compared to at least 4.96 for the other strategies), and it performed best.

Downs and Vogel (15) suggest multiple modes of operation for the Eastman Challenge Plant, including radical changes in the product mix (with G/H ratios ranging from 90:10 to 50:50 to 10:90). These modes require similarly diverse D/E feed ratios. The throughput manipulator and material balance controls for the control strategies in this research were designed to operate at the base case operating conditions, with a D/E feed ratio of approximately 1:1. If it is necessary to move among the three modes, the control would probably be more effective if the throughput manipulator were the total D+E feed. It is fairly straightforward to design a block whose inputs are the total D+E feed and D/E ratio, and whose outputs are the D feed setpoint and the E feed setpoint. Such a block will probably improve the control over the wide range of feed ratios necessary. Some feedforward compensation should also be included to adjust the feed ratios automatically based on changes in the product composition setpoint. Ricker (82) reports success with such a feedforward loop.

One interesting relationship uncovered in the variable set analysis was the equivalence between reaction temperature and product composition. The variables react at different speeds, however--the product composition is the more important variable, but the reactor temperature responds faster to the manipulated variable (E feed ratio) used to control that composition. These facts suggest the opportunity for a cascade configuration using the product composition as a master loop and the reactor temperature as the slave loop.

The results from the analysis of Module 2 suggested that it might be useful to manipulate reaction volume along with the TPM to adjust a plant's throughput rate. One must be careful in trying to applying this principle to the ECP, however. Consider the graphs shown in Figures 7.3 and 7.4. These graphs show the long-term response of one ECP control strategy to disturbance #1, introduced at a time of 60 hours. Under some sets of operating conditions, the E inventory in the plant may accumulate over a long period of time, as shown in Figure 7.4. For this same set of operating conditions, when the reactor level setpoint is lowered to 65% from its base-case value of 75%, the E inventory is self-regulating and the pressure is stable indefinitely. These graphs indicate that, given a set of operating conditions, there is a limited range of acceptable reactor liquid levels. Without more detailed kinetic data it is difficult to determine what the reactor level ought to be, but these results suggest that perhaps it is more important to adjust the reactor level based on the E inventory than on the desired throughput rate.

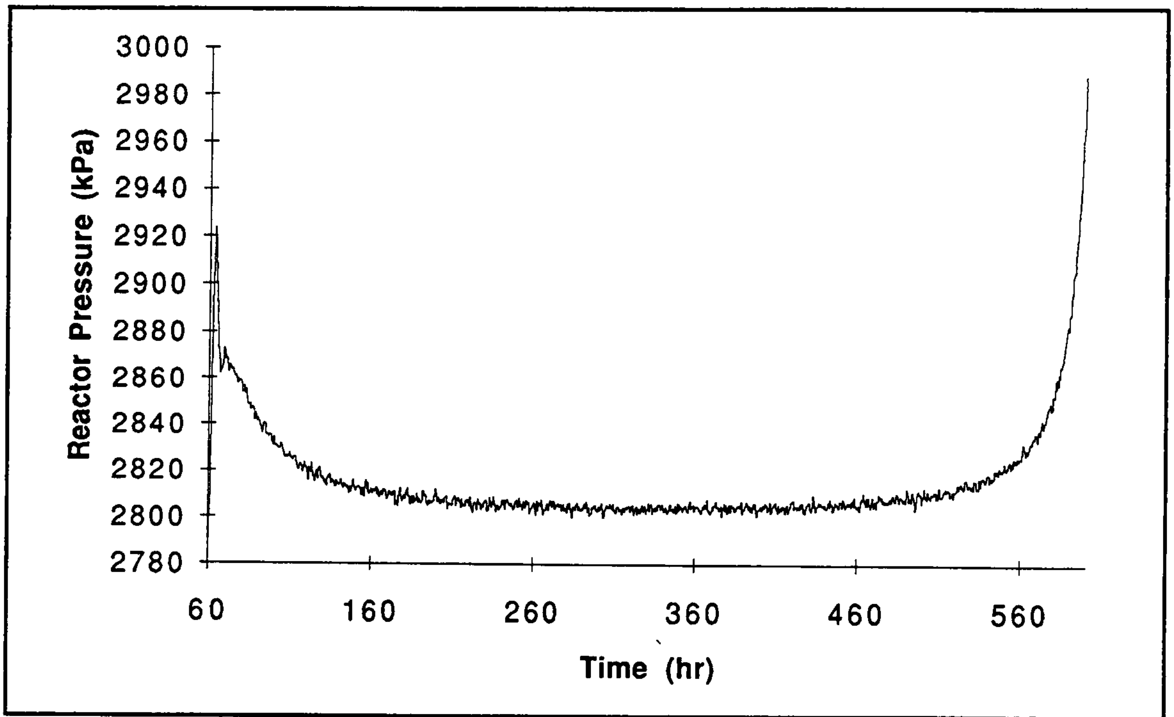

Figure 7.3: Long-Term Pressure Instability

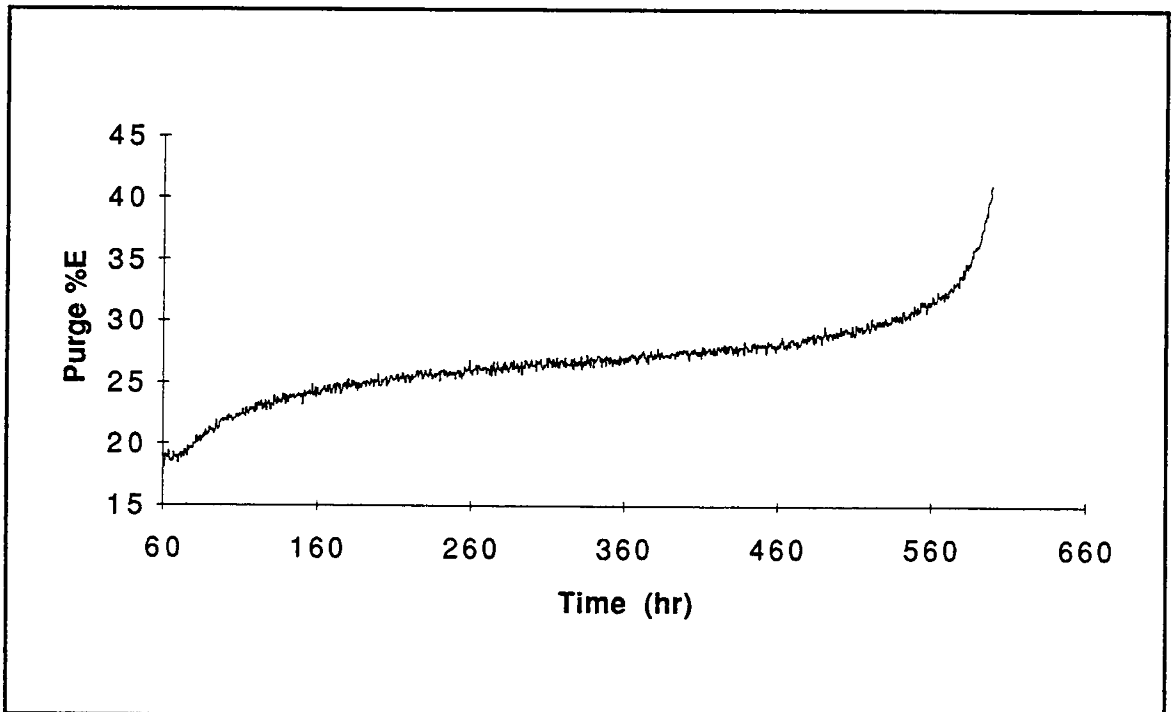

Figure 7.4: Long-Term E Inventory Buildup

7.7 Conclusions from the ECP

- The ECP has fewer real degrees of freedom than manipulated variables or measurements. An important first stage in control strategy development is deciding how many independent control loops can be applied.
- The variable selection procedure described in Chapter 6 provided a reasonably straightforward way to eliminate many of the variables as choices for control. The results of this analysis generally agree with the decisions of other authors about which variables are not useful for control. The ECP also illustrated how the variable selection procedure could suggest possible cascade configurations.
- Tuning the control strategies for the ECP was difficult because of the complex plant dynamics. It is sometimes difficult to distinguish problems with poor tuning from problems with poor control strategy design. Several of the plant disturbances have significant effects on the response characteristics of some of the control loops, requiring tighter or looser tuning to remain stable.
- Most disturbances for the ECP were rejected easily. The most dificult disturbances involve material balance upsets (feed composition changes or loss of A feed). Disturbance rejection seems to be controlled by plant dynamics; the best strategies are those that can deal quickly with upsets in the reactor pressure.
- Of the four control strategies identified in this research, Strategy 1 performed best. It was comparable to the best previously-published solutions, and better than several of them at rejecting all disturbances without reducing the throughput rate. Control of the purge %B had to be abandoned for Strategy 1 to continue functioning after the loss of the A feed.
- Control for these strategies might be improved by relaxing some of the control requirements (for example, the requirement to control the %B in the purge).

Chapter 8: Conclusions and Recommendations

8.1 Control Strategy Design Procedure

This work was intended to help define a systematic procedure for the design of control strategies for large plants. Heuristic and mathematical tools have been proposed for steps of the design process that have not been widely addressed in the literature. The procedure, as illustrated for the Eastman Challenge Plant, consisted of the following steps, which incorporate some results from Chapters 4-7:

- Become familiar with the objectives and operation of the plant. Successful control strategy design begins with thorough plant knowledge. "Thorough" knowledge includes information on operating conditions, common disturbances, response characteristics, and operating constraints.
- Choose a throughput manipulator (TPM), the variable to use for production rate control. With the exception of Price et al. (78, 79, 80), previous literature is scarce on this subject. The following conclusions from the Module 2 analysis should be considered:
  - The location of the TPM should permit a bulk inventory strategy that will direct material balance variation away from critical constraints and away from significant quality measurements.
  - A good throughput manipulator should not be subject to any disturbance that would significantly change the plant throughput at steady state.
  - It may be useful to manipulate reaction volume, not just flow rates, to adjust plant throughput.

- Identify non-self-regulating variables. Bulk inventories (such as liquid levels) are usually integrating. If the plant contains a material recycle stream, there is the potential for component inventories to integrate.
- Stabilize the plant by placing controls on these non-self-regulating variables. The following principles from the Module 2 analysis will be useful:
  - The bulk inventory strategy should be constructed to direct material balance variation away from critical constraints and away from significant quality measurements.
  - Component inventories are best controlled using composition measurements. Less direct measures should only be used with caution. Components present in excess might be preferred as inventory measures.
- Derive the physically scaled steady-state gain matrix for all possible inputs and measurements for the plant. If it is possible, derive gains for disturbances also.
- Use singular value analysis to decide how many degrees of freedom are available for control.
- If the plant does not have the same number of degrees of freedom as it has inputs and measurements, use the variable set reduction method described in Chapter 6 to eliminate nonessential variables and form groups of equivalent variables.
- Use engineering judgment to choose controlled and manipulated variables from the groups of equivalent variables. Preferred variables will have a quick dynamic response, will have a direct relationship to the operating objectives of the plant, or will be closely related to the plant constraints. If there are more available variable groups than degrees of freedom, choose the variable set based on stoichiometric considerations, maximum least singular values, speed of dynamic response, and propagation of variation. It is possible that more than one variable set will be a candidate after this step.

- Use steady-state tools, guided by engineering judgment, to pair controlled variables with manipulated variables in each candidate strategy.
- Evaluate each candidate strategy.

The steps above presume the development of a m-SISO control strategy. Multivariable strategies may be developed by a similar procedure, except that it will not be necessary to choose a single variable from groups of equivalent variables. Instead of pairing variables, a multivariable control law must be created to relate the controlled and manipulated variable sets.

8.2 Results from the Eastman Challenge Plant

- Principles from the Module 2 analysis were useful for selecting a throughput manipulator and designing the subsequent controls. Identifying the key process constraint was critical to the rest of the control strategy design.
- The previous literature does not provide a firm basis for choosing a controlled variable set. However, the variable set reduction method described in this study was useful toward that end. The method was able to predict some of the conclusions of other authors about which variables were not useful for control.
- After choosing a throughput manipulator and stabilizing the plant, the Eastman Challenge Plant was a 29x9 control problem, with over 5.5×10^{12} possible m-SISO control strategies. The step-by-step procedure in Section 8.1 above produced four candidate control strategies for evaluation.
- Careful stoichiometric analysis was able to predict some disturbances that were likely to cause severe operating problems.
- All of the four candidates perform comparably to other literature solutions offered for the ECP. With some minor changes in the operating point, and assuming that the purge %B is not a critical variable to control, one of these

strategies is able to reject all known plant disturbances without reducing the base-case steady-state throughput rate.

- Steady-state tools such as the RGA and singular value analysis were useful for designing the control strategies, but they were useful only when used by one with a good working knowledge of the plant characteristics. It is not possible at this point to create a completely objective set of mathematical rules to describe the design of effective control strategies. However, the best of the four candidate strategies did have the largest minimum singular value and the smallest maximum RGA value, suggesting that steady-state tools can help predict control performance.

8.3 Recommended Future Work

Anyone who spends much time tuning control loops and testing disturbance rejection properties for the Eastman Challenge Plant will quickly realize that, although the steady-state relationships can be unraveled with effort, the plant dynamics are more troublesome. The control strategy design procedure is a mix of mathematical and heuristic techniques, and dynamic considerations are currently handled only heuristically. There is an opportunity for theoretical development of techniques to incorporate dynamics into the design procedure. Such an approach, however, must address many different kinds of common dynamic responses (underdamped, inverse response, etc.) that cannot be modeled as first-order-plus-dead-time blocks, and it is often difficult to approximate such responses with simple transfer functions.

Nonlinearities in the Eastman Challenge Plant pose problems for controller design and tuning. Most of the current theoretical methods (including any techniques that rely on a steady-state gain matrix) assume linear systems. Linear tools have proven useful in designing control systems even for nonlinear plants, but they can break down if the disturbance responses require excursions far from the base-case operating point.

There was no opportunity to use the ECP to test steady-state tools that use disturbance gains, because the model permits only on/off control of the disturbances and does not provide information about the sizes of the disturbances. It would be useful to be able to use the model to calculate some disturbance gains as well as input gains.

The procedure described in Chapter 6 for identifying nonessential and equivalent variables was helpful for the ECP, but there is an opportunity for deriving more a detailed theoretical significance for the precise values of the indices (the diagonals of the matrices $\mathbf{KK}^{+}$ and $\mathbf{K}^{+}\mathbf{K}$). Currently one must be fairly conservative in choosing which variables to discard if there is no clear transition between the very small and large indices. It should also be possible to extend this same sort of pseudoinverse-based analysis to the disturbance gain matrix.

It is a struggle to find test problems that are complex enough to represent real-world challenges, yet simple enough to be analyzed in a straightforward fashion. Unfortunately, the complexity of the ECP limits its usefulness as a development tool. It is useful as a "testing ground" for control system design techniques, but the process interactions are so complex that it is difficult to analyze what physical mechanism is responsible when a control system fails in some aspect. There is a need for test plants of intermediate complexity for developing control design tools. Module 2 is a step in that direction--it is more complex than, for example, Module 1 or the Wood and Berry distillation column module, but is much easier to analyze than the ECP. It would be useful, however, to be able to gradually increase the complexity of Module 2 or a similar plant instead of moving directly from Module 2 to the ECP. The fact that the ECP is based on a real plant, however, highlights the need for design tools that can address some of the complexities found in that plant.

REFERENCES

1. Arkun, Y., and J. Downs, *A General Method to Calculate Input-Output Gains and the Relative Gain Array for Integrating Processes.* Comp. Chem. Eng., 14, #10, 1101-1110, 1990.

2. Arkun, Y., and V. Manousiouthakis, *The Dynamic Block Relative Gain and Its Role in the Design of Robust Decentralized Control.* Proceedings of the 1986 American Control Conference, 11-16, 1986.

3. Banerjee, A., and Y. Arkun, *Control Configuration Design Applied to the Tennessee Eastman Plant-Wide Control Problem.* Comp. Chem. Eng., 19(4), 453-480, 1995.

4. Barton, G. W., R. D. Johnston, and M. L. Brisk, *Comments on Determining Structural Controllability.* Int. J. Control, 38(5), 1081-1083, 1983.

5. Bekkers, P. and Rijnsdorp, J. E., *EPIC, A Software Tool for the Control Engineer in Early Process Design.* ESCAPE Symposium, Lyngby, Denmark, 1992.

6. Bristol, E.H., *On a New Measure of Interaction for Multivariable Process Control.* IEEE Transactions on Automatic Control, 133-134, 1966.

7. Bruns, D. D., and C. R. Smith, *Singular Value Analysis: A Geometrical Structure for Multivariable Processes.* Presented at the AIChE Winter Meeting, Orlando, Florida, 1982.

8. Buckley, P. S., Techniques of Process Control, John Wiley & Sons, Inc., New York, 1964.

9. Chang, J., and C. Yu, *Relative Disturbance Gain Array.* AIChE Journal, 38, #4, 521-534.

10. Chintapalli, P., and J. M. Douglas, *The Use of Economic Performance Measures to Synthesize Optimal Control Systems.* Ind. Eng. Chem. Fundam., 14, 1-10, 1975.

11. Chiu, M., and Y. Arkun, *Decentralized Control Structure Selection Based on Integrity Considerations.* Ind. Eng. Chem. Res., 29, 369-373, 1990.

12. Downs, J. J., *Distillation Control in a Plantwide Control Environment*, chapter 20 of Practical Distillation Control, edited by W. Luyben, Van Nostrand Reinhold Publishers, New York, 1990.

13. Downs, J. J., and J. E. Doss, *Present Status and Future Needs--A View from North American Industry.* In Chemical Process Control--CPCIV, Yaman Arkun and W. Harmon Ray, editors, CACHE, 53-77, 1991.

14. Downs, J. J., and B. Ogunnaike, *Design for Control and Operability: An Industrial Perspective*. Foundations of Computer Aided Process Design Conference, Snowmass, CO, July 11-15, 1994.

15. Downs, J. J., and E. F. Vogel, *A Plant-Wide Industrial Process Control Problem.* Comp. Chem. Eng., 17(3), 245-255, 1993.

16. Fisher, W. R., *Synthesis of Steady-state Control Structures for Complete Chemical Plants*. PhD Dissertation, University of Massachusetts, 1985.

17. Fisher, W.R., M. F. Doherty, and J. M. Douglas, *The Interface Between Design and Control* (Parts 1, 2, and 3). Ind. Eng. Chem. Res., 27, 597-605, 1988.

18. Fisher, W.R., M. F. Doherty, and J. M. Douglas, *Operating Heuristics for the Control of Complete Chemical Plants*. Proceedings of the 1985 American Control Conference, 293-298, 1985.

19. Fisher, W. R., M. F. Doherty, and J. M. Douglas, *Steady-State Control as a Prelude to Dynamic Control.* Chem. Eng. Res. Des., 63, 353-357, 1985.

20. Fonyo, Z., *Design Modifications and Proper Plantwide Control.* Comp. Chem. Eng., 18 (Supplement), S483-S492, 1994.

21. Friedly, J., *Use of the Bristol Array in Designing Noninteracting Control Loops. A Limitation and Extension.* Ind. Eng. Chem. Process Des. Dev., 23, 469-472, 1984.

22. Gagnepain, J., and D. Seborg, *Analysis of Process Interactions with Applications to Multiloop Control System Design.* Ind. Eng. Chem. Process Des. Dev., 21, 5-11, 1982.

23. Garcia, C., and M. Morari, *Internal Model Control: 1. A Unifying Review and Some New Results*. Ind. Eng. Chem. Process Des. Dev., 21, 308-323, 1982.

24. Garcia, C. E., and D. M. Prett, *Design Methodology Based on the Fundamental Control Problem Formulation*. In The Shell Process Control Workshop, Butterworth Publishers, 1987.

25. Govind, R., and G. J. Powers, *Control System Synthesis Strategies*. AIChE Journal, 28, 60-73, 1982.

26. Grosdidier, P., *Analysis of Interaction Direction with the Singular Value Decomposition*. Computers Chem. Engr., Vol 14, No. 6, pp 687-689, 1990.

27. Grosdidier, P., M. Morari, and B. Holt, *Closed-Loop Properties from Steady-State Gain Information*. Ind. Eng. Chem. Fundam., 24, 221-235, 1985.

28. Jafarey, A., and T. J. McAvoy, *Degeneracy of Decoupling in Distillation Systems*. Ind. Eng. Chem. Process Des. Dev., 17, 485-490, 1978.

29. Jensen, N., *Comments on "Some Results on Dynamic Interaction Analysis of Complex Control Systems".* Ind. Eng. Chem. Process Des. Dev., 24, 228-229, 1985.

30. Johnston, R. D., and G. W. Barton, *Control Objective Reduction in Single-input Single-output Control Schemes.* Int. J. Control, 40(2), 265-270, 1984.

31. Johnston, R. D., and G. W. Barton, *Design and Performance Assessment of Control Systems Using Singular Value Analysis.* Ind. Eng. Chem. Res., 26, 830-839, 1987.

32. Johnston, R. D., G. W. Barton, and M. L. Brisk, *Determination of the Generic Rank of Structural Matrices.* Int. J. Control, 40(2), 257-264, 1984.

33. Johnston, R. D., G. W. Barton, and M. L. Brisk, *Single-Input--Single-Output Control System Synthesis. Parts 1 and 2.* Comp. Chem. Eng., 9(6), 547-566, 1985.

34. Joshi, S. K., and J. M. Douglas, *Avoiding Accumulation of Trace Components.* Ind. Eng. Chem. Res., 31, 1502-1509, 1992.

35. Kanadibhotla, R. S., and J. B. Riggs, *Nonlinear Model Based Control of a Recycle Reactor Process.* Comp. Chem. Eng., 19(8), 933-948, 1995.

36. Kapoor, N., T.J. McAvoy, and T.E. Marlin, *Effect of Recycle Structure on Distillation Tower Time Constants.* AIChE Journal, 32, #3, 411-418.

37. Keller, J. P., and D. Bonvin, *Selection of Input and Output Variables as a Model Reduction Problem.* Proceedings of the 1990 American Control Conference, 2115-2119, 1990.

38. Klema, V., and A. Laub, *The Singular Value Decomposition: Its Computation and Some Applications.* IEEE Transactions on Automatic Control, 25(2), 164-176, 1980.

39. Koppel, L. B., *Conditions Imposed by Process Statics on Multivariable Process Dynamics.* AIChE Journal, 31, 70-75, 1985.

40. Kresta, J. V., J. F. MacGregor, and T. E. Marlin, *Multivariate Statistical Monitoring of Process Operating Performance.* Canadian Journal of Chemical Engineering, 69, 35-47, 1991.

41. Krzanowski, W. J., *Selection of Variables to Preserve Multivariate Data Structure, Using Principal Components.* Appl. Statist., 36(1), 22-33, 1987.

42. Lau, H., J. Alvarez, and K. F. Jensen, *Synthesis of Control Structures by Singular Value Analysis: Dynamic Measures of Sensitivity and Interaction.* AIChE Journal, 31, 427-439, 1985.

43. Lau, H. and K. F. Jensen, *Evaluation of Changeover Control Policies by Singular Value Analysis--Effects of Scaling*. AIChE Journal, 31, 135-146, 1985.

44. Luyben, W., *Dynamics and Control of Recycle Systems: 1. Simple Open-Loop and Closed-Loop Systems; 2. Comparison of Alternative Process Designs; 3. Alternative Process Designs in a Ternary System.* Ind. Eng. Chem. Res., 32, 466-486, 1993 (Part 3 Submitted).

45. Luyben, W., Process Modeling, Simulation, and Control for Chemical Engineers, McGraw-Hill Publishing, New York, 1990.

46. Luyben, W., *Sensitivity of Distillation Relative Gain Arrays to Steady-State Gains*. Ind. Eng. Chem. Res., 26, 2076-2078, 1990.

47. Luyben, W., *Snowball Effects in Recycle Systems*. Submitted to Ind. Eng. Chem. Res., 1993.

48. Luyben, W., *Steady-State Energy Conservation Aspects of Distillation Column Control System Design*. Ind. Eng. Chem. Fundam., 14, 321-325, 1975.

49. Luyben, W., *The Concept of "Eigenstructure" in Process Control*. Ind. Eng. Chem. Res., 27, 206-208, 1990.

50. Manousiouthakis, V., R. Savage, and Y. Arkun, *Synthesis of Decentralized Process Control Structures Using the Concept of Block Relative Gain*. AIChE Journal, 32, 991-1003, 1986.

51. Marino-Galarraga, M., T. J. McAvoy, and T. E. Marlin, *Shortcut Operability Analysis. 3. Short-Cut Methodology for the Assessment of Process Control Designs*. Ind. Eng. Chem. Res., 26, 521-531, 1987.

52. McAvoy, T. J., *Connection Between Relative Gain and Control Loop Stability and Design*. AIChE Journal, 27(4), 613-619, 1981.

53. McAvoy, T.J., *Interacting Control Systems: Steady State Treatment of Dual Composition Control in Distillation Columns*. ISA Transactions, 16, 83-90, 1977.

54. McAvoy, T. J., *Response to Comments on "Some Results on Dynamic Interaction Analysis of Complex Control Systems"*. Ind. Eng. Chem. Process Des. Dev., 24, 229-230, 1985.

55. McAvoy, T.J., *Some Results on Dynamic Interaction Analysis of Complex Control Systems*. Ind. Eng. Chem. Process Des. Dev., 22, 42-49, 1983.

56. McAvoy, T. J., and N. Ye, *Base Control for the Tennessee Eastman Problem*. Comp. Chem. Eng., 18(5), 383-413, 1994.

57. McAvoy, T., N. Ye, and C. Gang, *An Improved Base Control for the Tennessee Eastman Problem*. Proceedings of the 1995 American Control Conference, Vol 1, 240-244, 1995.

58. McCabe, G., *Principal Variables*. Technometrics, 26, 137-144, 1984.

59. Mellichamp, D. A., *Industrial Challenge Problems in Process Control*. Comp. Chem. Eng., 17(3), v-ix, 1993.

60. Mijares, G., J. D. Cole, N. W. Naugle, H. A. Preisig, and C. D. Holland, *A New Criterion for the Pairing of Control and Manipulated Variables*. AIChE Journal, 32(9), 1439-1449, 1986.

61. Moore, B. C., *Principal Component Analysis in Linear Systems: Controllability, Observability, and Model Reduction*. IEEE Trans. Automatic Control, AC-26(1), 17-31, 1981.

62. Moore, C. F., *A New Role for Engineering Process Control Focused on Improving Quality*, chapter 19 of Competing Globally Through Customer Value, edited by M. J. Stahl and G. M. Bounds, Quorum Books, New York, 1991.

63. Moore, C. F., *Application of Singular Value Decomposition to the Design, Analysis, and Control of Industrial Processes*. Proceedings of the 1986 American Control Conference, 643-650, 1986.

64. Moore, C. F., *Selection of Controlled and Manipulated Variables*, chapter 8 of Practical Distillation Control, edited by W. Luyben, Van Nostrand Reinhold Publishers, New York, 1990.

65. Moore, C. F., J. Hackney, and D. Canter, *Selecting Sensor Location and Type for Multivariable Processes*. In The Shell Process Control Workshop, Butterworth Publishers, 1987.

66. Morari, M., *Design of Resilient Processing Plants--III. A General Framework for the Assessment of Dynamic Resilience*. Chem. Eng. Sci., 38(11), 1881-1891, 1983.

67. Morari, M., Y. Arkun, and G. Stephanopoulos, *Studies in the Synthesis of Control Structures for Chemical Processes*. AIChE Journal, 26, 220-246, 1980.

68. Morari, M., and E. Zafiriou, Robust Process Control, Chapter 14, Prentice Hall, Englewood Cliffs, N.J., 1989.

69. Narraway, L. T., and J. D. Perkins, *Selection of Process Control Structure Based on Linear Dynamic Economics*. Ind. Eng. Chem. Res., 32, 2681-2692, 1993.

70. Niederlinski, A., *A Heuristic Approach to the Design of Linear Multivariable Interacting Control Systems*. Automatica, 7, 691-701, 1971a.

71. Niederlinski, A., *Two-Variable Distillation Control: Decouple or Not Decouple.* AIChE Journal, 17, 1261-1263, 1971b.

72. Norris, G. A., and R. E. Skelton, *Selection of Dynamic Sensors and Actuators in the Control of Linear Systems.* Journal of Dynamic Systems, Measurement, and Control, 111, 389-397, 1989.

73. Papadourakis, A., M. F. Doherty, and J. M. Douglas, *Relative Gain Array for Units in Plants with Recycle.* Ind. Eng. Chem. Res., 26, 1262-1264, 1987.

74. Papastathopoulou, H., and W. Luyben, *Potential Pitfalls in Ratio Control Schemes.* Ind. Eng. Chem. Res., 29, 2044-2053, 1990.

75. Papastathopoulou, H., and W. Luyben, *A New Method for the Derivation of Steady-State Gains for Multivariable Processes.* Ind. Eng. Chem. Res., 29, 366-369, 1990.

76. Perkins, J. D., G. W. Barton, C. Gannavarapu, and J. B. Lear, *Choosing Control Structures for Continuous Processes Based on Economics.* Proceedings of the 1990 American Control Conference, p. 2142-2143, 1990.

77. Prett, D. M., C. E. Garcia, and B. L. Ramaker, The Second Shell Process Control Workshop: Solutions to the Shell Standard Control Problem, Butterworth Publishers, 1990.

78. Price, R. M., *Design of Plant-Wide Regulatory Control Systems.* PhD dissertation, Lehigh University, 1993.

79. Price, R. M., and C. Georgakis, *Plantwide Regulatory Control Design Procedure Using a Tiered Framework.* Ind. Eng. Chem. Res., 32, 2693-2705, 1993.

80. Price, R. M., P. R. Lyman, and C. Georgakis, *Throughput Manipulation in Plantwide Control Structures.* Ind. Eng. Chem. Res., 33, 1197-1207, 1994.

81. Ricker, N. L., *Optimal Steady-state Operation of the Tennessee Eastman Challenge Process.* Comp. Chem. Eng., 19(9), 949-959, 1995.

82. Ricker, N. L., *Decentralized Control of the Tennessee Eastman Challenge Process.* Submitted to Journal of Process Control, 1995.

83. Ricker, N. L., and J. H. Lee, *Nonlinear Model Predictive Control of the Tennessee Eastman Challenge Process.* Comp. Chem. Eng., 19(9), 961-981, 1995.

84. Ricker, N. L., and J. H. Lee, *Nonlinear Modeling and State Estimation for the Tennessee Eastman Challenge Process.* Comp. Chem. Eng., 19(9), 983-1005, 1995.

85. Rijnsdorp, J. E., and P. Bekkers, *Early Integration of Process and Control Design.* submitted for IFAC-Workshop on Interactions Between Process and Control Design, London, 1992.

86. Skogestad, S., and M. Hovd, *Use of Frequency-Dependent RGA for Control Structure Selection.* Proceedings of the 1993 American Control Conference, 2133-2139, 1990.

87. Skogestad, S., and M. Morari, *Design of Resilient Processing Plants--IX. Effect of Model Uncertainty on Dynamic Resilience.* Chem. Eng. Sci., 42(7), 1765-1780, 1987.

88. Skogestad, S., and M. Morari, *Effect of Disturbance Directions on Closed-Loop Performance.* Ind. Eng. Chem. Res., 26, 2029-2035, 1987.

89. Skogestad, S., and M. Morari, *Implications of Large RGA Elements on Control Performance.* Ind. Eng. Chem. Res., 26, 2323-2330, 1987.

90. Stanley, G., M. Marino-Galarraga, and T. J. McAvoy, *Shortcut Operability Analysis. 1. The Relative Disturbance Gain.* Ind. Eng. Chem. Process Des. Dev., 24, 1181-1188, 1985.

91. Stephanopoulos, G., *Synthesis of Control Systems for Chemical Plants--A Challenge for Creativity.* Comp. Chem. Eng., 7(4), 331-365, 1983.

92. Strang, G., Linear Algebra and Its Applications, 3rd ed., Harcourt Brace Jovanovich, San Diego, 1988.

93. Thurston, C. W., *Computer-aided Design of Distillation Column Controls.* Hydrocarbon Processing, July, 1981, 125-130, 1981.

94. Tung, L.S., and T.F. Edgar, *Analysis of Control-Output Interactions in Dynamic Systems.* AIChE Journal, 27, #4, 690-693, 1981.

95. Tyréus, B., *Multivariable Control System Design for an Industrial Distillation Column.* Ind. Eng. Chem. Process Des. Dev., 18, 177-182, 1979.

96. Van de Wal, M., and B. Jager, *Control Structure Design: A Survey.* Proceedings of the 1995 American Control Conference, Vol 1, 225-229, 1995.

97. Vinson, D., C. Georgakis, and J. Fossy, *Studies in Plant-Wide Controllability Using the Tennessee Eastman Challenge Problem, The Case for Multivariable Control.* Proceedings of the 1995 American Control Conference, Vol 1, 250-254, 1995.

98. Waller, J. B., M. F. Sågfors, and K. V. Waller, *Ill-Conditionedness and Process Directionality--The Use of Condition Numbers in Process Control.* IFAC Symposium on Advanced Control of Chemical Processes, Kyoto, Japan, p. 463-468, 1994.

99. Williams, T. J., and R. E. Otto, *A Generalized Chemical Processing Model for the Investigation of Computer Control.* AIEE Transactions, 79, Part I, 458, 1960.

100. Williams, V. J., *Management of Variation in Plant-Wide Process Control System Design and Analysis.* PhD Thesis, University of Tennessee, Knoxville, 1993.

101. Wise, B. M., N. L. Ricker, D. F. Veltkamp, and B. R. Kowalski, *A Theoretical Basis for the Use of Principal Component Models for Monitoring Multivariate Processes.* Process Control and Quality, 1, 41-51, 1990.

102. Witcher, M. F., and T. J. McAvoy, *Interacting Control Systems: Steady-State and Dynamic Measurement of Interaction.* ISA Transactions, 16, 35-41, 1977.

103. Wold, S., K. Esbensen, and P. Geladi, *Principal Component Analysis.* Chemometrics and Intelligent Laboratory Systems, 2, 37-52, 1987.

104. Wood, R. K., and M. W. Berry, *Terminal Composition Control of a Binary Distillation Column.* Chem. Eng. Sci., 28, 1707, 1973.

105. Woolverton, P., *How to Use Relative Gain Analysis in Systems with Integrating Variables.* InTech, 27, 63-65, 1980.

106. Yan, M. and N.L. Ricker, *Multi-Objective Control of the Tennessee Eastman Challenge Process.* Proceedings of the 1995 American Control Conference, Vol 1, 245-249, 1995.

107. Yu, C., and W. Luyben, *Design of Multiloop SISO Controllers in Multivariable Processes.* Ind. Eng. Chem. Process Des. Dev., 25, 498-503, 1986.

108. Yu, C., and W. Luyben, *Robustness with Respect to Integral Controllability.* Ind. Eng. Chem. Res., 26, 1043-1045, 1987.

APPENDICES

APPENDIX A

An Introduction to Some Basic Plant-Wide Control Concepts

An Introduction to Some Basic Plant-Wide Control Concepts

The material on the following pages is a brief introduction to some of the basic concepts behind plant-wide process control. These ideas are not new, original, or startling; many of them reflect material presented at the short course titled "Plant-Wide Design, Operation, and Control," conducted at Lehigh University from May 3-7, 1993. This presentation is appropriate for an audience with some limited background in process control, and it has been used as the basis for teaching an introduction to plant-wide control in a section of the undergraduate introductory control course here at the University of Tennessee.

It is common for control engineers, whether academic or industrial, to consider control loops one at a time. The academic approach to process control, especially at the undergraduate level, often encourages this single-loop approach to control problems. When an industrial process exhibits control difficulties, the first response is usually to try to discover which loop is causing the problem. Sometimes, however, the problem is difficult to diagnose because it is the fault not of one loop but of the overall design of the multi-loop strategy. The challenge for the control engineer is to recognize which control strategy designs are bad before they are implemented on an actual process. This study of plant-wide control issues presents a few simple ways to distinguish between good and bad designs.

The principles outlined in this discussion are based partly on presentations from a May, 1993, chemical engineering short course entitled "Plant-Wide Design, Operation, and Control," conducted at Lehigh University by representatives of several universities and corporations interested in plant-wide control. One of the course's primary points of emphasis was the importance of evaluating steady-state material balances. Plant stability requires that inventory controls in a plant (control loops that regulate amounts of material) must allow material to enter, leave, or accumulate as necessary to reach a new operable steady state after a disturbance or a setpoint change; therefore, it is paramount to understand the material balances in any plant before designing a control strategy for it. The design principles and the accompanying examples in this paper will point out the need to evaluate material balances for bulk materials and for individual components within the plant.

Generally speaking, "good" control strategies can stabilize a plant at all points within some desired operating range. "Bad" strategies cannot stabilize the plant everywhere within that range, often because of a failure to satisfy some material balance for the plant. "Marginal" strategies can stabilize the plant, but cannot recover quickly or smoothly from certain disturbances. The objective of this study is to provide some guidance in deciding which strategies are likely to be good ones. Once control engineers have identified a set of potential control strategies for a particular plant, they must use some more rigorous method, such as a series of simulations, to decide which strategy is best, based on some set of operating objectives. Identification of good control strategy designs involves addressing a number of fundamental issues:

- What are the objectives of this plant?
- What are the quality and performance measures for the plant?
- Where are the production bottlenecks and constraints?
- What variable should control the production rate?
- How should bulk inventories be controlled?
- Where are the recycle systems?
- How are component inventories controlled in these recycle systems?
- Would process design changes improve the operation and control?
- What are the primary sources of variation?
- How does variation propagate through the plant?
- Can this variation be redirected to less harmful locations?

It is obvious from this list that the set of good control strategies will be different for every plant. This paper will present six general control design principles that apply to nearly all plant configurations, although evaluation of individual strategies requires detailed knowledge of the specific plant operations. The first three principles presented here will be illustrated very simply with tanks, flow controllers (FC), and level controllers (LC). The last three are presented during discussion of the inventory control for a small chemical plant.

The tank in Figure A.1(a) illustrates a simple problem involving material balances. The tank has one inlet stream with a flow controller and one outlet stream with a flow controller. It is easy to see the fatal flaw in the design of this control strategy: If the setpoints on the two flow control loops are different, the tank will run dry or overflow. Even if the setpoints are identical, it is not possible for the flow control strategy to measure and dispense exactly equal flows. Thus, if the setpoints on the loops are both exactly 1 gallon per minute (gpm), the inlet flow might average 1.001 gpm while the exit flow averages 0.999 gpm. Because the inlet flow exceeds the outlet flow, the tank will eventually overflow even though the flow measurements are very precise.

If a real process actually employed the control strategy shown in Figure A.1(a), an alert process engineer would quickly spot and diagnose the problem. However, it is usually easier and much less costly to identify problems in the design phase instead of waiting to analyze failures in the real-time operation of the process. A control engineer who takes the time to evaluate the material balance for this tank will spot the flaw. The general material balance equation is

$$\text{Input - Output + Generation - Consumption = Accumulation.} \tag{1}$$

Because no material is generated or consumed in this tank, and because there is only one inlet stream and one outlet stream, the balance simplifies to

$$\text{(Flow in) - (Flow out) = Accumulation.} \tag{2}$$

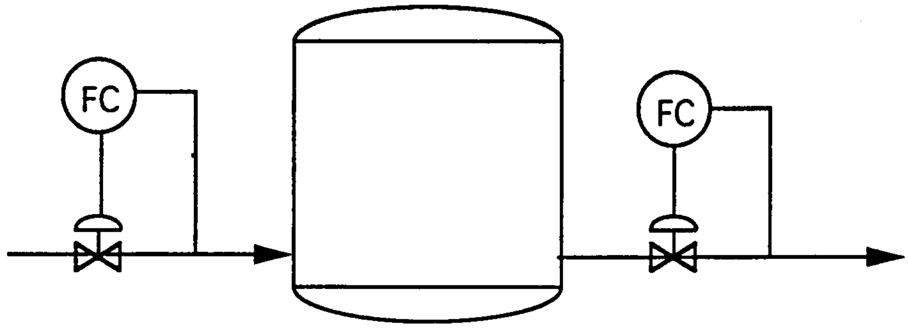

(a) Control system without inventory control

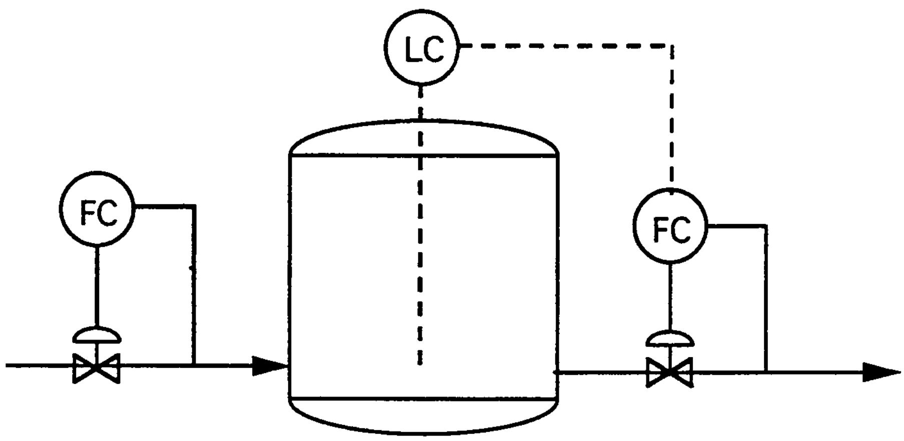

(b) Control system with inventory control

Figure A.1: Control Systems for a Simple Tank

In order for the tank to achieve a steady state, the accumulation of material must be exactly zero. Thus, if both flow rates are fixed, the flow rate into the tank must be exactly the same as the flow rate out; if the rates differ by even the smallest fraction, the tank will overflow or run dry.

An easy solution to this material balance problem, as shown in Figure A.1(b), is to add a level controller that sets the setpoint to one of the flow controllers. The level in the tank is a measure of the inventory (amount of material) within the tank. Because the level controller increases or decreases the outlet flow (in this case) in response to inventory changes, this control strategy satisfies the steady-state material balance. If the level is constant, no material accumulates inside the tank.

This material balance analysis suggests the first of the principles that should guide the design of a control strategy for a plant:

> Design Principle #1: *A control strategy cannot independently set all the inputs and outputs to a process. It must measure every non-self-regulating inventory in the process and manipulate inputs or outputs to maintain each of these inventories.*

In our simple example in Figure A.1, one flow rate controls the tank level, which is a measure of the inventory of material in the tank. The same analysis is also necessary for large plants, however. The same overall material balance applies to multi-unit plants with multiple inlets and outlets, and the same kind of problem may arise. The control strategy must have some measurement of the inventory in each system and must have a manipulated variable available to control that inventory at some level. Otherwise it will be impossible for the process to reach steady state.

Notice the emphasis on "non-self-regulating" inventories in Design Principle #1. A self-regulating inventory is one that reaches a new steady state after a change in one of the flow rates. For example, if exit stream from the tank in Figure A.1(a) were set by gravity flow, then it would be self-regulating; a small increase in the inlet flow would lead to a higher tank level, and the outlet flow would automatically increase until the two flows were equal and the tank level was constant. A controller might be used for a self-regulating inventory, but it is not absolutely essential.

For the tank in Figure A.1(b), it makes no difference which of the two process streams controls the tank level. In general, though, some manipulated variables are better choices than others for bulk inventory control. The best choice depends on the exact design of the plant. For example, consider the tank shown in Figure A.2(a), a slightly more complicated configuration than the one in Figure A.1. This tank has two inlet streams, two outlet streams, and one inventory measurement. From only the information given in this figure, it is impossible to say which stream is best to use for level (inventory) control. The control engineer must have additional information from the process designer. Suppose the following information is available, as in Figure A.2(b):

Stream 1: Controllable; design flow rate of 1000 gpm.

Stream 2: Wild (not controllable); flow rate fluctuates from 0 to 200 gpm with average of 110 gpm.

Stream 3: Controllable; purge stream with design flow rate of 10 gpm.

Stream 4: Controllable; design flow rate of 1100 gpm.

The major disturbance to the material balance is Stream 2, which may vary as much as 110 gpm from the design condition. In order to satisfy the steady-state material balance, the stream used for level control must be able to change as much as Stream 2 changes. If

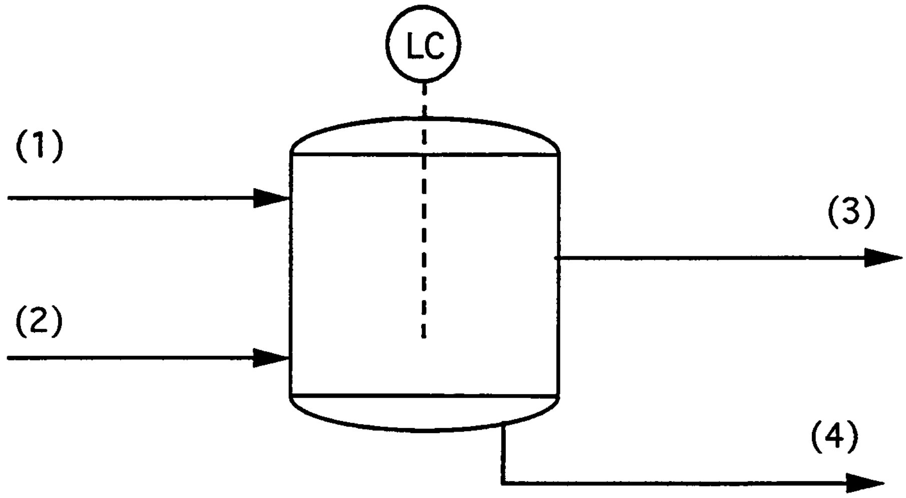

(a) Not enough process information to choose inventory controller

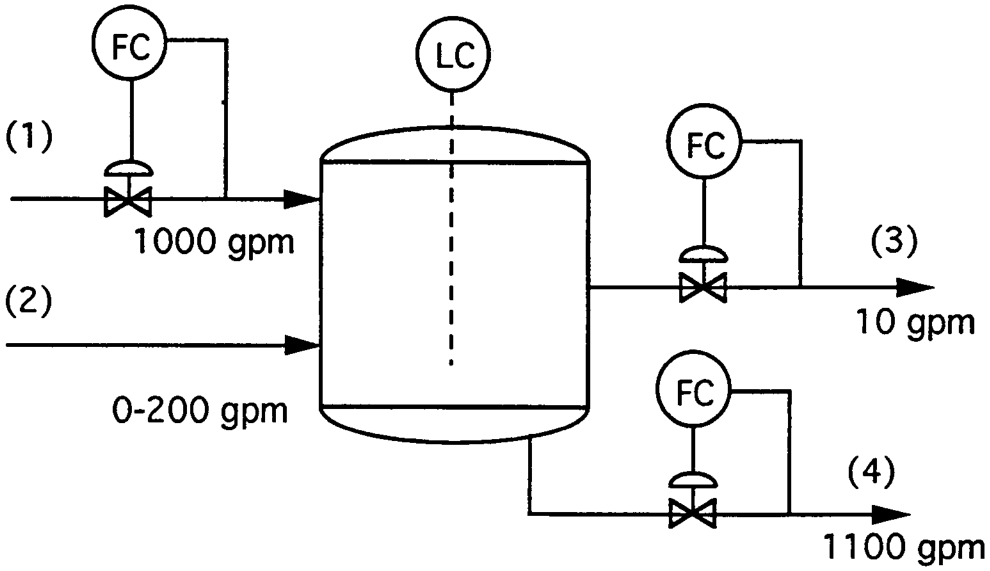

(b) Essential information for inventory control design

Figure A.2: Choosing the Best Stream for Inventory Control

Stream 2 suddenly drops from 110 gpm to 20 gpm, some other stream must change its flow by a corresponding 90 gpm to counter the effect of the disturbance. The piping for Stream 3, originally designed to transport only 10 gpm, is probably too small to carry an additional 90 gpm if the flow rate of Stream 2 should increase by that amount. Likewise, if the flow rate of Stream 2 drops by 50 gpm, the flow rate of Stream 3 cannot drop by that amount to maintain the tank level. Therefore, Stream 3 is probably a poor choice for level control. Streams 1 and 4 are large enough to compensate for the expected flow rate changes, so the level controller should manipulate the flow rate setpoint for one of those two streams. Furthermore, these large streams will give quicker level correction than small streams would. This decision suggests our second design principle:

> Design Principle #2: *Streams used for inventory control must be large enough to compensate for the expected range of material balance disturbances.*

As before, the same analysis applies for large plants. The tank in Figure A.2 might instead be a distillation column and reactor; the level controller could just as easily be some other measure of inventory in the larger system. The same design principle still applies for the same reason: the steady-state material balance cannot be satisfied unless the inventory control stream can vary just as widely as the disturbances.

It is very important to observe that this second principle does not require that the streams used for inventory control have the largest flow rates in the system. The streams must only be sufficiently large to compensate for the expected range of material balance disturbances. In Figure A.3 there are two reactors with a recycle stream. The control strategy uses the largest streams, Streams 2 and 4, to control the tank inventories.

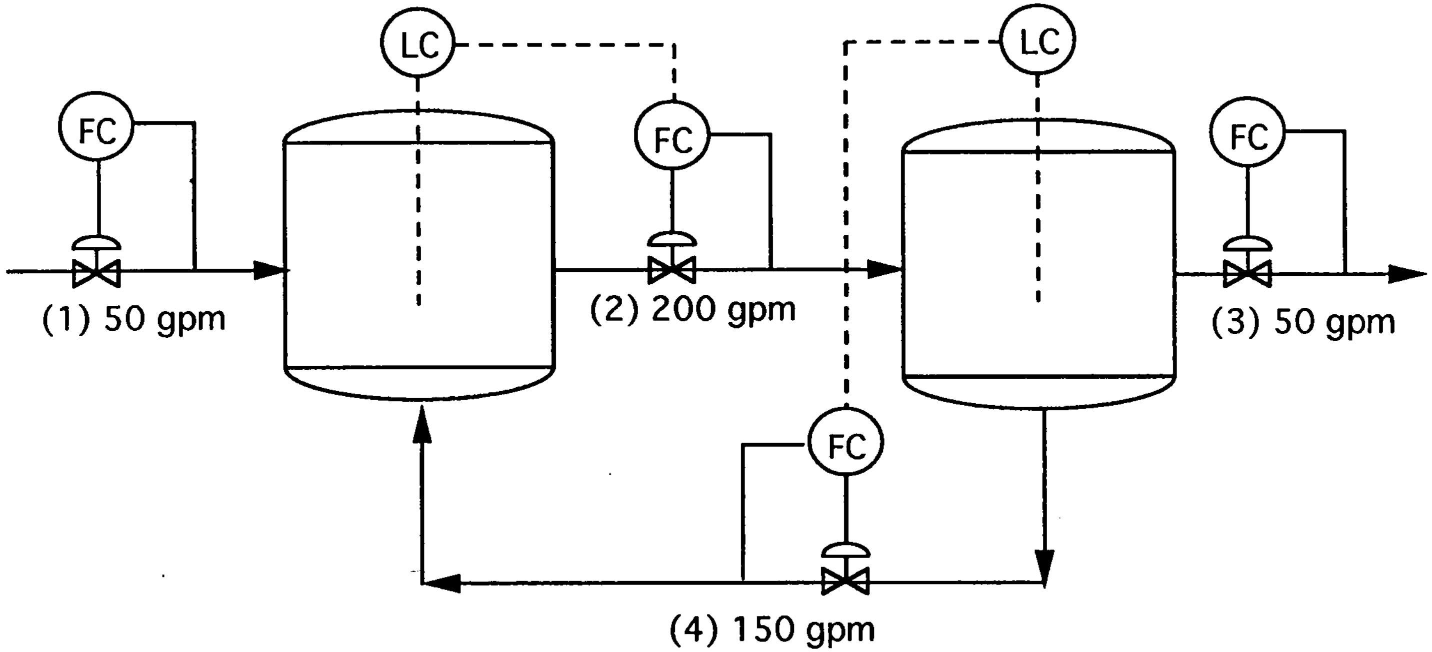

Figure A.3: Unworkable Inventory Control Scheme

However, the control strategy shown in Figure A.3 cannot possibly work for the same reason that the strategy in Figure A.1 could not work--the input and output rates for the whole system are independently set. Thus, if the average flow rates of Streams 1 and 3 are actually 50.01 and 49.99 gpm, respectively, material will accumulate somewhere in this system until one of the tanks overflows. The flaw in this control strategy is another that should be found and eliminated during the process design phase, not during operation. However, the control engineer must be willing to take the time to evaluate the overall material balance and observe that the control strategy fixes both the input and output of the whole system, even though there are individual inventory controls on each of the units. This example illustrates the reason that we emphasize control design from a plant-wide perspective: although each individual unit operation in Figure A.3 has a control strategy that looks feasible, the strategy will not work when the units are connected!

All of our examples so far have illustrated the most basic purpose of a control loop: to move variation from a controlled variable to a manipulated variable. In Figure A.2, for example, the level control loop reduces the variation in tank level by causing variation in one of the flow rates. The heat exchanger shown in Figure A.4 also illustrates this principle. The flow and temperature transmitters (FT and TT) show that the inlet conditions of the process fluid are not constant under normal operating conditions. If the steam supply rate is fixed and there is no control strategy, as in Figure A.4(a), the outlet temperature of the process fluid will also vary to reflect those changes. A temperature controller (TC) on the outlet temperature, as in Figure A.4(b), will cause the steam flow rate to vary instead of the outlet temperature. The addition of any control strategy moves variation from the controlled variable to the manipulated variable. This phenomenon is called "propagation of variation" or "transformation of variation."

Figure A.5 demonstrates how an inventory control strategy propagates variation throughout a plant. In Figure A.5, operators set the flow rate setpoint for Stream 1. The

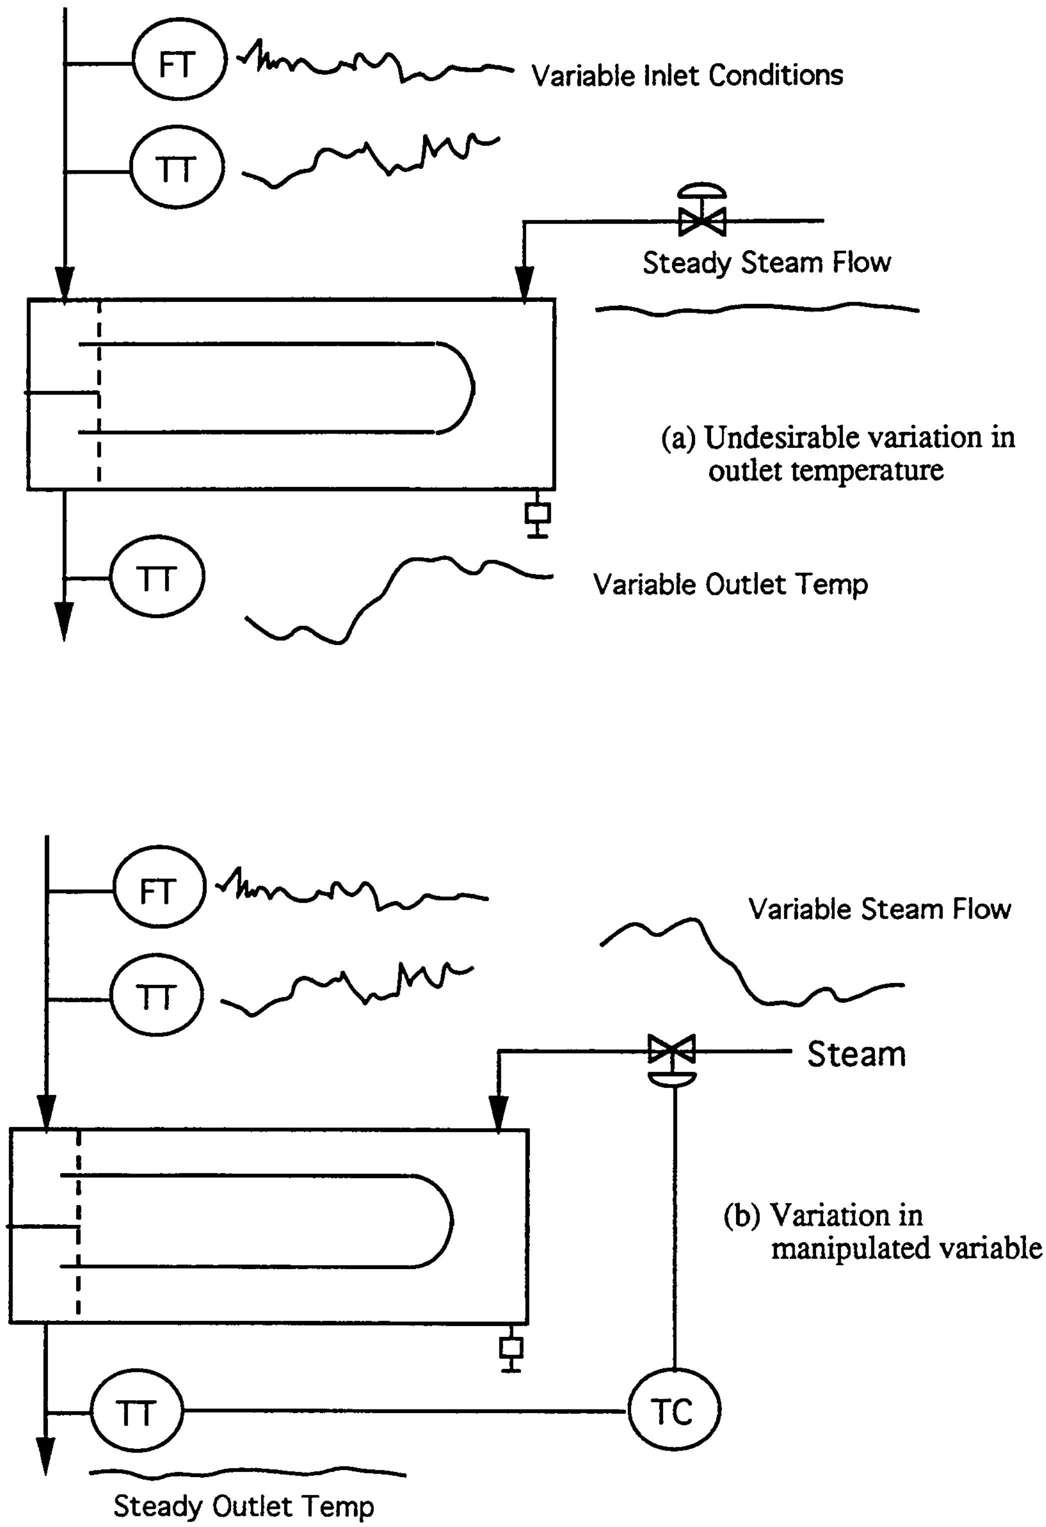

Figure A.4: Transformation of Variation in a Heat Exchanger

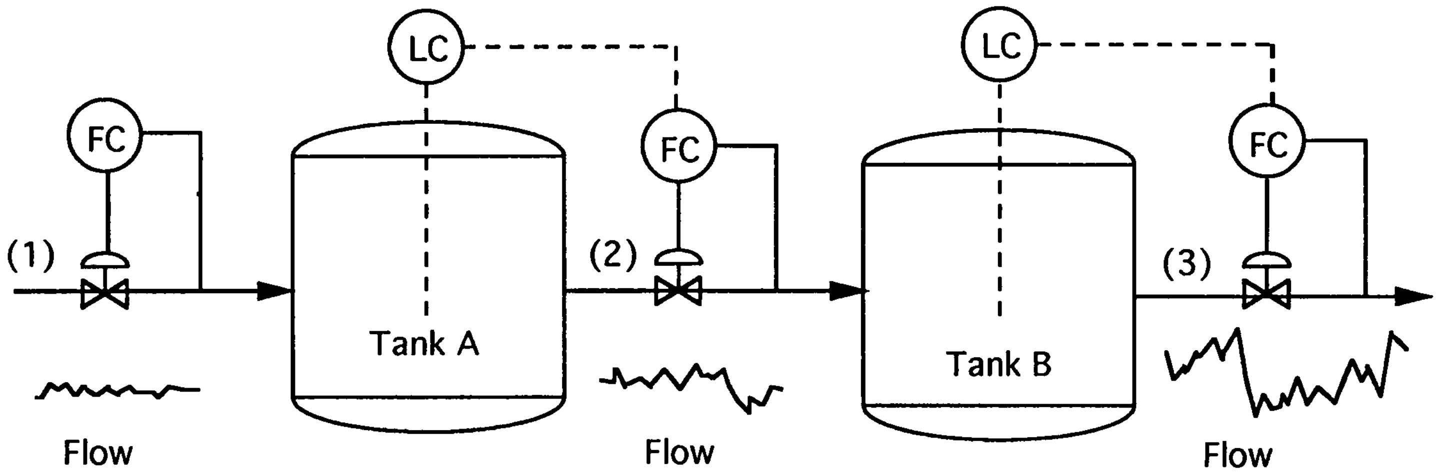

Figure A.5: Variation Propagated Forward by Level Control System control strategy controls tank inventories by setting the flow rate setpoints for Streams 2 and 3. Because no control loop moves variation into the setpoint of Stream 1, the flow rate of that stream will be fairly constant. Any disturbance to the material balance for Tank A will cause variation in the flow rate for Stream 2. Disturbances to the material balance for Tank B (including variation in Stream 2, the inlet stream for that tank) will cause variation in the flow rate for Stream 3. Because the level control loop for Tank A moves variation into the inlet stream for Tank B, disturbances in Tank B or in Tank A will cause variation in Stream 3. In general, when successive streams in a flow path are used for inventory control, the flow rate variation will be greatest for the streams farthest from a flow rate which is set manually, as is the flow rate of Stream 1 in Figure A.5.

Control engineers must understand this transformation of variation in order to design stable control strategies. Certain operations in a chemical process are very sensitive to flow rate changes. Those operations will be most stable when their feed flow rates are constant. If Tank A were a highly exothermic reactor instead of a tank, it would be unwise to use the inventory control strategy shown in Figure A.6 because that strategy causes the reactor feed stream to be the most variable stream in the process. These observations support our next design principle:

> Design Principle #3: *Arrange inventory control loops to direct variation away from the unit operations that are most sensitive to variation in flow rates.*

Simple tanks were sufficient to illustrate the importance of the first three design principles, but the next three require a more complicated process model. The process

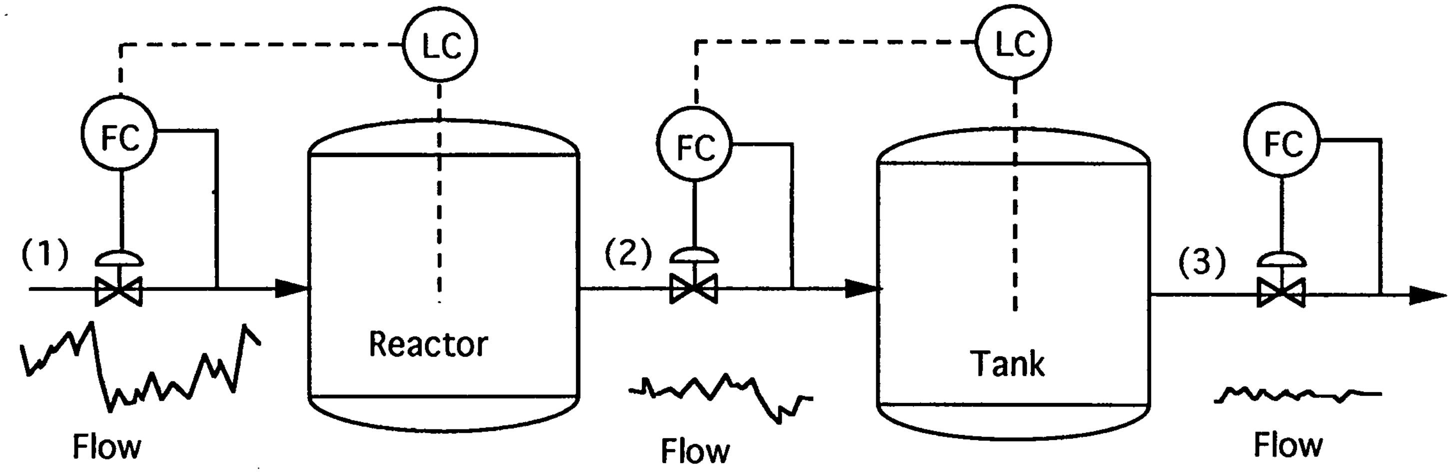

Figure A.6: Variation Propagated Toward Sensitive Unit Operation

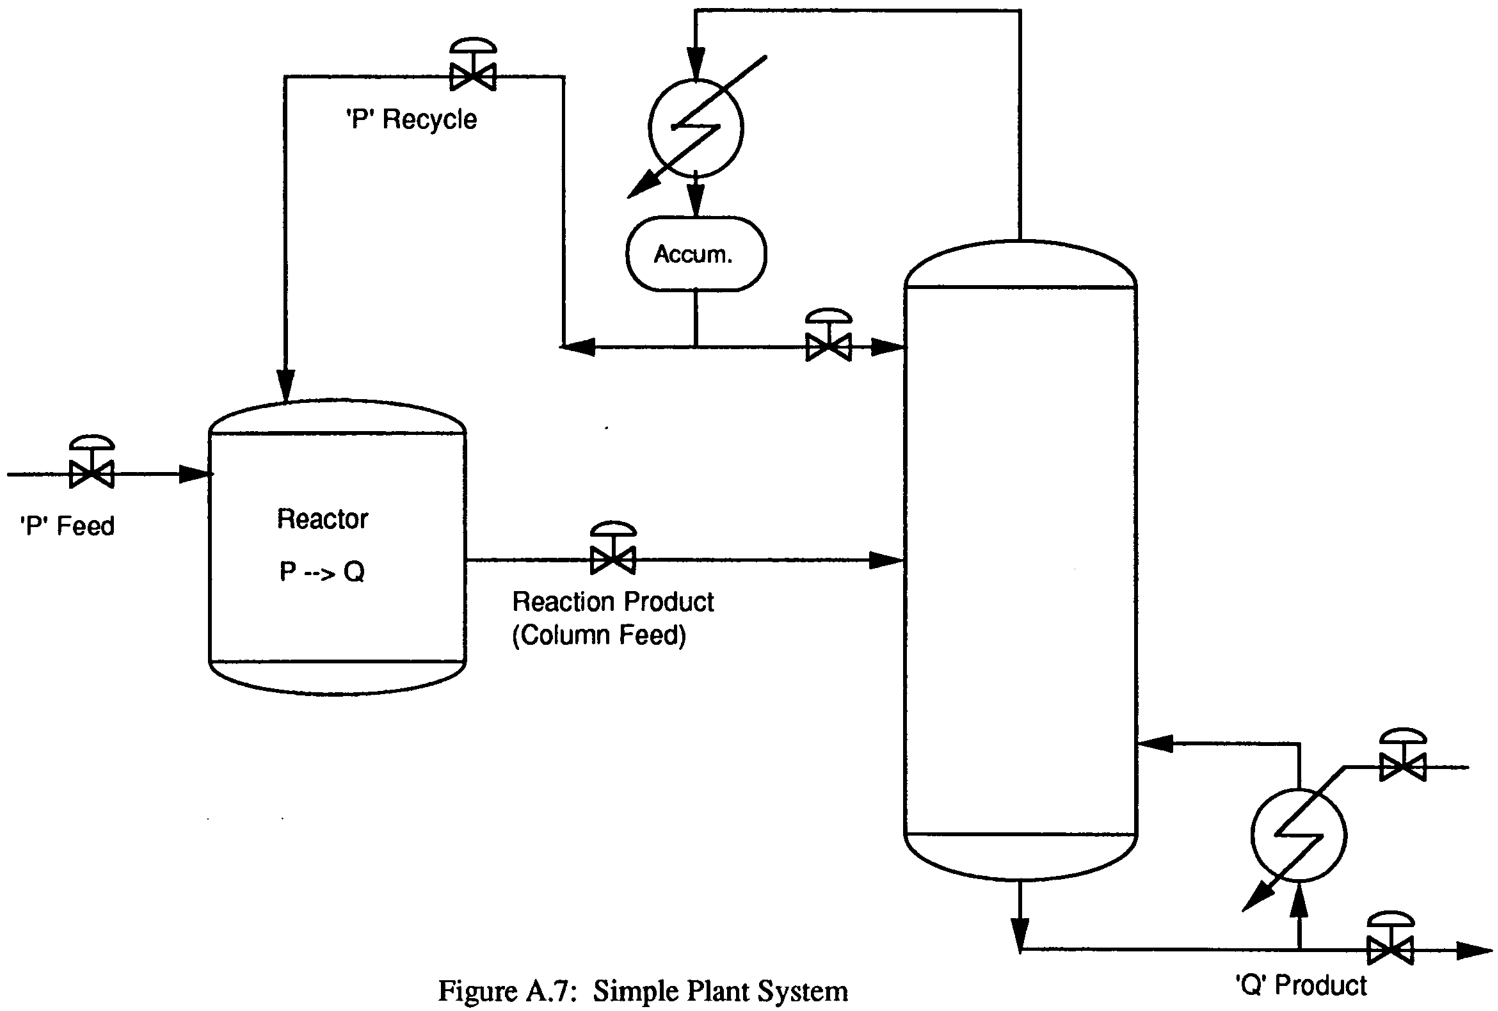

Figure A.7: Simple Plant System

shown in Figure A.7, a reactor and distillation column, is typical of a small industrial process. Certain assumptions will make this process easy to analyze:

1. The feed stream to the reactor consists of pure P, a feed stock.
2. A single, elementary reaction, P-->Q, takes place within the reactor with a rate coefficient of Arrhenius form.
3. The amount of component P in the product stream is small compared to the amounts in other process streams; virtually all of component P that enters the distillation column exits in the distillate.
4. Temperature control in the reactor is precise over a reasonable range of reactor temperatures.
5. No reaction takes place in the distillation column.

The most subtle control problem in this plant is probably the inventory control for component P, so this analysis will cover only those control loops in the primary process path for component P (feed stream, reactor, column feed stream, and distillate/recycle stream). The control engineer must choose some set of controlled variables such as these:

1) plant throughput rate;
2) reactor level (because level is a measure of bulk inventory);
3) reactor temperature;
4) accumulator level (another bulk inventory);
5) some measure of a component inventory.

This plant requires other control loops to control the product composition, maintain the reboiler level, and perform other functions, but these other loops do not directly concern inventory control for component P.

A well-intentioned engineer might try to meet these objectives with a control strategy such as the one shown in Figure A.8. This design appears to perform all the necessary functions: the P feed flow controller (FC) is available to change the production rate; there

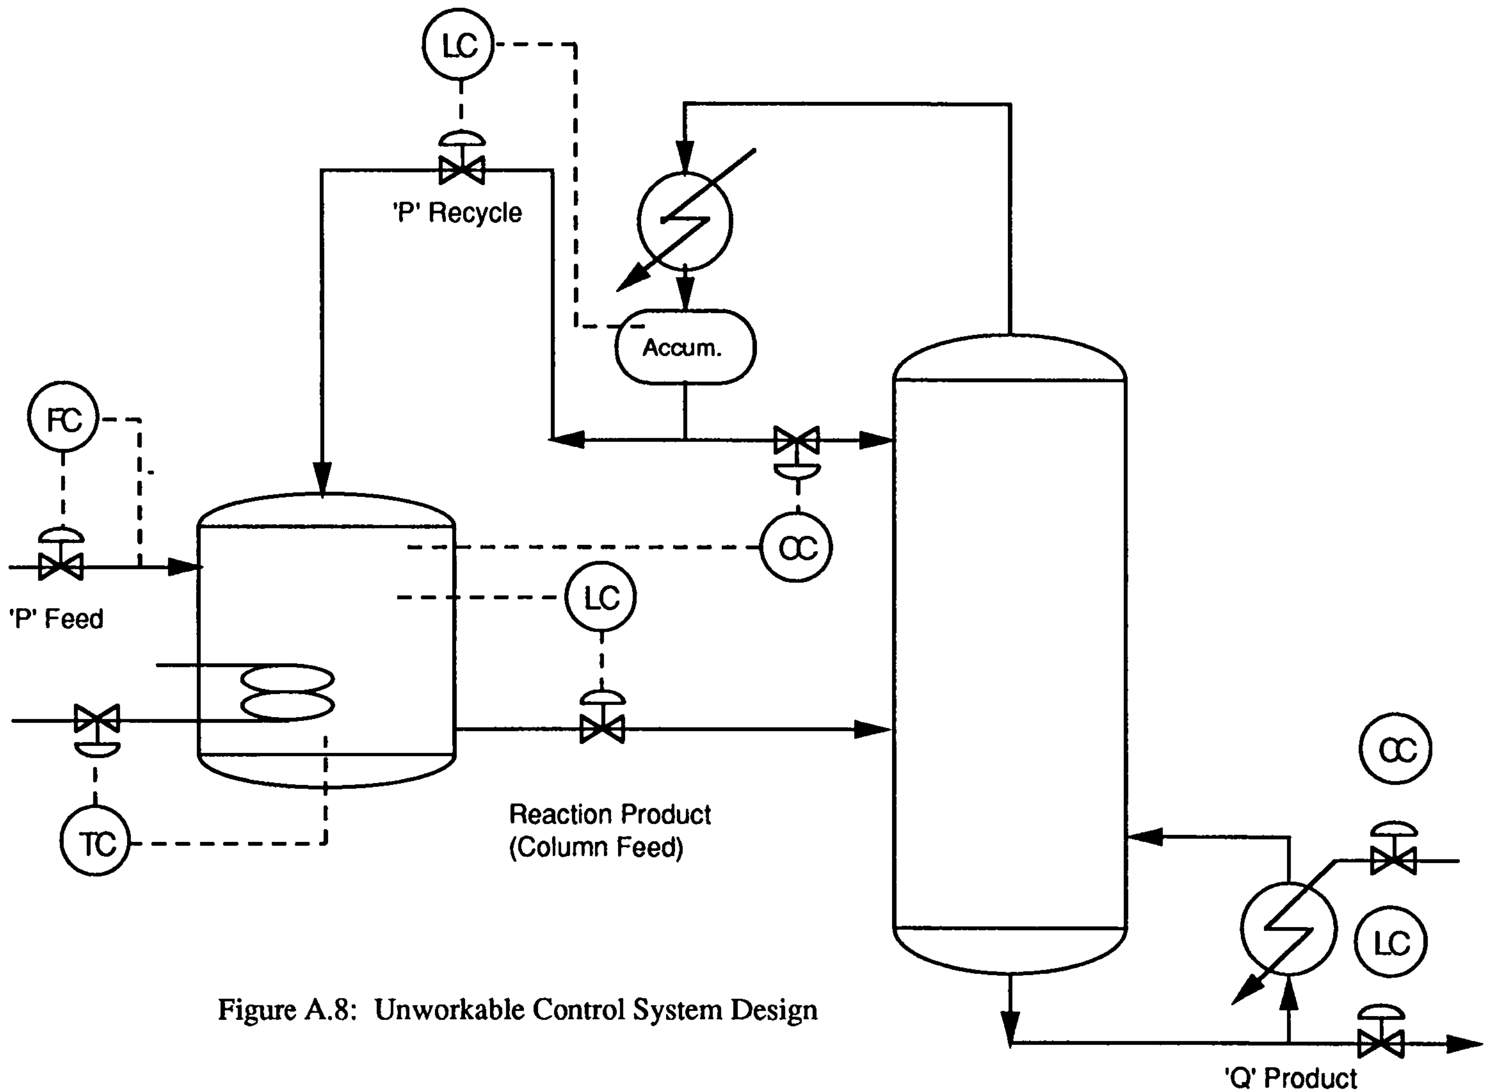

Figure A.8: Unworkable Control System Design are viable level controls (LC) on the reactor and accumulator; there is a heating (or cooling) medium for temperature control (TC); and there is a composition controller (CC) to maintain a constant inventory of component P in the reactor. There is nothing obviously wrong with this control strategy design, and it might even keep the process running for a while. Inevitably, though, it will fail, and engineers who look carefully at the process data will observe steadily worsening symptoms leading up to the failure. One possible scenario for this plant is that the composition control loop will call for successively more (or less) reflux, then drift away from its setpoint after the reflux valve is completely open (or closed). However the problem may manifest itself, there is an underlying cause that might not be immediately obvious to a process engineer because it is not directly measured by the process.

The flaw in the control strategy in Figure A.8 is once again a material balance problem. This time, the problem is the accumulation of a single component instead of a bulk amount of material. Equation 1, the fundamental material balance equation, applies to each component of a system as well as to the bulk:

$$\text{Input - Output + Generation - Consumption = Accumulation.} \qquad (1)$$

The amount of component P entering this plant is simply the rate of feed. Only a negligible amount (close to zero) exits in the product stream. The reaction does not generate any amount of P, but it consumes an amount equal to the total rate of reaction

$$R = V{*}k_0{*}\exp(-E/RT){*}X_P \qquad (3)$$

where

R = rate of reaction and consumption of P;

V = reactor volume;

k_0 = pre-exponential factor for the reaction rate constant;

E = reaction activation energy;

R = universal gas constant;

T = absolute temperature;

X_P = mole fraction of P in the reaction mixture.

Thus the material balance equation for component P over the entire plant simplifies to

$$(\text{Feed}) - V*k_0*\exp(-E/RT)*X_P = \text{Accumulation}. \tag{4}$$

The problem with this material balance is that every term on the left hand side of the equation is fixed. A simple flow controller sets the feed rate; a level controller fixes the reactor volume; k0, E, and R are physical constants; a temperature controller holds the reaction temperature constant; and a composition controller fixes the mole fraction of P in the mixture. Because of instrumentation limitations, it is impossible for those terms to combine in such a way that the accumulation of component P in the system is exactly zero. The amount of component P in this system will slowly build up or deplete, just as if this whole system were the tank from Figure A.1(a), with all its inputs and outputs independently set. If, over time, the amount of component P in the feed stream is even slightly more or less than the amount that reacts, there is no way for the system to get rid of the excess or make up the deficit.

The control engineer working on the process design may be the only person well qualified observe that the control strategy in Figure A.8 will not work, even though it looks good at first glance. He must be able to look at the process design and think about component inventories, though, before spotting the problem. Because every process with a separation unit and a recycle stream has the potential to accumulate unmeasured, recycled components, the following design principle is especially important for processes with recycle loops:

Design Principle #4: *A control strategy must allow bulk AND component inventories to regulate themselves within operable limits. Check the material balance for every component around every unit and recycle loop to make sure there are mechanisms by which the system can compensate for disturbances or small imbalances.*

Let us assume for simplicity that the feed rate of component P is slightly higher than the reaction rate. How can the inventory of component P build up within the system when the whole purpose of the composition control loop in Figure A.8 is to prevent such an occurrence? Unfortunately, instruments and computer logic cannot countermand the laws of nature; if component P cannot accumulate within the reactor (where it is monitored) and if the control strategy does not allow it to leave, it must accumulate in the distillation column instead.

The composition controller has the unfortunate effect of fixing the overall reaction rate. If there were no composition controller, an excess of P would cause the mole fraction X_P in Equation (3) to rise, and the reaction rate (R) would rise accordingly until the amount of P consumed was equal to the amount fed. Within some limits, the system would regulate itself to compensate for an imbalance between the rates of feed and reaction. Without this self-regulating effect, there is no way for the system to rid itself of the excess.

If the composition controller in Figure A.8 causes so many difficulties, perhaps it should be removed to let the system regulate its own composition. The control strategy in Figure A.9 is identical to the one in Figure A.8 except that it has no composition controller

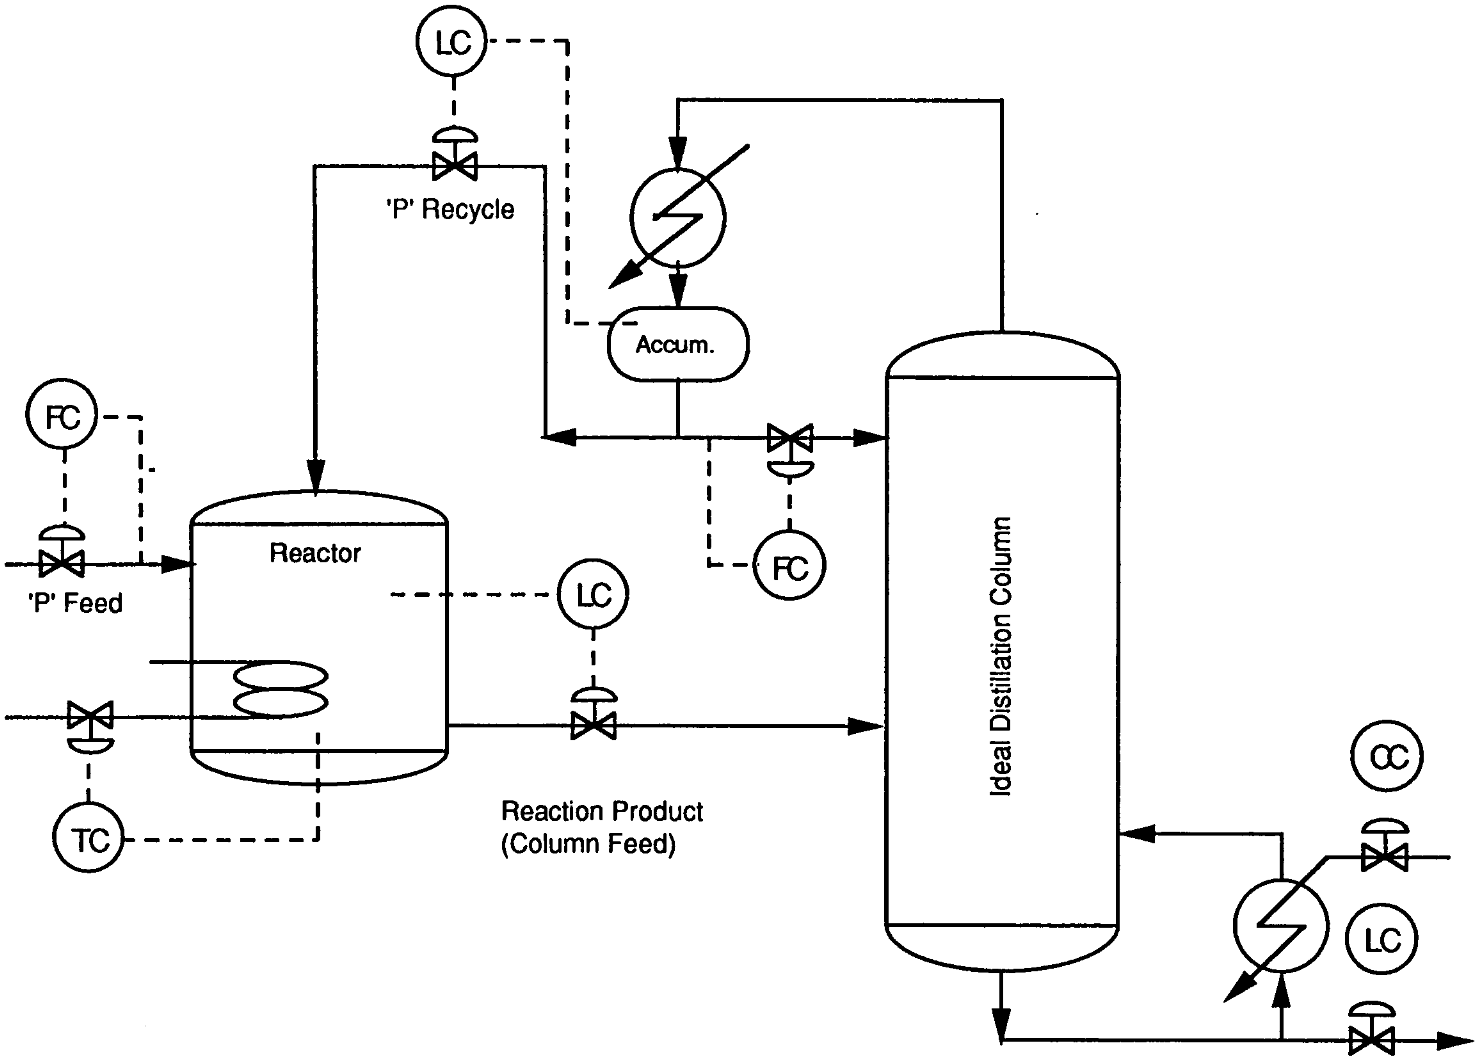

Figure A.9: Marginally Acceptable Control System Design on the reactor. This control strategy satisfies the component material balance by relying on fluctuations in X_P (and, therefore, in the reaction rate) to eliminate small excesses or deficits in the reactor feed. The inventory of component P regulates itself even without a controller.

Unlike the control strategy shown in Figure A.8, the strategy in Figure A.9 will operate satisfactorily at the design condition (the base-case operating point for which the process was originally designed). In the real world, though, it is commonly necessary for operators to change a plant's operating point to make more or less product. In this control strategy, an increase in the rate of fresh feed will eventually produce an increase in the overall production rate. However, there is an upper limit to the amount of fresh feed that can be added. The reaction rate shown in Equation (3) is proportional to the mole fraction of component P in the reactor. At any fixed temperature and volume, there is a theoretical upper limit to the rate of consumption of P when the mole fraction X_P is 1.0 (its maximum value):

$$R_{max} = V*k_0*\exp(-E/RT). \tag{5}$$

Because R_{max} is the maximum rate at which component P can react, it is also the maximum rate at which P can be added to the system. (Remember that no P exits the system through the product stream, so all that enters must be reacted.) In practice, it is not possible to achieve the theoretical maximum rate R_{max} given by equation 5. To understand why not, first consider how the other plant variables change when the reactor feed rate is increased:

1) The reactor outlet rate increases to maintain the reactor level.
2) Distillate and bottoms rates increase in an amount proportional to the increase in the column feed.

3) X_P increases in the reactor because the amount of component P entering the system is greater than the amount being reacted.

4) Reaction rate increases by an amount proportional to the increase in X_P.

5) The rate of distillate flow increases further as the additional P in the column feed exits in the distillate.

Two separate effects increase the distillate rate: the increase in the column feed flow rate and the increase in the mole fraction of P in the column feed. Thorough analysis of all the steady-state material balance equations reveals that the relationship between recycle rate and plant feed rate has the general shape shown in Figure A.10. When the feed rate is equal to R_{max}, the required distillate flow rate is infinite. Because the distillation column has a limited vapor capacity, the actual maximum production rate is limited to some rate lower than Rmax. Figure A.10 shows that a small increase in feed rate may or may not result in a large percentage change in the distillate rate, depending on how close the feed rate is to the theoretical maximum. If the design feed rate is at the point marked "X" on Figure A.10, an increase in the feed rate will cause a much larger percentage increase in the distillate rate.

Luyben (44, 47) uses the term "snowball effect" to describe this large change in distillate rate that results from a small change in the production rate. The snowball effect is a result of the reaction kinetics--a different reaction volume will shift the curve in Figure A.10, though the curve will always have this general shape. If the desired fresh feed rate is at the point marked "X" on figure A.10, perhaps a larger reactor would be appropriate so that small production rate shifts will not have such a drastic effect on the column operating conditions.

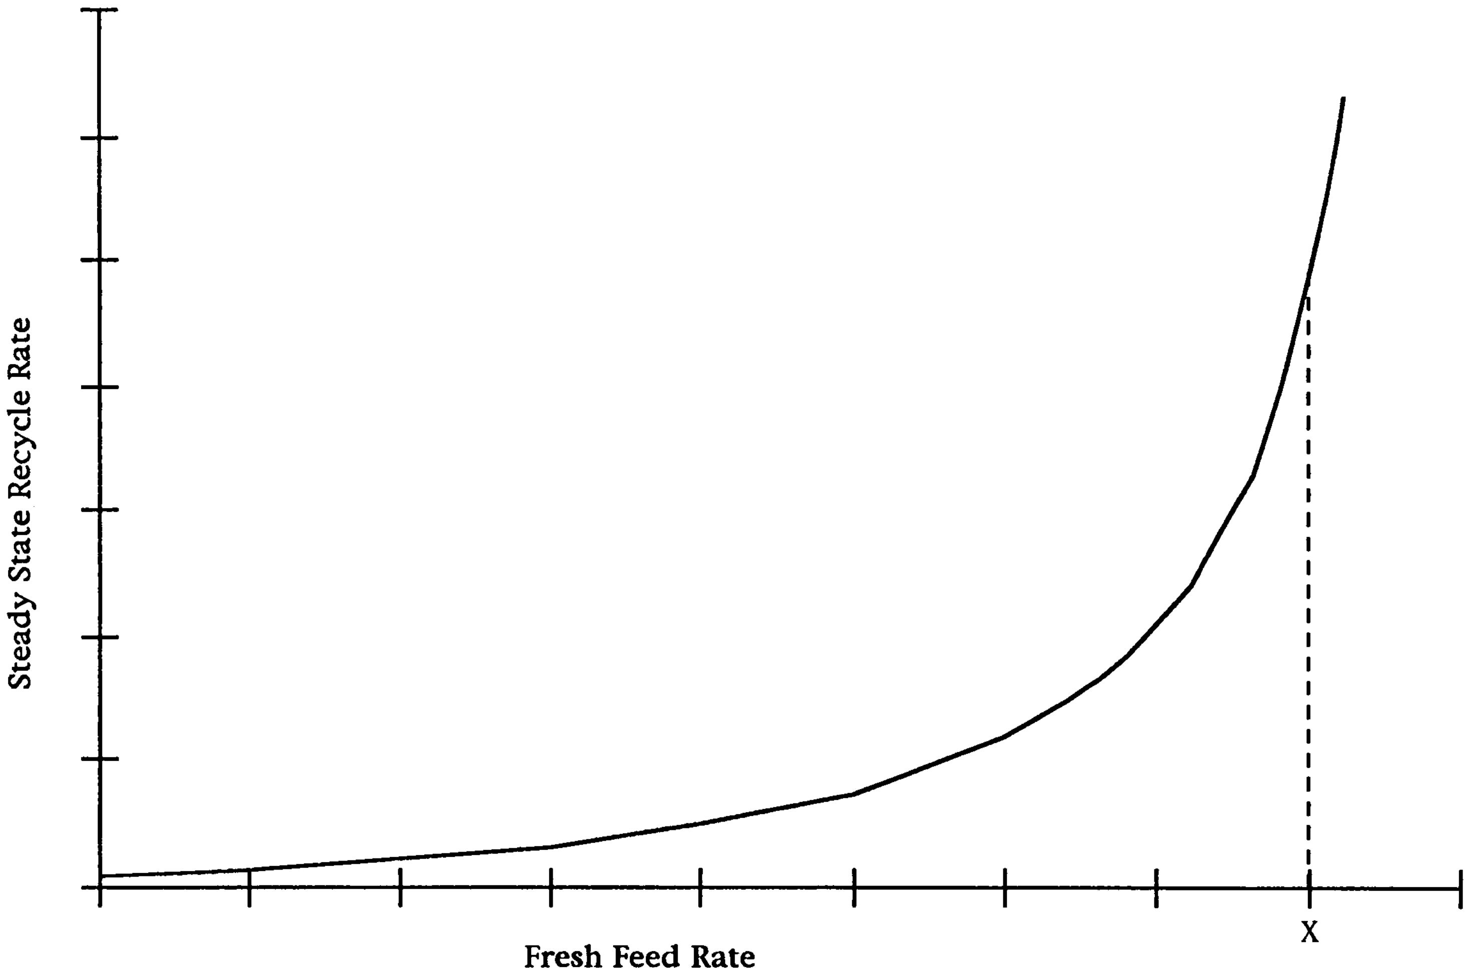

Figure A.10: Snowball Effect in Distillate

There is an important difference between the flaw in the control strategy in Figure A.8 and the flaw in the control strategy in Figure A.9. In Figure A.8, the control strategy violates a fundamental material balance constraint for component P. Because component P will either accumulate or deplete, the control strategy cannot possibly hold the plant at the desired operating point. The control strategy in Figure A.9 will hold the plant at the original design operating point. However, a change in the throughput rate will cause the distillation column to fail if the column cannot maintain the resulting steady state distillate rate. (Either too much or too little vapor will cause column failure.) The material balance violation in Figure A.8 will cause the plant to fail; the snowball effect in Figure A.9 might cause it to fail.

Luyben suggests that the following practice will eliminate the snowball effect:

> Design Principle #5: *If a plant contains a liquid recycle loop, put a simple flow controller somewhere within the recycle path. (This principle applies to systems whose yield is limited by reaction rate, not by thermodynamic equilibrium.)*

Luyben's principle suggests that perhaps the control strategy in Figure A.11 is a better way to control the plant. The flow controller on the column feed has two beneficial effects: it prevents the flow rate in the recycle loop from increasing beyond reasonable bounds, and it ensures that there is little variation in the flow rate entering the column.

What is the best way to increase the production rate using the control strategy in Figure A.11? It is natural for an engineer to want to change a flow rate to vary the throughput, but Equation (3) suggests two other possibilities not immediately obvious from

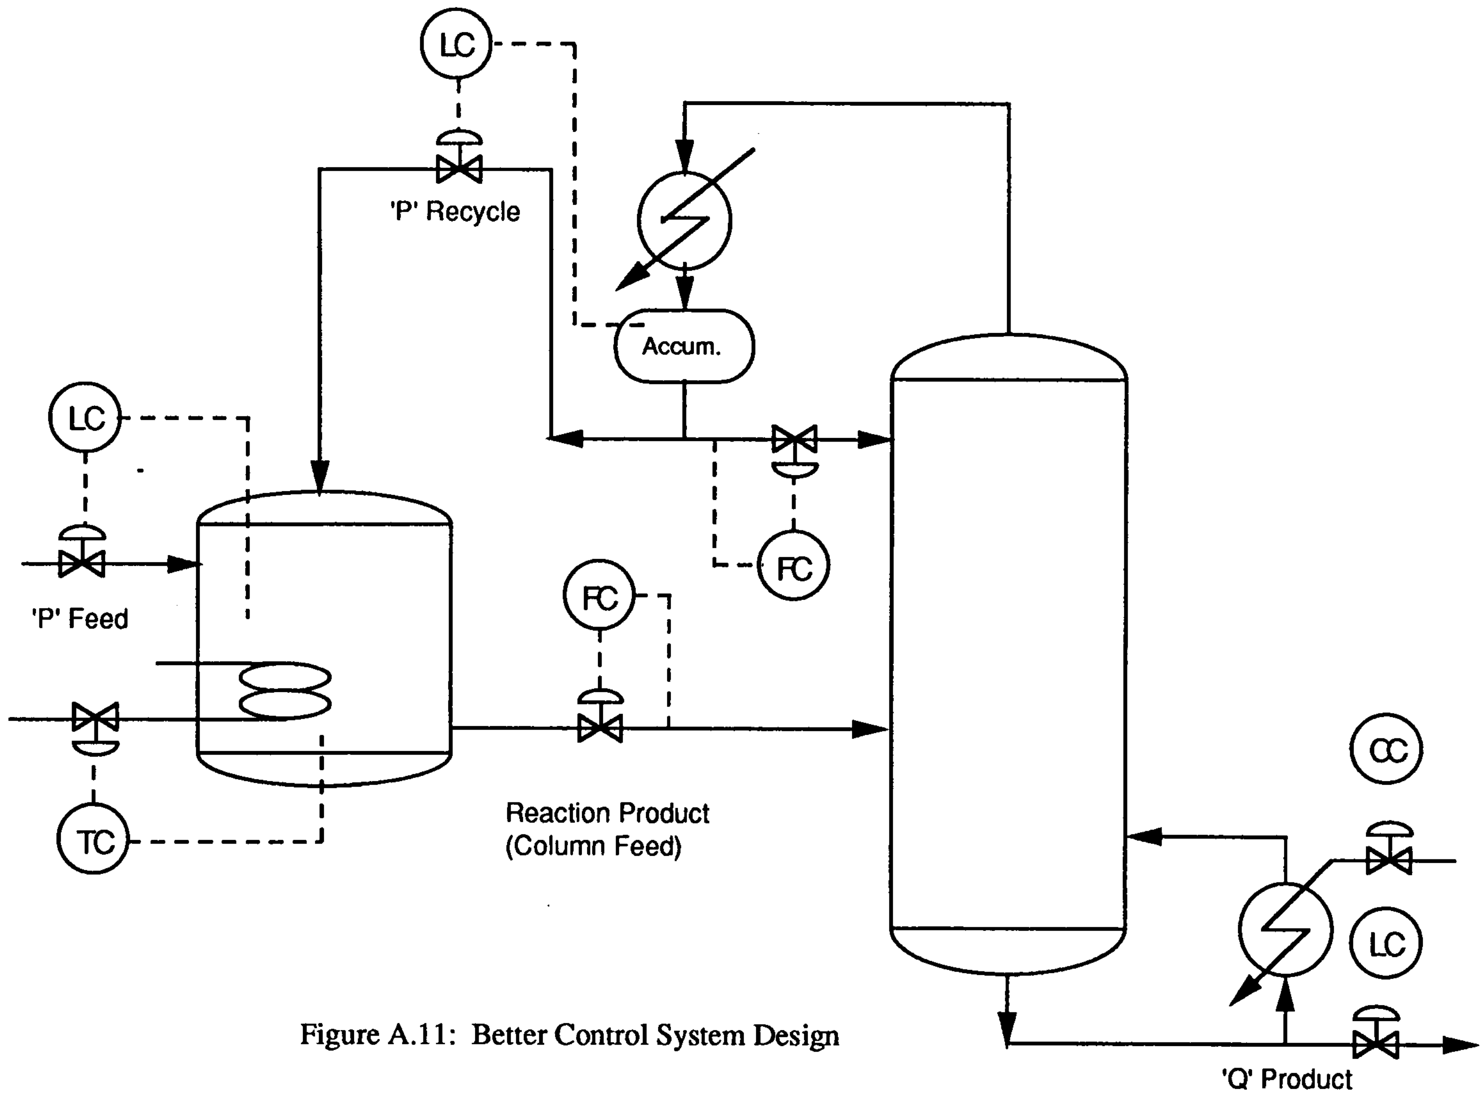

Figure A.11: Better Control System Design the figure. Raising the temperature setpoint or increasing the reaction volume (by raising the level setpoint) will increase the rate of generation of the desired product Q. Neither of these changes will result in any snowball effect. A more complicated way might be to increase the column feed rate and the reactor level setpoint by the same percentage. (This dual setpoint change prevents the snowball effect by maintaining a constant reactor residence time.) Our final design principle reminds us to use cascaded loops when appropriate:

> Design Principle #6: *Whenever a control loop is implemented, its setpoint becomes a potential manipulated variable.*

Although there are control loops in Figure A.11 on the reactor temperature, the column feed rate, the reflux rate, and the reactor level, the control engineer and process engineer must still determine setpoints for these loops. Those setpoints may be used just like any other manipulated variable to affect the operation of the plant. Physical understanding of the process may suggest novel ways to use these setpoints, such as simultaneously changing the column feed rate and reactor level setpoints to change the production rate.

The development of the control strategy in Figure A.11 is one example of a series of control strategy and process design decisions that can dramatically improve the dynamic behavior of a chemical plant. For many years chemical engineers have recognized that a plant's behavior depends on the tuning of individual control loops. It is apparent, however, that dynamic behavior also depends strongly on how the loops are chosen, not just how they are tuned. The accompanying research proposal describes efforts to determine specific methods for selecting the best inventory, throughput, and quality controls for a plant.

These design principles and the accompanying proposal emphasize the need for control engineers to have detailed knowledge of plant operations before designing a control strategy. Control engineers who understand the process and evaluate the correct material balances can avoid many of the control strategy design flaws that plague large plant-wide systems. Processes whose control strategies are designed without thought for these principles may suffer chronic costly shutdowns or require expensive, troublesome retrofits. An insightful control engineer can save a company significant expense by investing time in designing control strategies instead of simply operating them.

Appendix B: Dynamic Responses for Module 2 Control Strategies

The rankings in Chapter 4 for the Module 2 control strategies were based on numbers calculated from dynamic responses to disturbances. The criteria themselves were presented in Tables 4.3-4.5; this appendix shows the dynamic responses themselves.

All of the figures in this appendix show dynamic responses for the same 8 variables:

- Product flow rate
- Total composition of light components in product
- 'D' composition in product
- 'A' composition in first reactor (and feed to second reactor)
- 'B' composition in first reactor (and feed to second reactor)
- Ratio of steam flow to bottoms flow
- 'C' composition in distillate
- Pressure drop through distillation column

Figures B.1-B.12 show the dynamic responses of strategies 1-12 to a feed composition change of 10% 'A' introduced into the 'B' feed. (Strategies 1B, 2B, 6B, and 12B respond identically to their respective parent strategies following a feed composition change, so there are no separate graphs for them.) All of the strategies except for Strategy 3 rejected this composition change easily; Strategy 3 responded poorly because of factors described in Chapter 4.

Figures B.13-B.28 show the dynamic responses of all 16 strategies to a throughput increase of 5%. Note that maintaining constant steady-state reactor residence time (as in Strategies 1b, 2b, 6b, and 12b) improves the response for almost every variable.

Unlike the other strategies, Strategies 9 and 10 were unable to reject large feed composition changes involving excesses of 'B' instead of 'A.' Figure B.29 shows the dynamic response of Strategy 9 when 9% 'B' was introduced into the 'A' feed. The composition controller for this strategy evaluates the reactant ratio by monitoring the ratio of steam flow to product flow. As described in Chapter 4, the relationship between the controlled and manipulated variables for this loop is not monotonic because there are two feasible operating modes for this plant--an 'A'-rich mode and a 'B'-rich mode. Reactant 'A' is in excess at the base-case operating point for Module 2. As more 'B' is introduced via the feed composition disturbance, though, the plant moves toward the 'B'-rich region, for which the composition controller gain has the wrong sign. Note in Figure B.29 that the plant responds slowly during the first 25 hours or so, while the controller is taking the 'correct' control action. As the plant moves into the 'B'-rich region, however, the controller begins to take the wrong action--increasing the 'B' feed and sending the plant quickly out of control. The problem is a dynamic one--because the dynamics of the composition response are so slow, the controller must respond slowly to the disturbance. A larger controller gain would permit rejection of larger disturbances without moving into the 'B'-rich operating mode, but larger controller gains are unstable for smaller disturbances. Because of the slow loop dynamics, there is no controller gain that permits rejection of both large and small feed composition disturbances.

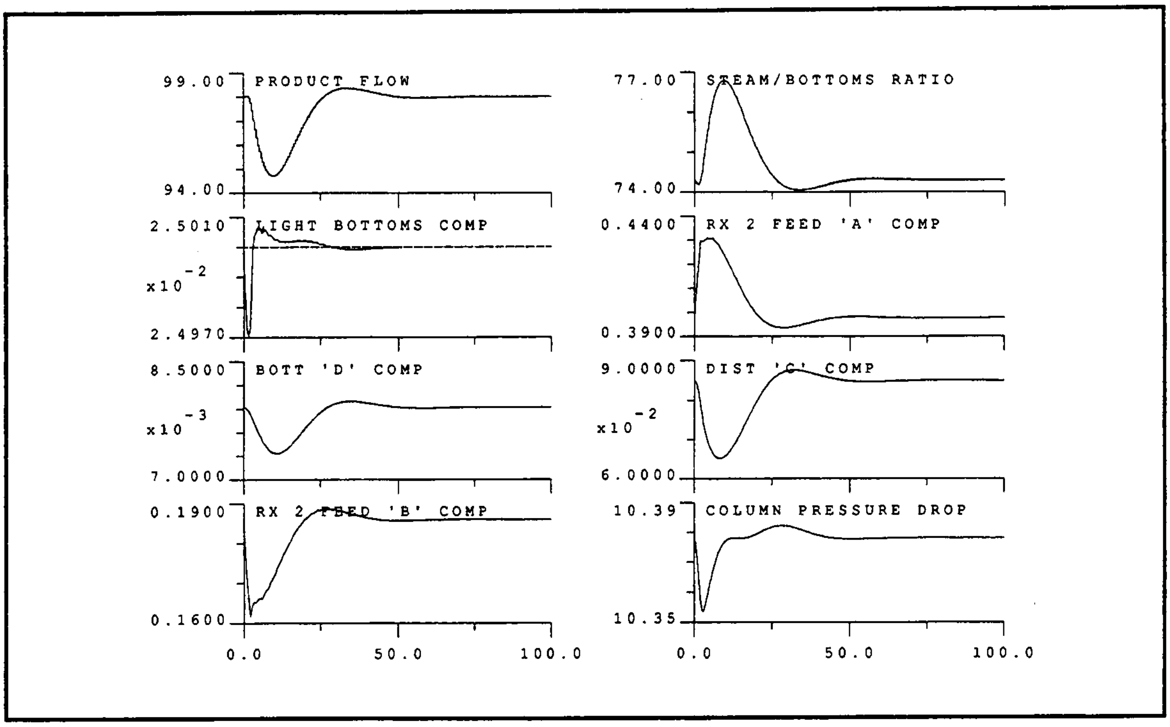

Figure B.1: Strategy 1; 10% A in B Feed

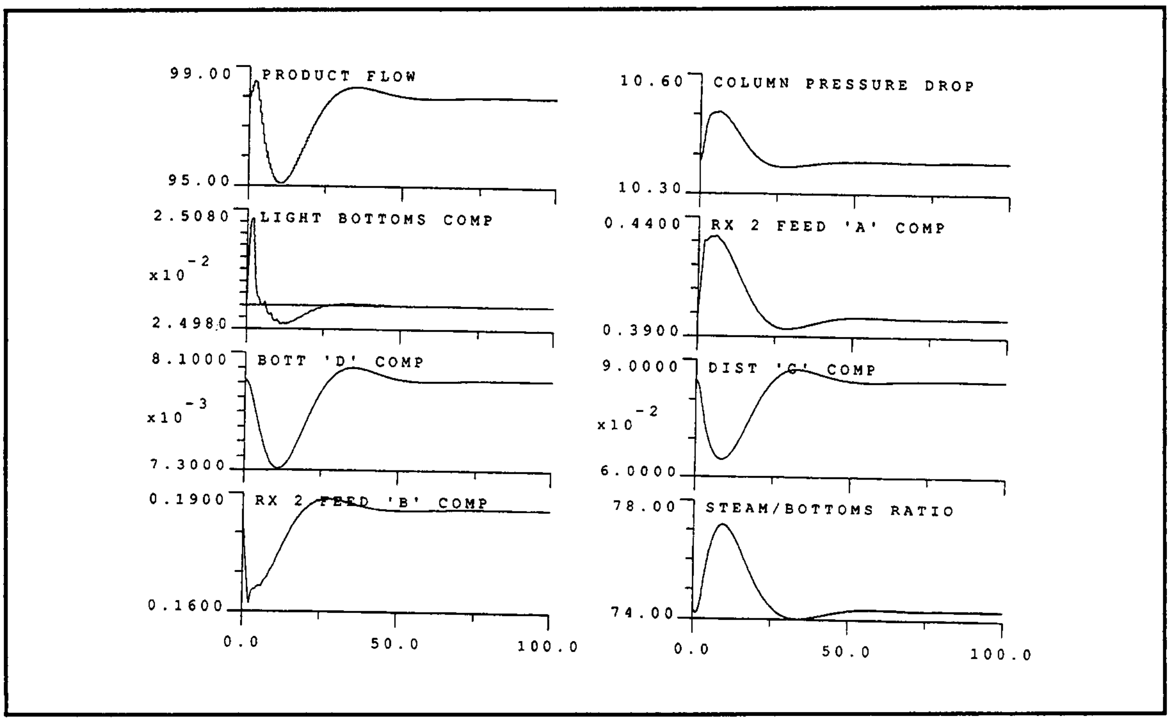

Figure B.2: Strategy 2; 10% A in B Feed

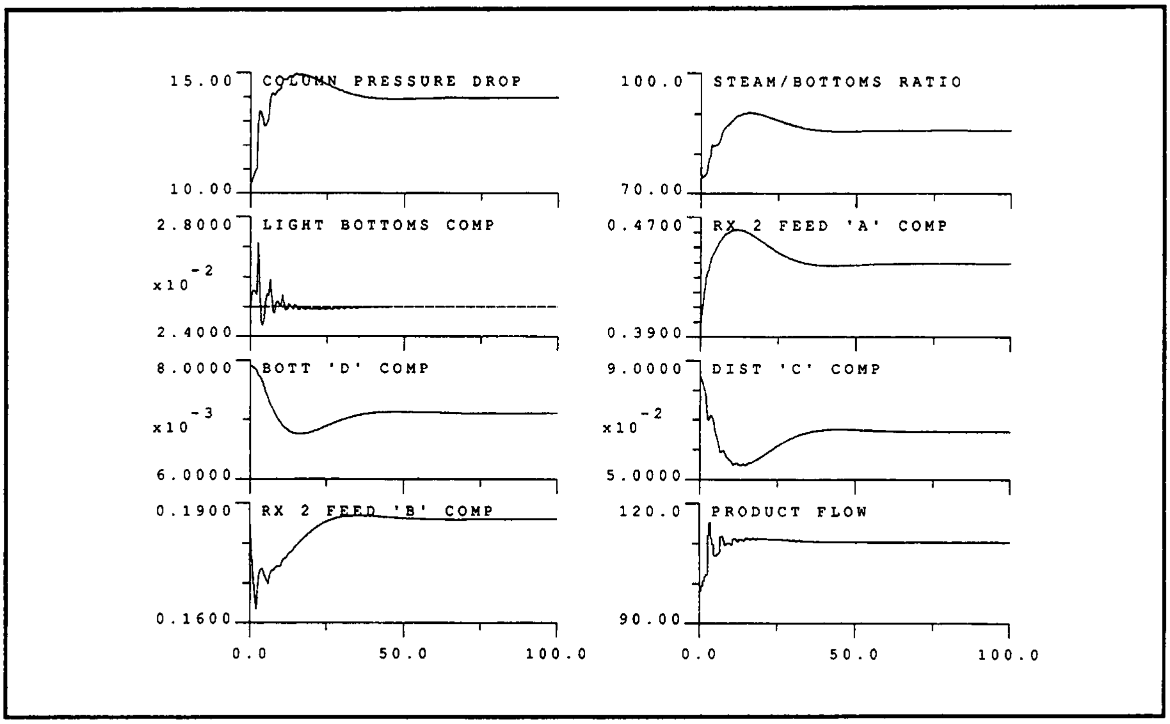

Figure B.3: Strategy 3; 10% A in B Feed

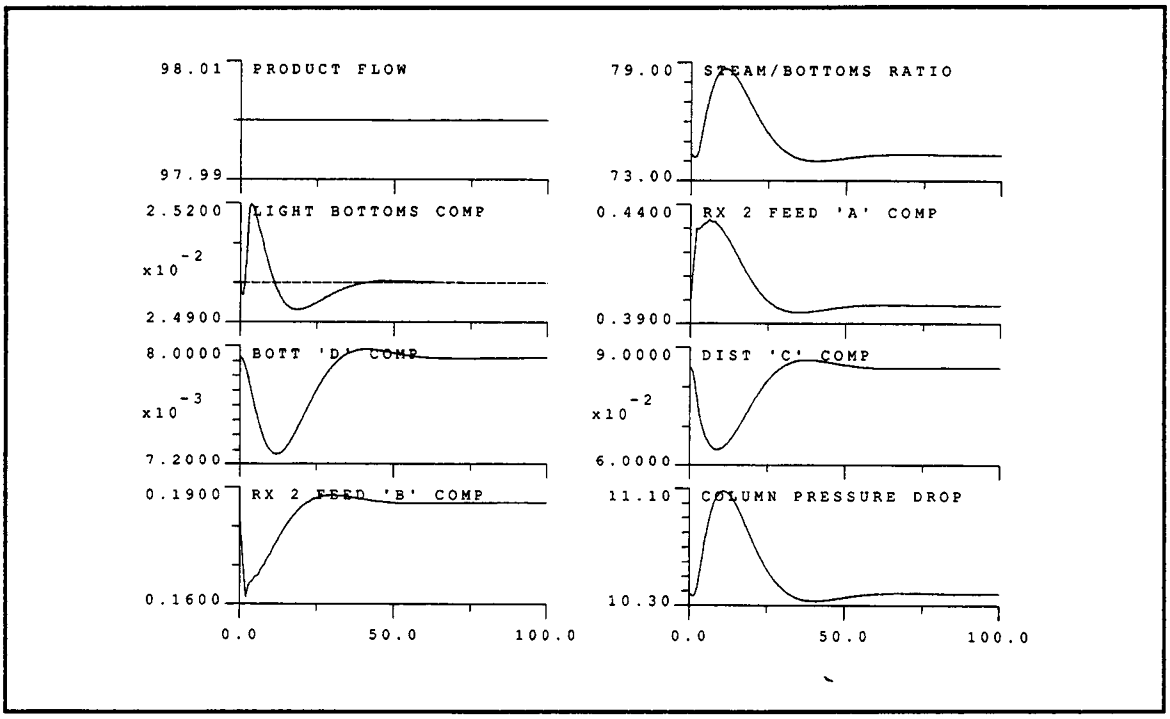

Figure B.4: Strategy 4; 10% A in B Feed

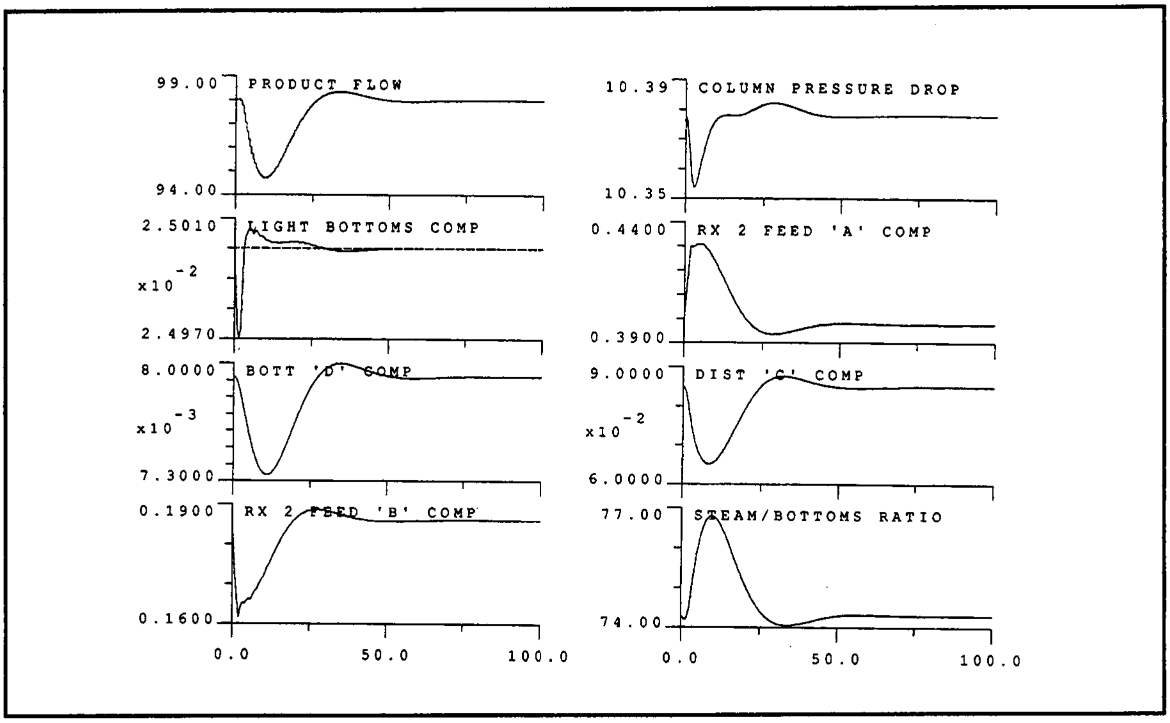

Figure B.5: Strategy 5; 10% A in B Feed

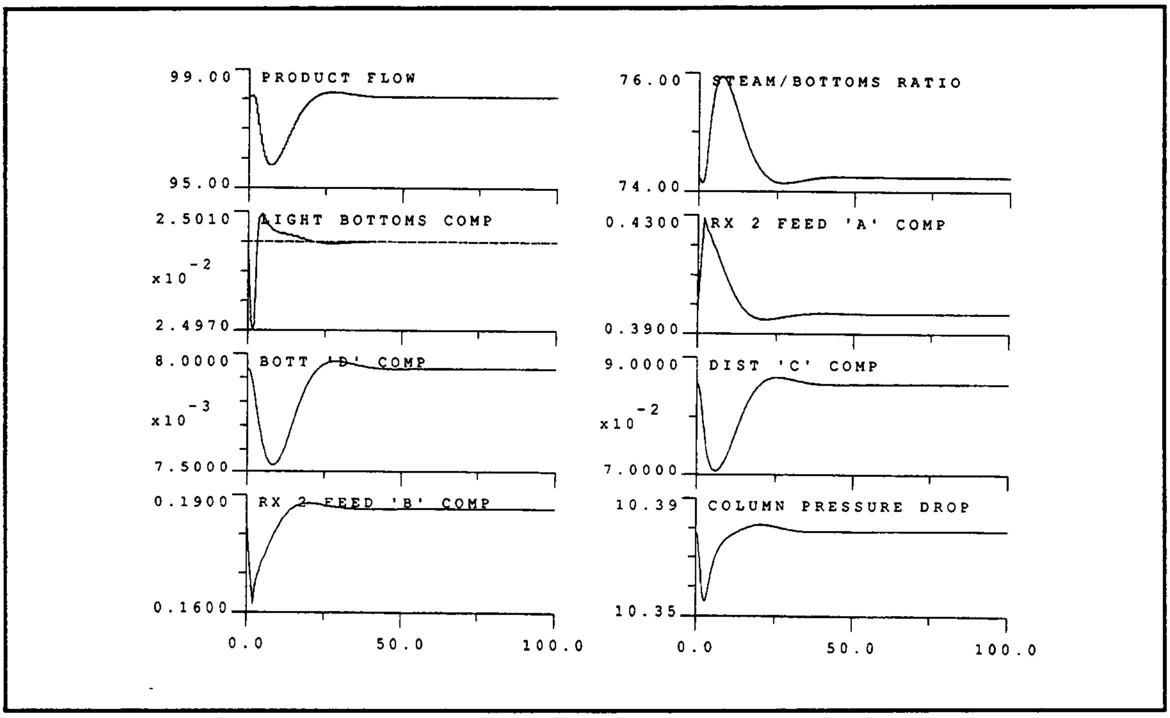

Figure B.6: Strategy 6; 10% A in B Feed

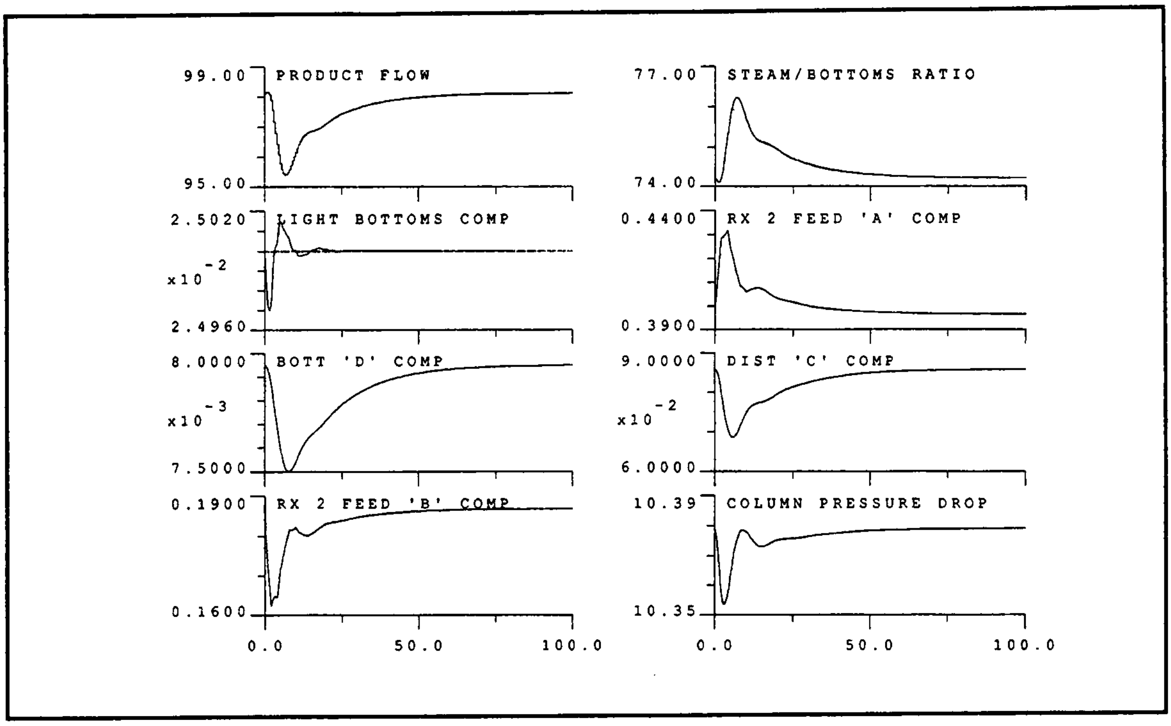

Figure B.7: Strategy 7; 10% A in B Feed

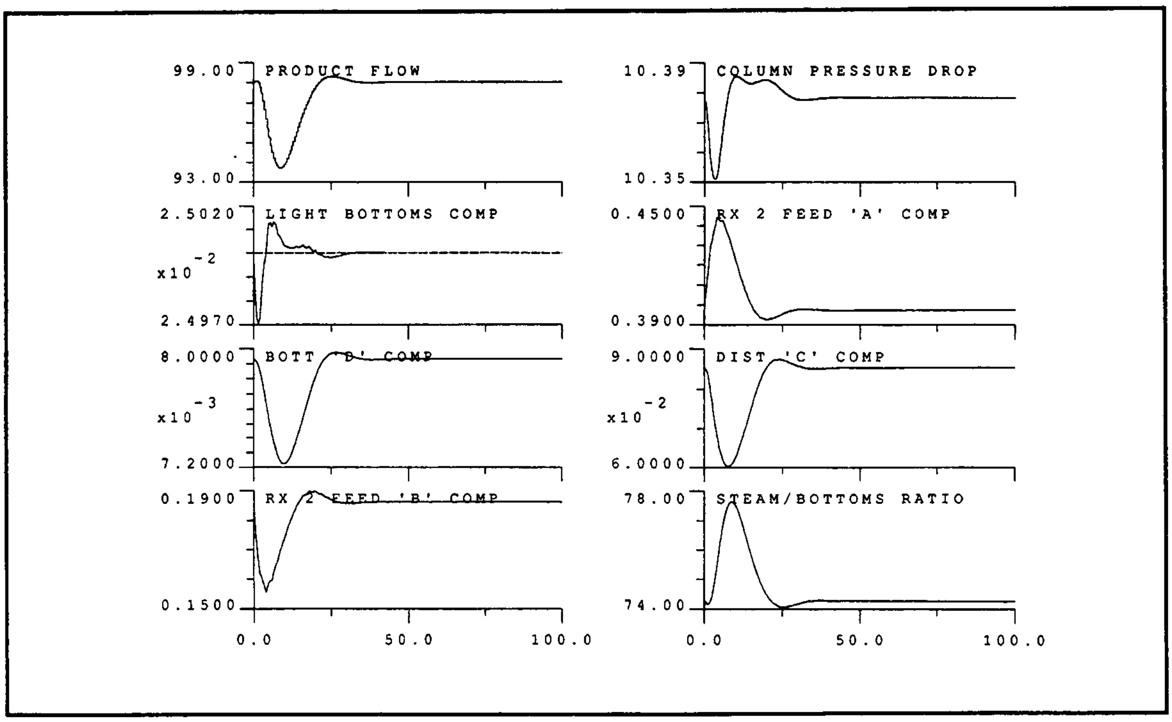

Figure B.8: Strategy 8; 10% A in B Feed

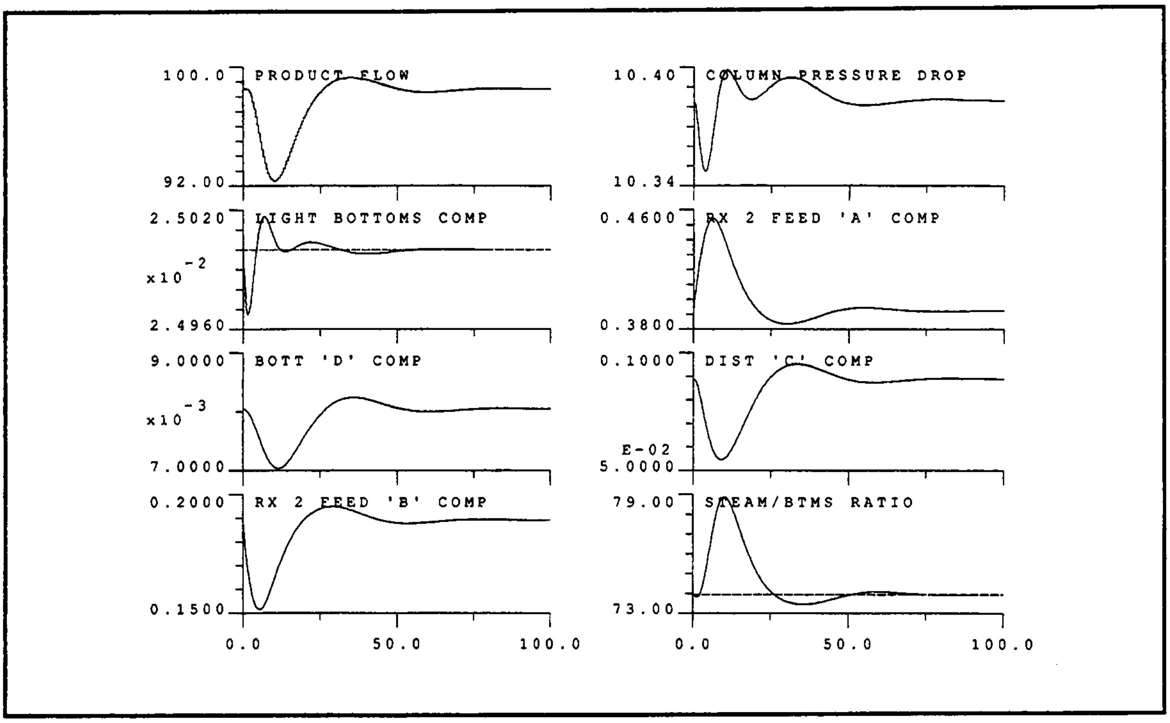

Figure B.9: Strategy 9; 10% A in B Feed

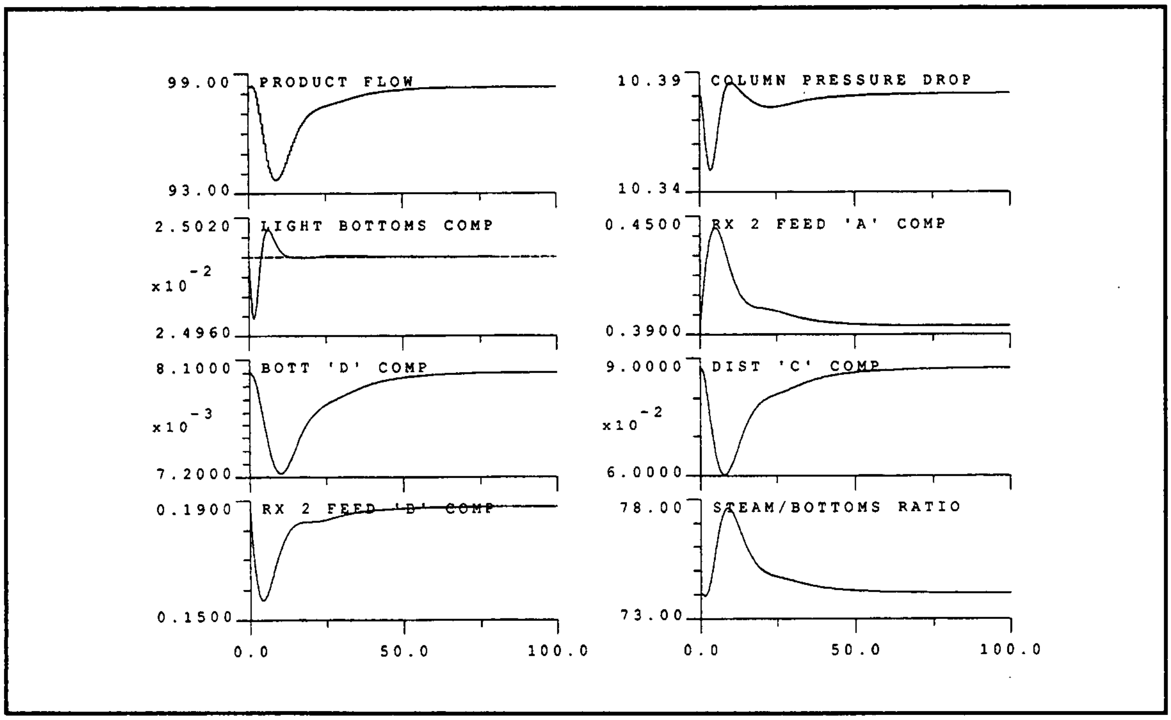

Figure B.10: Strategy 10; 10% A in B Feed

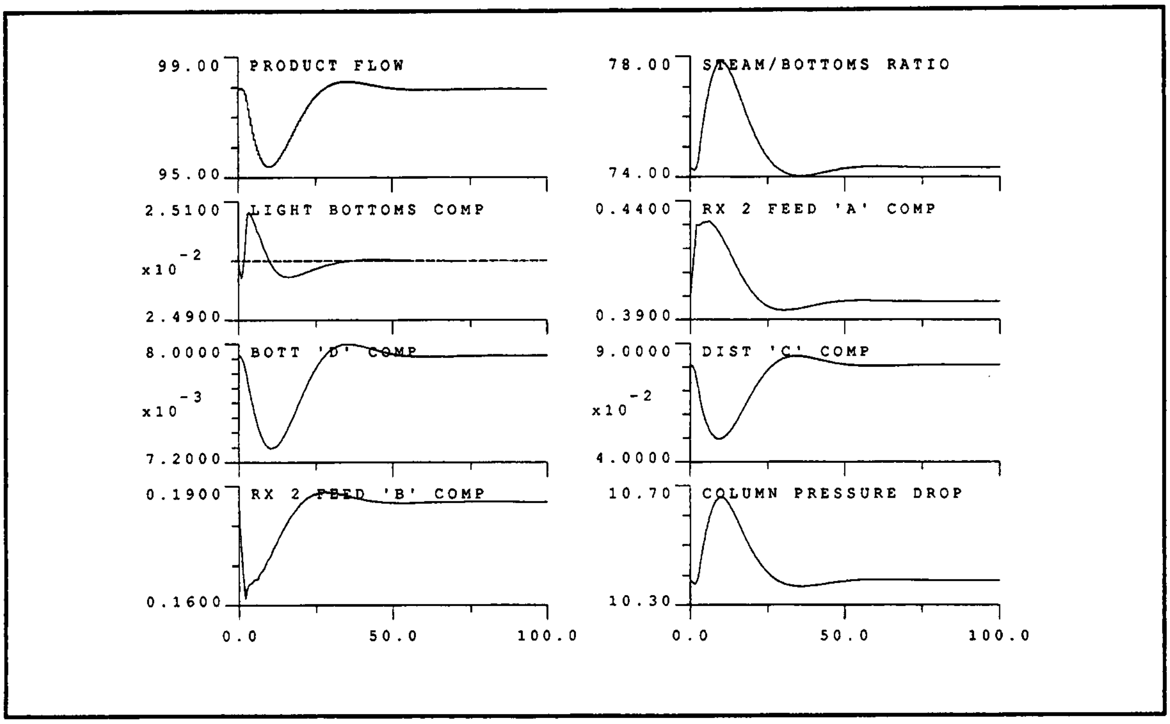

Figure B.11: Strategy 11; 10% A in B Feed

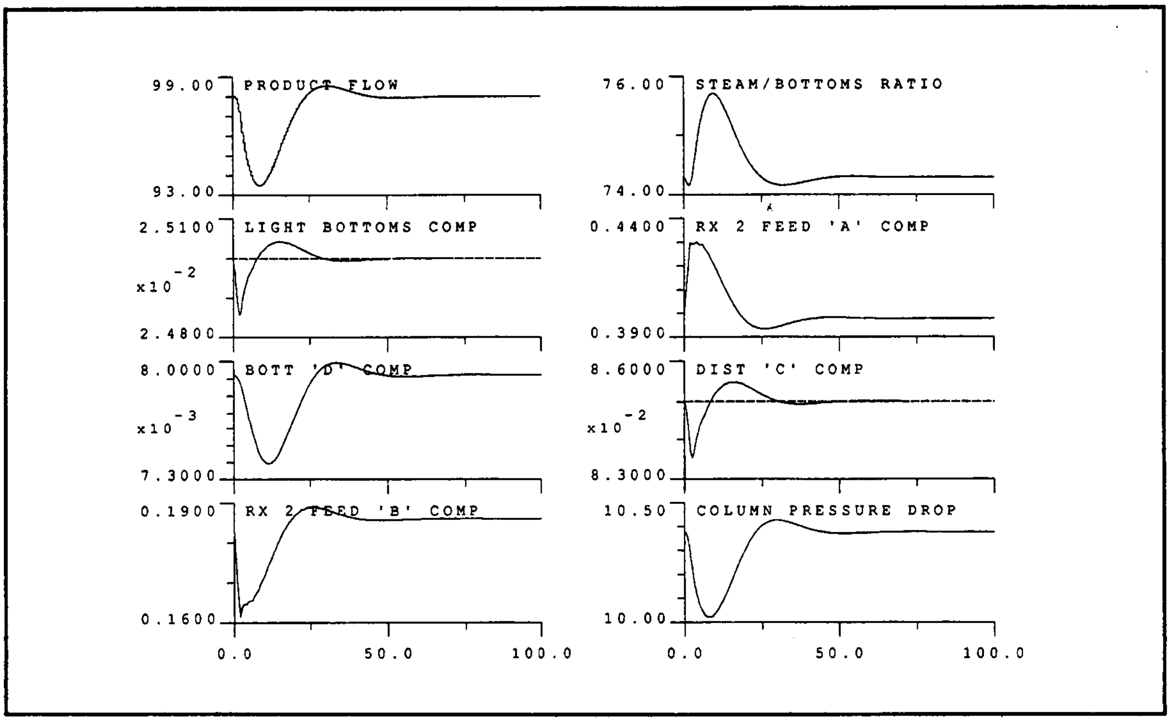

Figure B.12: Strategy 12; 10% A in B Feed

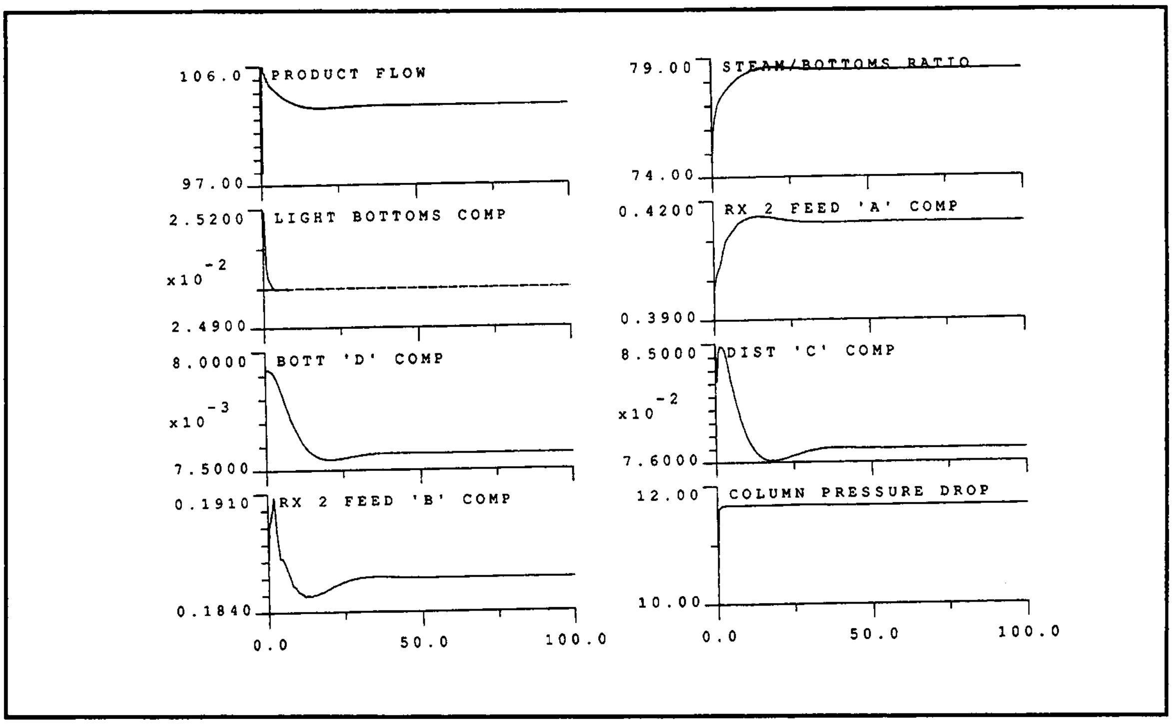

Figure B.13: Strategy 1; 5% Throughput Increase

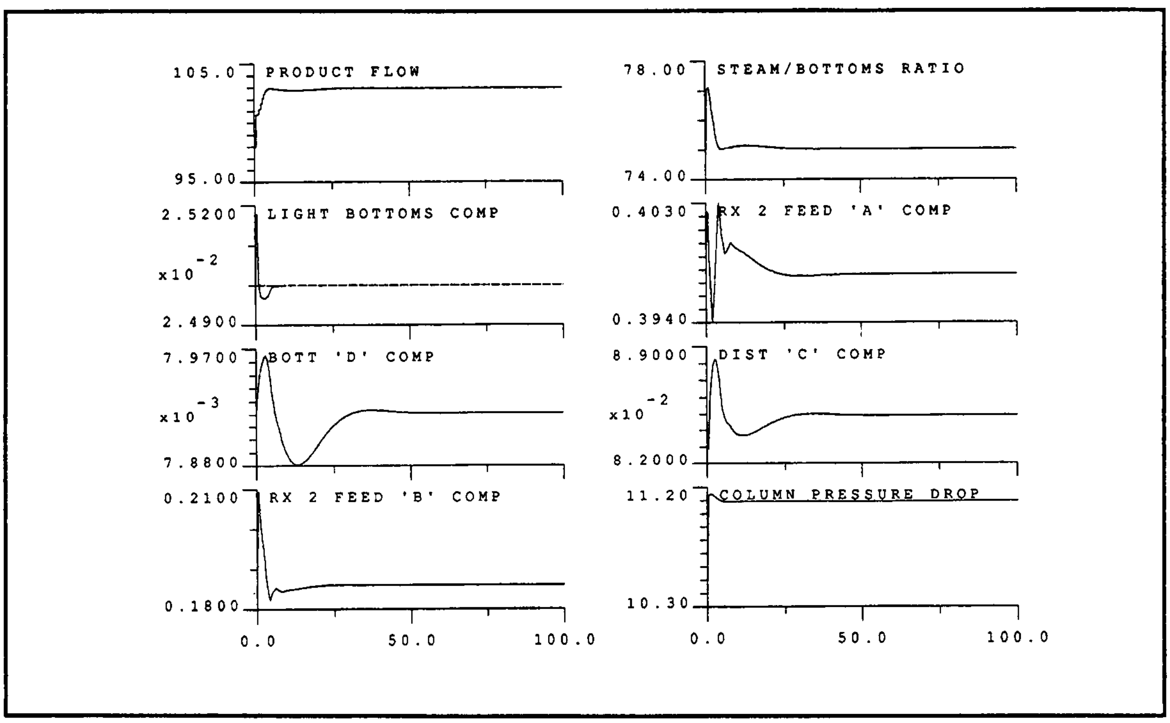

Figure B.14: Strategy 1b; 5% Throughput Increase

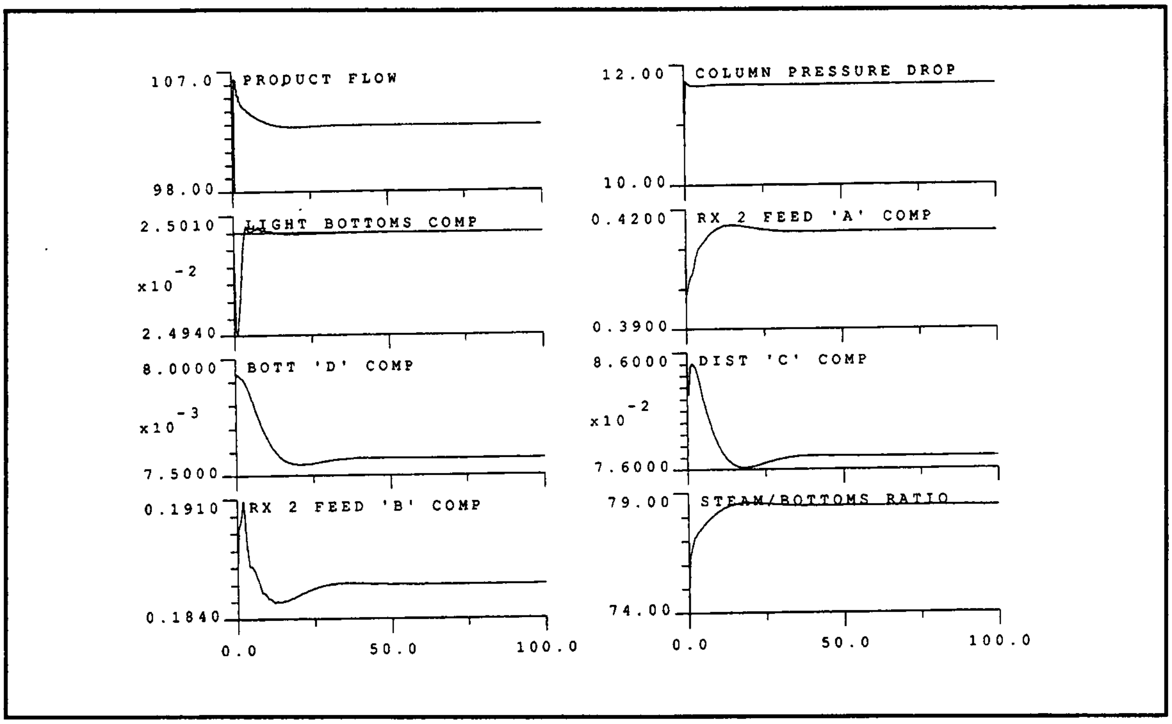

Figure B.15: Strategy 2; 5% Throughput Increase

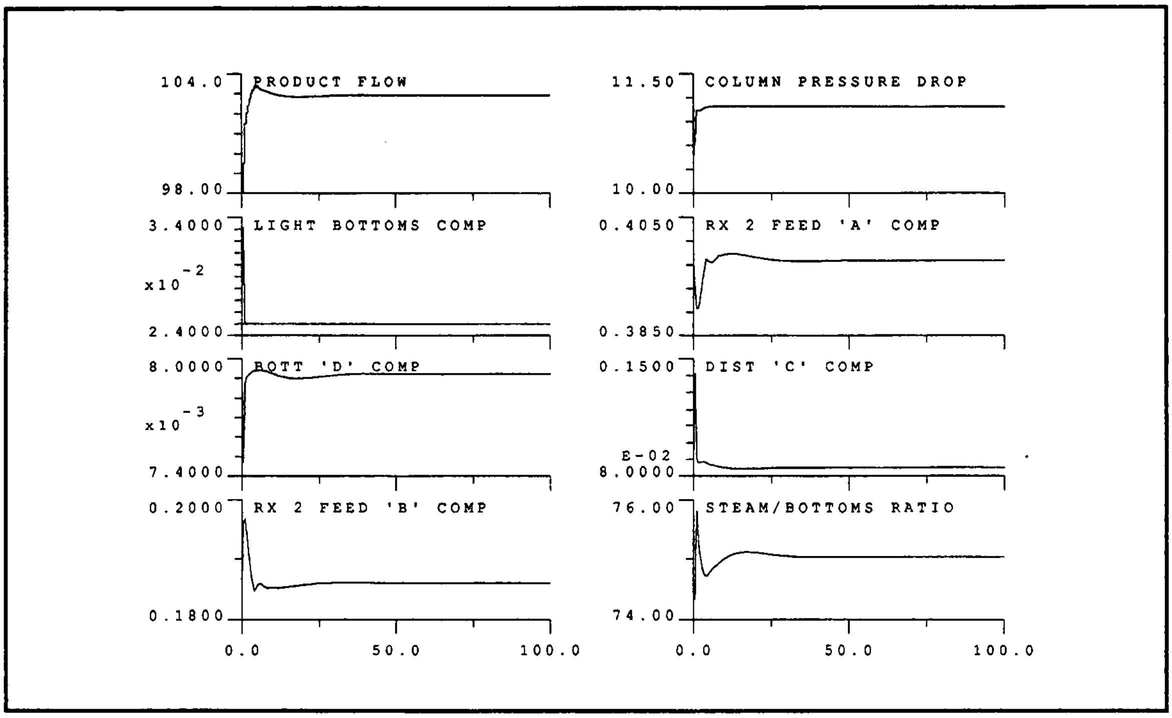

Figure B.16: Strategy 2b; 5% Throughput Increase

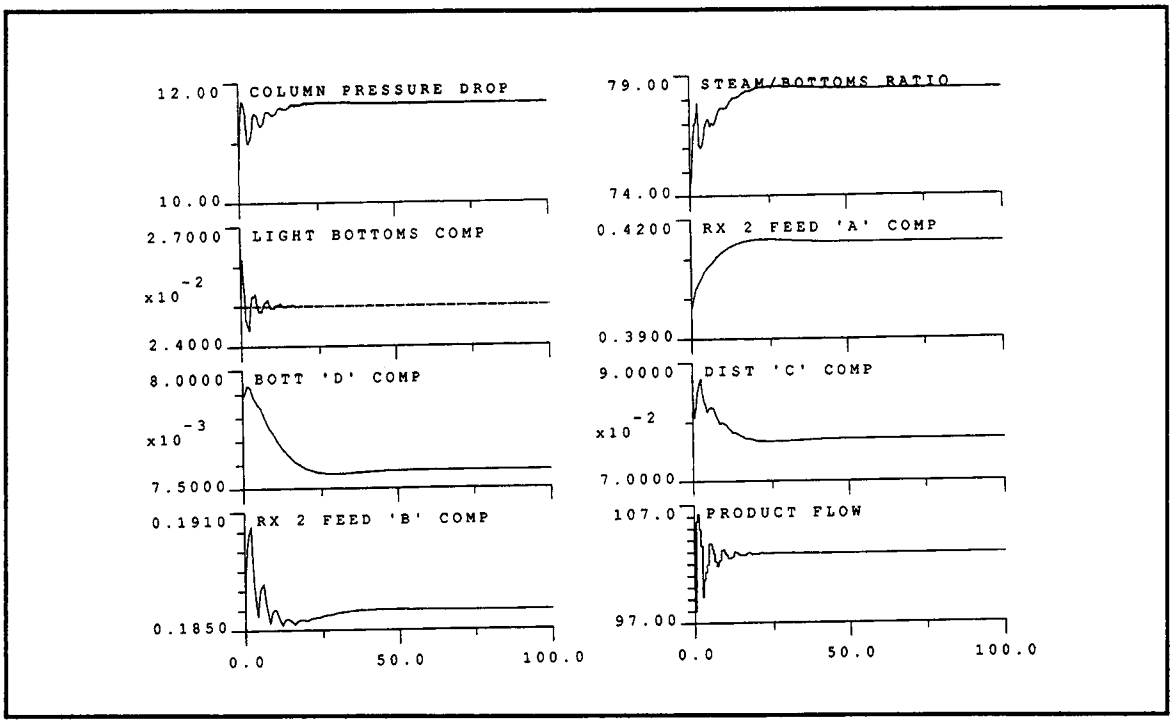

Figure B.17: Strategy 3; 5% Throughput Increase

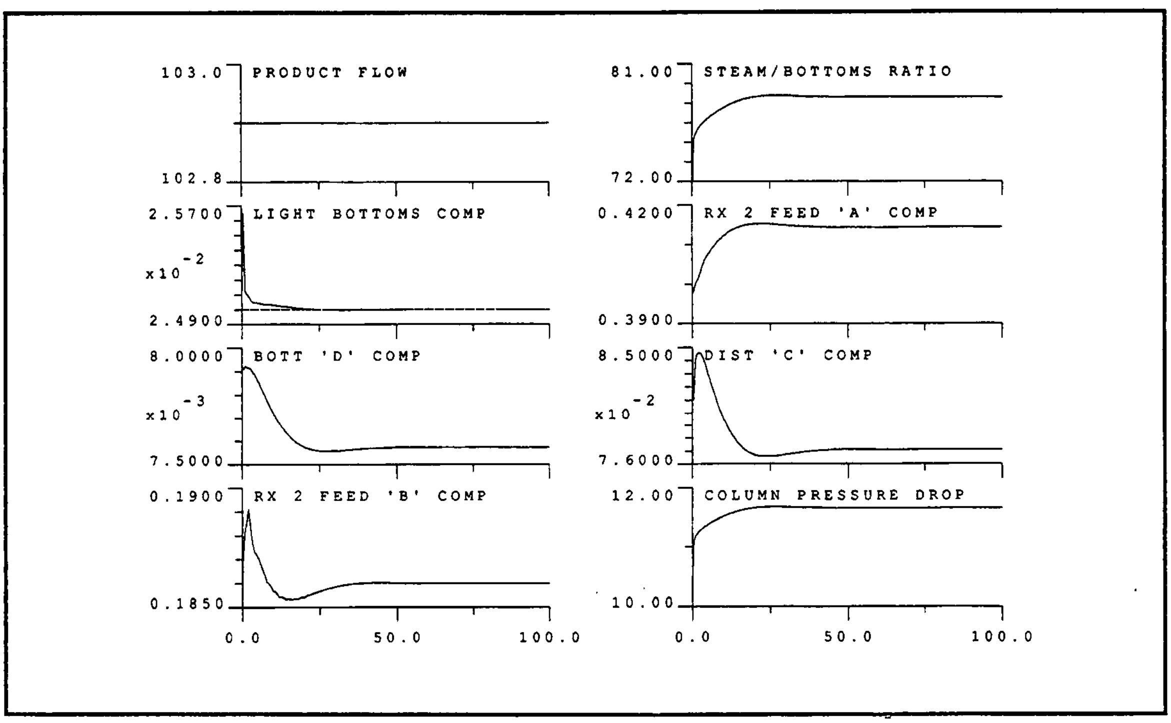

Figure B.18: Strategy 4; 5% Throughput Increase

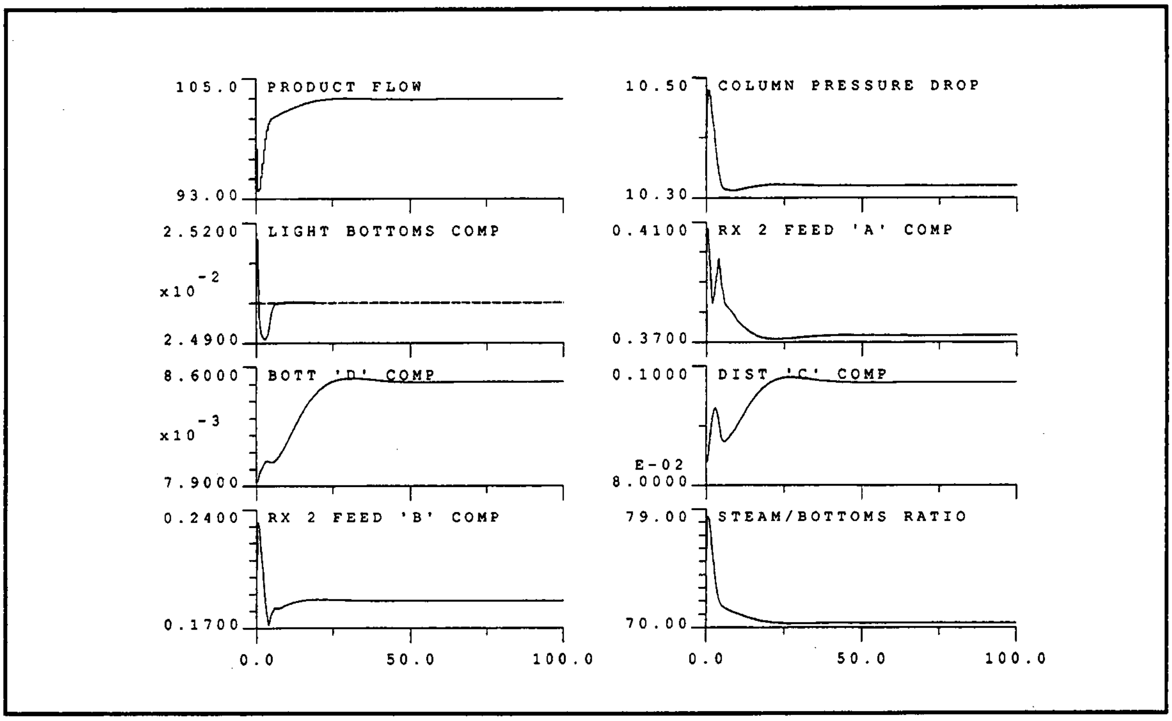

Figure B.19: Strategy 5; 5% Throughput Increase

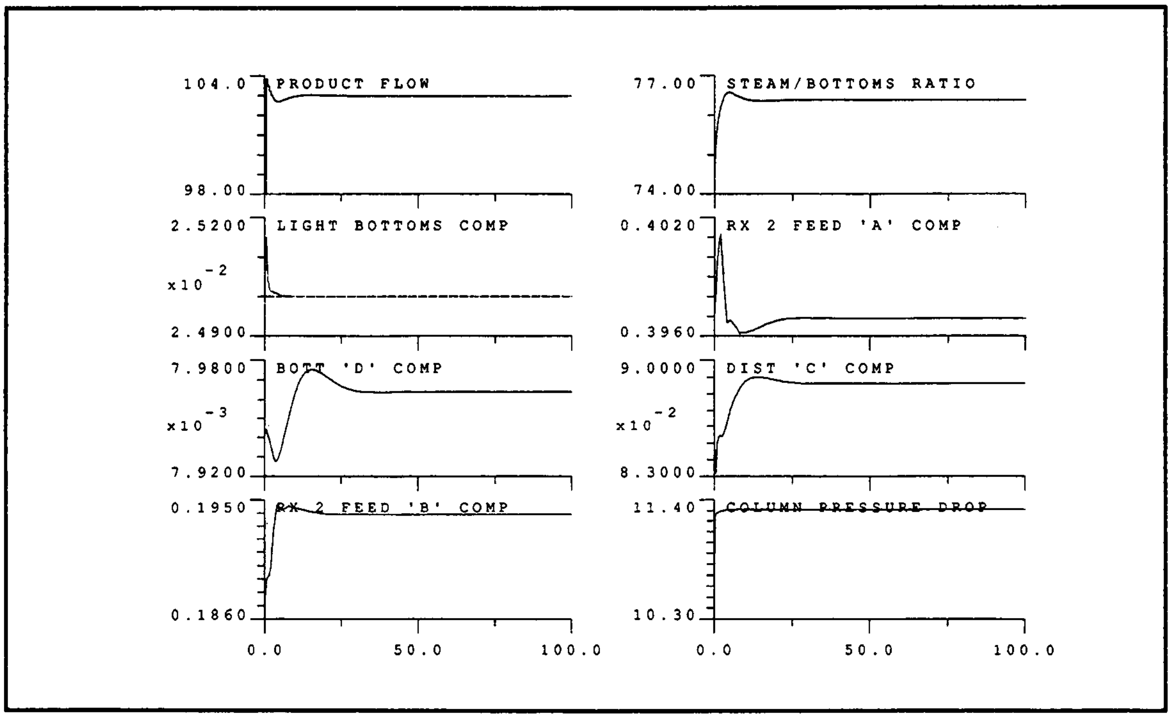

Figure B.20: Strategy 6; 5% Throughput Increase

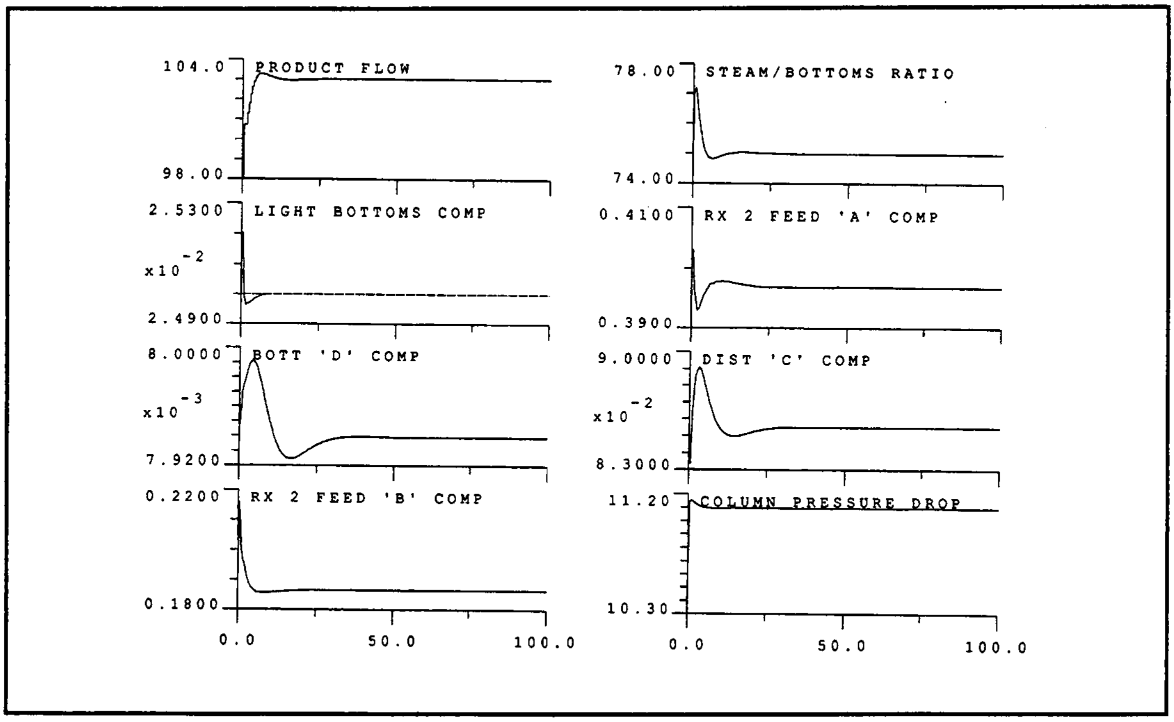

Figure B.21: Strategy 6b; 5% Throughput Increase

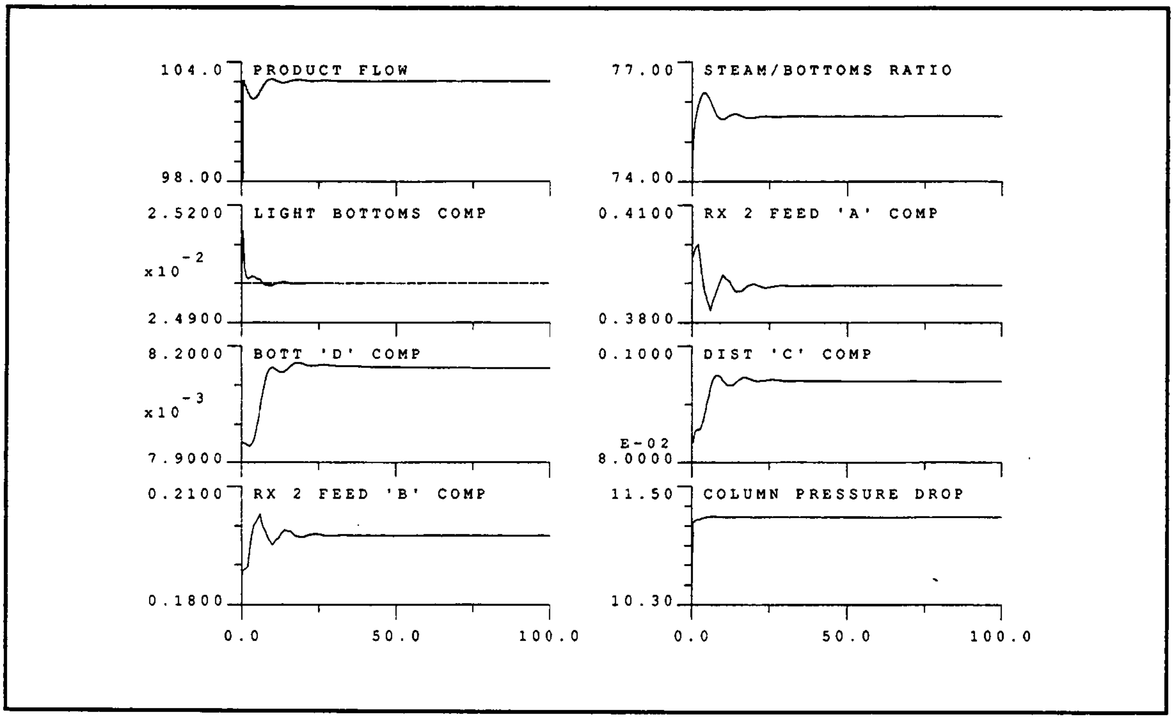

Figure B.22: Strategy 7; 5% Throughput Increase

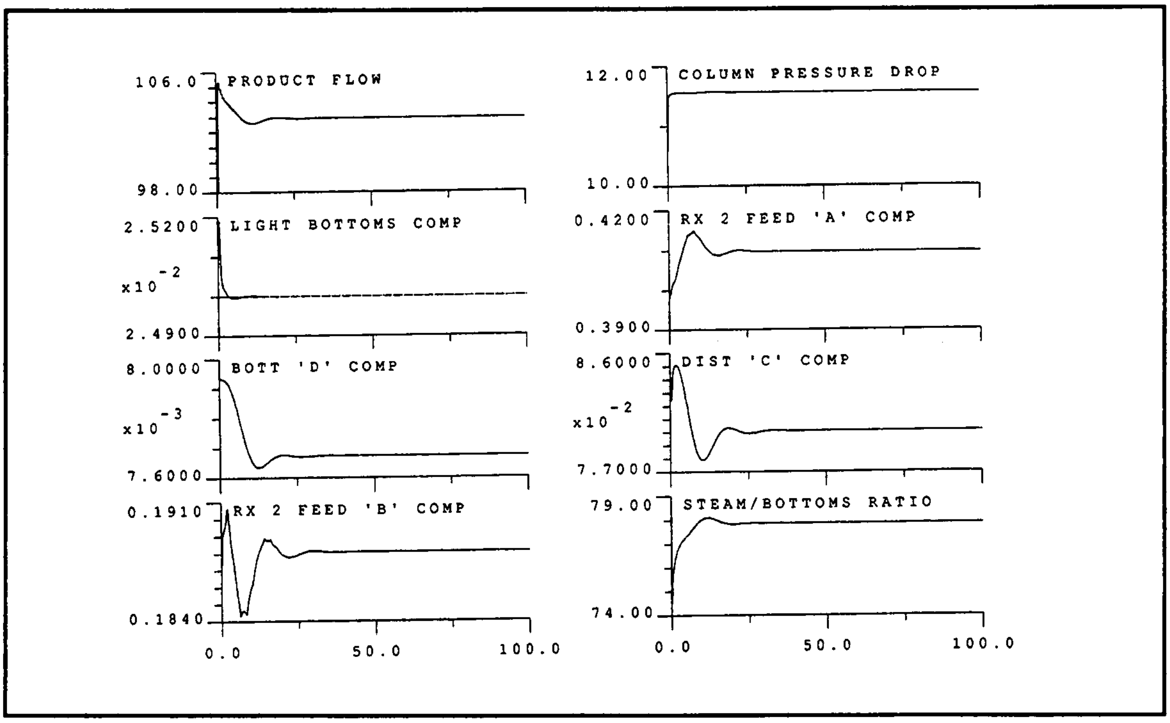

Figure B.23: Strategy 8; 5% Throughput Increase

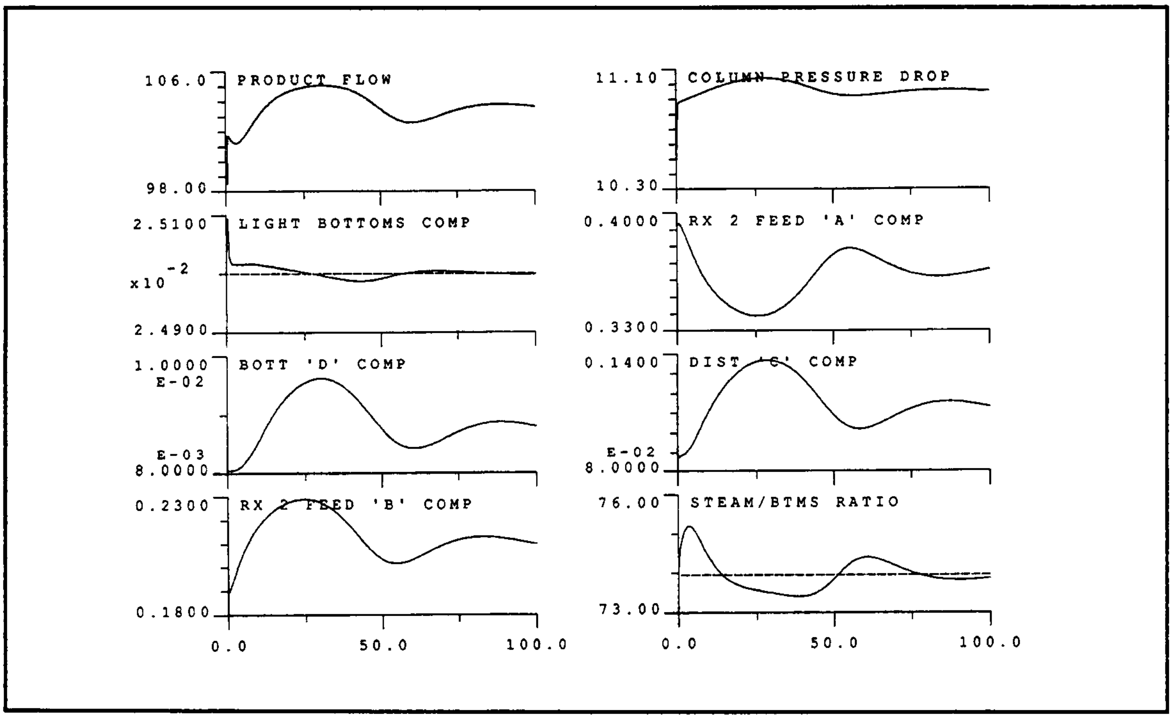

Figure B.24: Strategy 9; 5% Throughput Increase

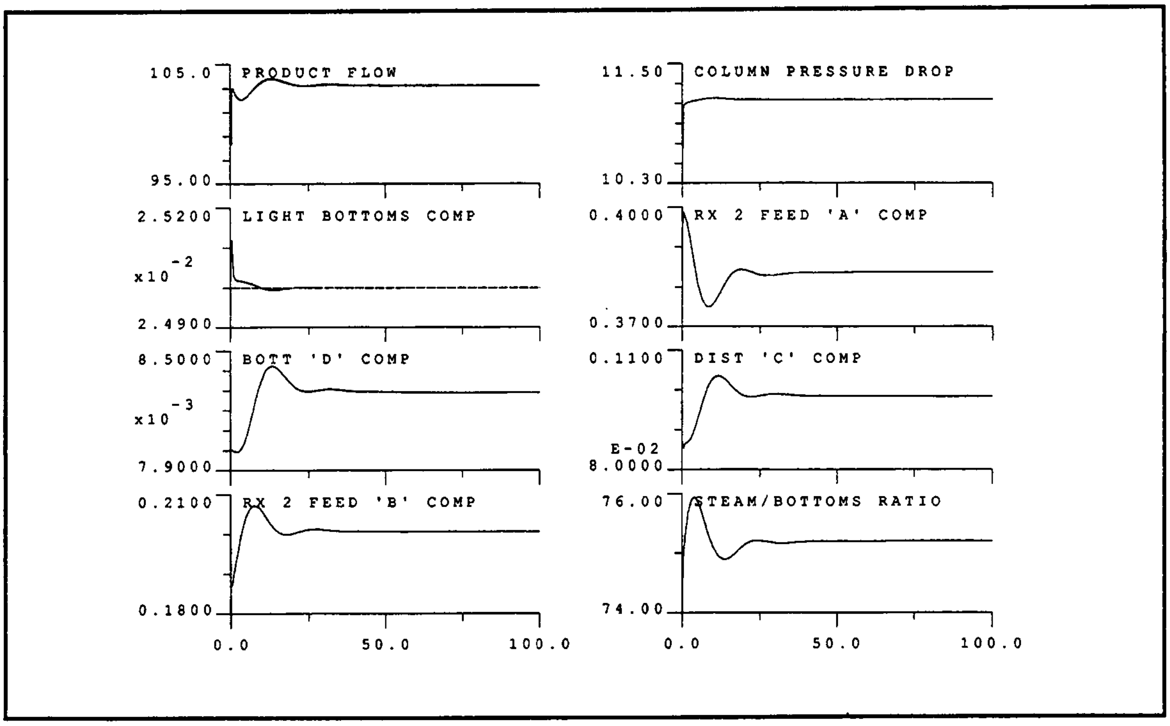

Figure B.25: Strategy 10; 5% Throughput Increase

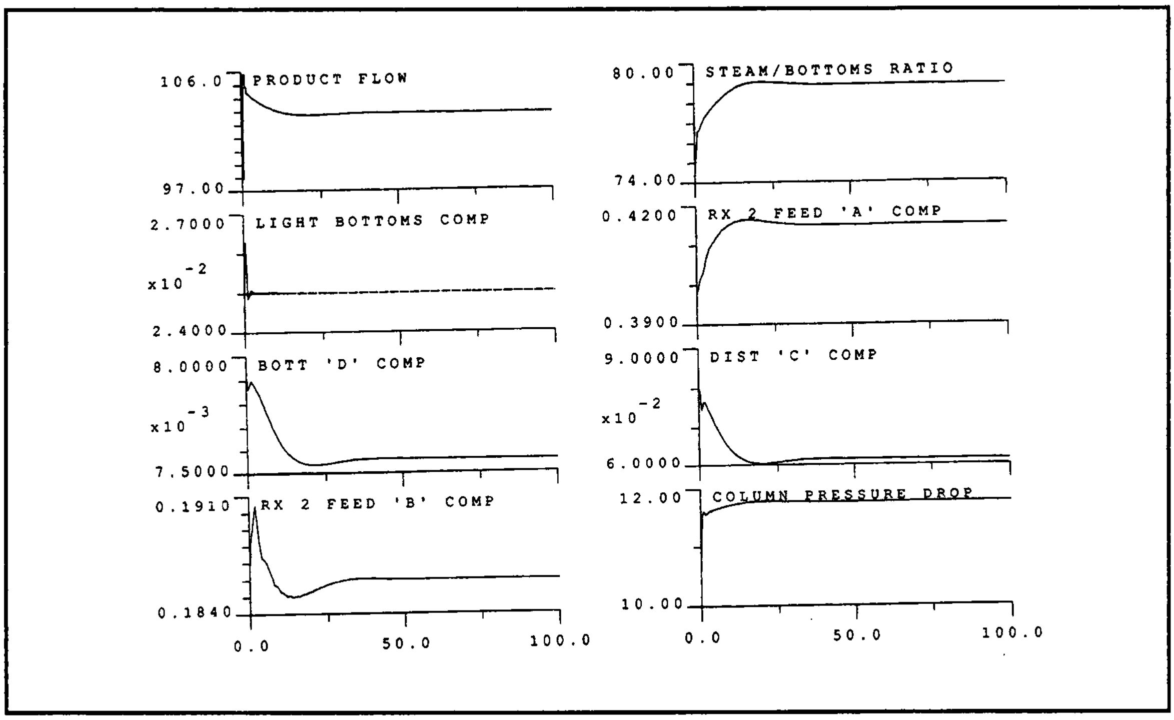

Figure B.26: Strategy 11; 5% Throughput Increase

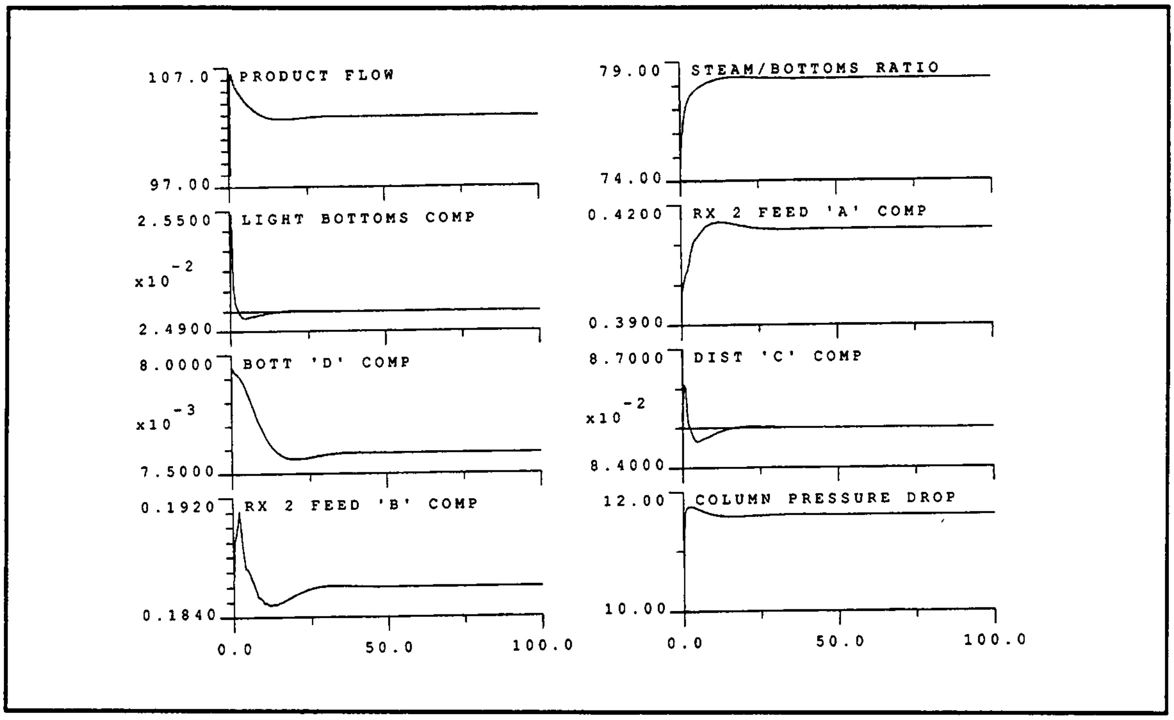

Figure B.27: Strategy 12; 5% Throughput Increase

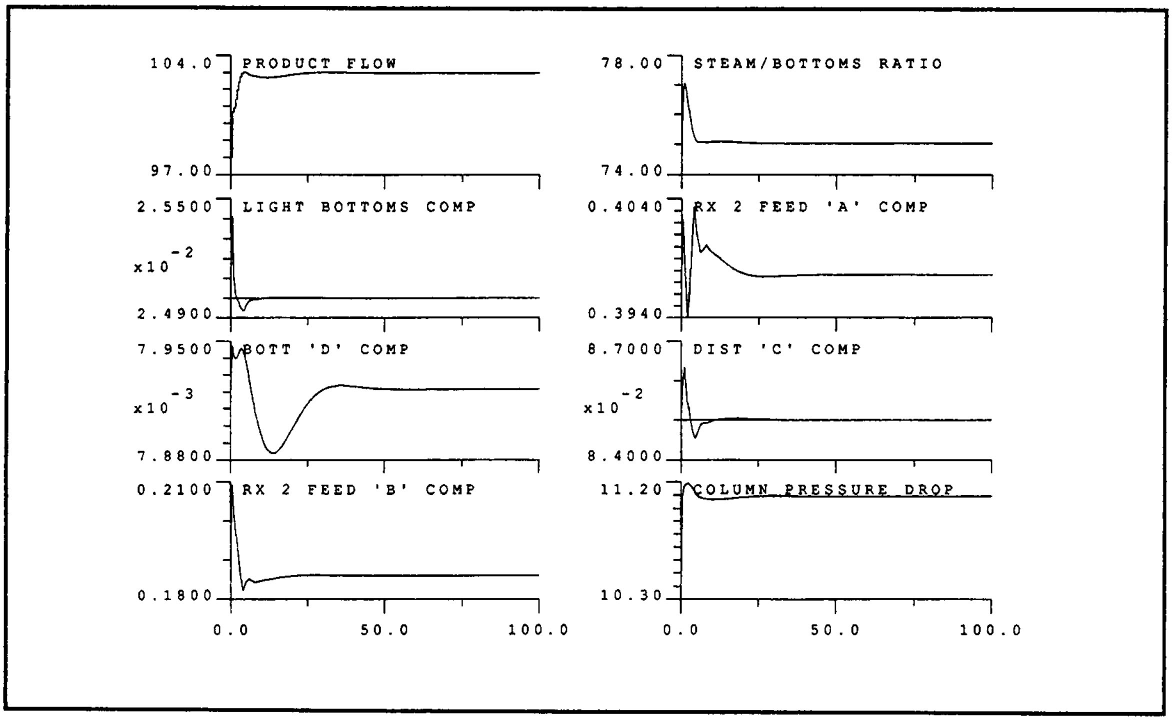

Figure B.28: Strategy 12b; 5% Throughput Increase

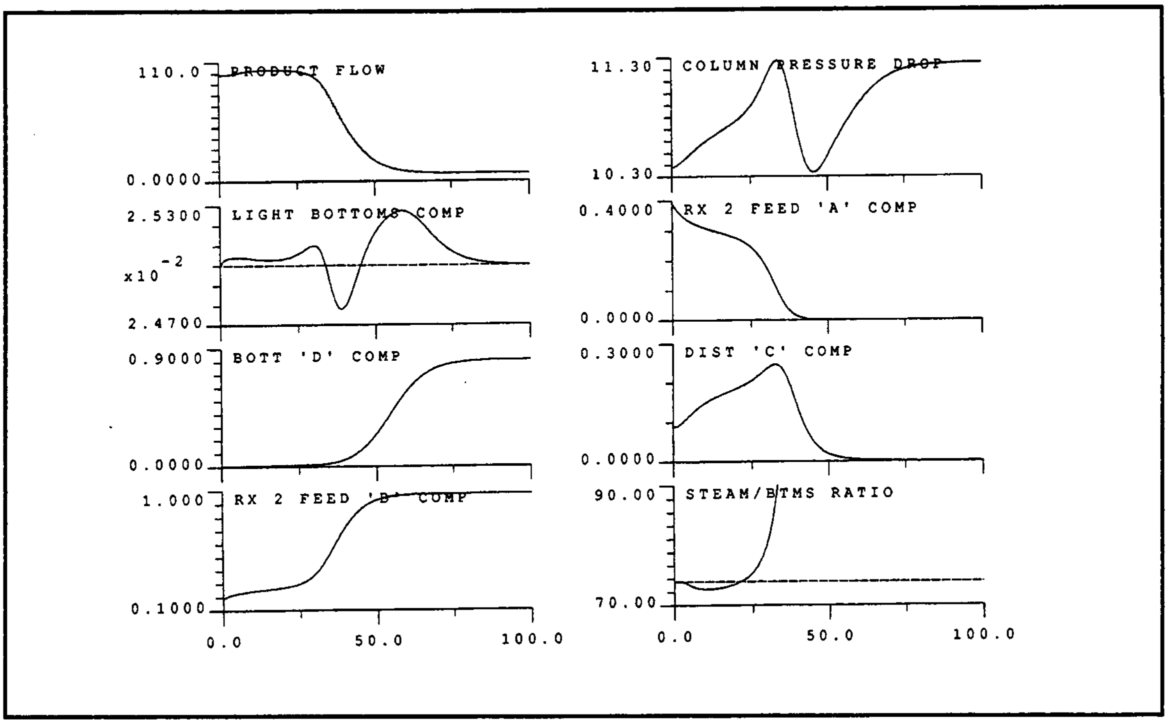

Figure B.29: Strategy 9; 9% B in A Feed

Appendix C: Variable Selection

Chapter 6 introduces a systematic procedure for variable set selection. That procedure uses the diagonals of the matrices $\mathbf{KK}^+$ and $\mathbf{K}^+\mathbf{K}$ to indicate which measurements and inputs are nonessential. The purpose of this appendix is to give some justification for that decision.

C.1 Identifying Nonessential Variables

Assume some linear system described by

(1) $$\mathbf{Kx} = \mathbf{c}$$

where

- $\mathbf{K}$ = steady state gain matrix, scaled by physical valve and instrument spans, with dimension n x m and rank r, with $r \leq \min(m,n)$;
- n = number of measurements;
- m = number of inputs (usually valves);
- $\mathbf{x}$ = vector of input values;
- $\mathbf{c}$ = vector of measurement values.

As stated in Chapter 6, the purpose of the control system is to find the appropriate $\mathbf{x}$ such that $\mathbf{c}$ approaches the setpoint vector. If $\mathbf{K}$ is square and invertible (requiring that $m = n = r$), then $\mathbf{x}$ must be equal to $\mathbf{K}^{-1}\mathbf{c}$ at steady state.

Difficulty arises when $\mathbf{K}$ is not square, or if it is square but not invertible. In that case, the optimum least-squares solution to $\mathbf{Kx} = \mathbf{c}$ is $\mathbf{x}^* = \mathbf{K}^+\mathbf{c}$, where $\mathbf{K}^+$ denotes the pseudoinverse of $\mathbf{K}$, and $\mathbf{x}^*$ is the "shortest" set of inputs that will cause the process to move as close as possible (in a least-squares sense) to the vector of setpoints. Substituting this least-squares solution back into equation (1) gives

(2) $$\mathbf{KK^+c} = \mathbf{c^*}$$

The interpretation of equation (2) is that the matrix $\mathbf{KK^+}$ maps the setpoint vector ($\mathbf{c}$) into another vector ($\mathbf{c^*}$) that is "reachable" by the manipulated variables and that is as close as possible to the given setpoints.

Before going farther, it will be useful to review briefly how the pseudoinverse of a matrix is constructed using the singular value decomposition. Even if it is not square, $\mathbf{K}$ may be decomposed by the SVD into three matrices $\mathbf{U}$, Σ, and $\mathbf{V}$, such that $\mathbf{K} = \mathbf{U}\Sigma\mathbf{V}^T$. $\mathbf{U}$ is an *nxn* orthonormal matrix, $\mathbf{V}$ is an *mxm* orthonormal matrix, and Σ is an *nxm* diagonal matrix. If $\mathbf{K}$ were square and invertible, we could construct its inverse by calculating $\mathbf{V}\,\Sigma^{-1}\mathbf{U}^T$, where Σ^{-1} is simply the diagonal matrix whose entries are the reciprocals of the entries of Σ. When $\mathbf{K}$ is not invertible, for whatever reason, Σ may have several different forms. Figure C.1 shows some possible forms for Σ. In any of these cases, the procedure for forming $\mathbf{K}^+$ requires first forming Σ^+ in this manner:

1. Transpose Σ.
2. Determine the rank *r*, the number of degrees of freedom to be represented by the reduced system. (In many cases, *r* is simply the number of non-zero singular values. However, if there is noise in the system or if some inputs have very little effect on the system, there may be some very small singular values that should not count toward the rank of the system. The effective rank is the number of sufficiently large singular values.)
3. Replace the first *r* non-zero singular values σ_i with $1/\sigma_i$.
4. Replace any remaining singular values with 0.

$$\begin{bmatrix} \sigma_1 & 0 & \cdots & 0 & 0 & 0 \\ 0 & \sigma_2 & \cdots & 0 & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & \cdots & \sigma_n & 0 & 0 \end{bmatrix}$$

Case 1: Σ has more columns than rows; rank = n.

$$\begin{bmatrix} \sigma_1 & 0 & \cdots & 0 \\ 0 & \sigma_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \sigma_m \\ 0 & 0 & \cdots & 0 \\ 0 & 0 & \cdots & 0 \end{bmatrix}$$

Case 2: Σ has more rows than columns; rank = m.

$$\begin{bmatrix} \sigma_1 & 0 & \cdots & 0 & 0 & \cdots & 0 \\ 0 & \sigma_2 & \cdots & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \cdots & \sigma_r & 0 & \cdots & 0 \\ 0 & 0 & \cdots & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & 0 & \cdots & 0 \end{bmatrix}$$

Case 3: Rank < $\min(m,n)$

Figure C.1: Singular Value Matrices for Non-Invertible **K**.

Once Σ^+ is formed, $\mathbf{K}^+$ is simply $\mathbf{V}\,\Sigma^+\mathbf{U}^T$. Note that

$$(3) \qquad \Sigma\Sigma^+ = \Sigma^+\Sigma = \begin{bmatrix} 1 & 0 & - & 0 & 0 & - & 0 \\ 0 & 1 & - & 0 & 0 & - & 0 \\ | & | & \backslash & | & | & - & | \\ 0 & 0 & - & 1 & 0 & - & 0 \\ 0 & 0 & - & 0 & 0 & - & 0 \\ | & | & - & | & | & \backslash & | \\ 0 & 0 & - & 0 & 0 & - & 0 \end{bmatrix}$$

Thus $\Sigma\Sigma^+$ and $\Sigma^+\Sigma$ are both "pseudo-identity" matrices.

The matrix $\mathbf{KK}^+$ (hereafter called κ) has several useful properties stated here without proof:

1. The diagonal elements are positive numbers between 0 and 1.
2. The trace is equal to r, the number of ranks of $\mathbf{K}$ retained.
3. The diagonal elements are the square roots of the Euclidian vector lengths. I.e., $(\kappa_{ii})^2$ is the Euclidian length of the ith column vector of κ.

This third property implies that the size of a diagonal element of κ puts an upper limit on the size of any other element in the same column or in the same row (because κ is symmetric). If κ_{ii} is the ith diagonal element of κ, then no other single element in the same vector can be larger than $\sqrt{\kappa_{ii} - \kappa_{ii}^2}$. This fact is useful if κ_{ii} is large (close to 1) or if it is small (close to 0); in either case, all the other elements of that vector of κ must be close to 0. Table C.1 shows the maximum possible values for non-diagonal elements of κ.

How should we interpret these diagonal elements? Consider Equation (2) again. Suppose there is a step change in the ith setpoint, so that the new setpoint ($\mathbf{c}$) vector is

Table C.1: Diagonal Elements in κ Matrix

Diagonal Element	Largest Possible Non-Diagonal Element in Same Vector
0.00	0.00
0.01	0.10
0.05	0.22
0.10	0.30
0.20	0.40
0.30	0.46
0.50	0.50
0.70	0.46
0.80	0.40
0.90	0.30
0.95	0.22
0.99	0.10
1.00	0.00

$[0\ 0 \ldots \mathbf{c}_i \ldots 0]^T$. (Assuming that the system is defined using deviation variables, a setpoint of 0 corresponds to the original process operating point.) If the system could be perfectly controlled, it would eventually reach a new steady state for which $\mathbf{c}^* = [0\ 0 \ldots \mathbf{c}_i \ldots 0]^T = \mathbf{c}_i[0\ 0 \ldots 1 \ldots 0]$ also. However, in the general case, this point cannot be attained. The best we can do is apply some "optimal least-squares controller" that will get as close as possible to the new setpoint vector with the least possible movement in the inputs. This controller can only reach the operating point $\mathbf{c}^* = \mathbf{c}_i[\kappa_{1i}\ \ \kappa_{2i} \ldots \kappa_{ii} \ldots \kappa_{ni}]$. That is, the new steady-state is simply a reflection of the *i*th column of κ. How well the new steady state mimics the vector of setpoints depends on the size of κ_{ii} thus:

- If κ_{ii} is close to 1, then all other elements κ_{ji} must be close to 0. Then $\mathbf{c}^*$ is close to $\mathbf{c}_i[0\ 0 \ldots 1 \ldots 0]$, and the new steady state is close to the setpoint vector. Thus, an optimum least-squares controller can control the *i*th measurement without greatly disturbing the other measurements.
- If κ_{ii} is close to 0, then all other elements κ_{ji} must be close to 0. Then $\mathbf{c}^*$ is close to $\mathbf{c}_i[0\ 0 \ldots 0 \ldots 0]$; that is, it is very close to the original operating point. Thus, the optimum least-squares controller will not control the *i*th measurement with the given inputs. In the terminology of Chapter 6, this is a nonessential measurement.
- If κ_{ii} is somewhere between 0 and 1, but not close to either one, then some other elements κ_{ji} also have significant values, and the $\mathbf{c}^*$ vector will have more than one significant non-zero entry at the new steady state. These measurements are coupled--the controller cannot control one of them independent of the others, but it could control them if their setpoints were moved together. In the terminology of Chapter 6, these are "equivalent" measurements.

Specifically, it is important to note that small κ_{ii} values indicate measurements that the optimal least-squares controller cannot match to their setpoints without massive changes in other setpoints. Such measurements are poor choices for controlled variables in <u>any</u> control scheme, because there is no input or combination of inputs that significantly affects those measurements. Thus, even if we are using a m-SISO controller, any process measurement that corresponds to a small value of κ_{ii} can be confidently eliminated from the set of potential controlled variables. Similarly, process measurements that correspond to large values of κ_{ii} are good choices for controlled variables. When several measurements each have "medium-sized" values of κ_{ii}, they are coupled at steady state, and further analysis must be done to determine which variables are linked with which others. Two different methods for such analysis will be presented later in this appendix.

A similar reasoning process applies to the matrix $\mathbf{K}^{+}\mathbf{K}$ (instead of $\mathbf{K}\mathbf{K}^{+}$) to analyze the set of inputs. Similar conclusions follow (with κ'_{ii} denoting the diagonal elements of $\mathbf{K}^{+}\mathbf{K}$):

- If κ'_{ii} is close to 1, then the optimum least-squares controller uses the *i*th input to reach the new steady state following a setpoint change, and that input should be included in the manipulated variable set.
- If κ'_{ii} is close to 0, then the optimum least-squares controller does not require significant changes in the *i*th input for any setpoint change, and that input should not be included in the manipulated variable set. In the terminology of Chapter 6, this is a nonessential input.
- If κ'_{ii} is between 0 and 1, but not close to either one, then the optimum least-squares controller uses the *i*th input in tandem with certain other inputs. That set

of coupled inputs should be used as a single manipulated variable. In the terminology of Chapter 6, these are equivalent inputs.

Example C-1: Identifying Nonessential Variables

Consider the gain matrix shown in Figure C.2. The diagonals of κ are 0.085, 0.991, and 0.925. The second and third measurements correspond to diagonal elements close to 1, so these measurements are essential. The first measurement corresponds to a diagonal element close to 0, so this measurements is nonessential. Similarly, the diagonals of κ' are 0.9520, 0.9345, and 0.1135. The first and second inputs are essential, while the third is nonessential. (Note: the singular value decomposition of the matrix shown in Figure C.2 reveals that the third degree of freedom in this system is very weak; the singular vectors of this weak degree of freedom are dominated by the first measurement and the third input, in agreement with our conclusions here.)

$$\text{Original Gain Matrix:} \quad \begin{bmatrix} 0.42 & 0.90 & 0.18 \\ 0.84 & 3.49 & 0.75 \\ -1.14 & -1.93 & -0.25 \end{bmatrix}$$

Figure C.2: Identifying Nonessential Variables

C.2 Identifying Equivalent Variables

When the diagonal elements of κ or κ' are not close to either 0 or 1, then some degree of freedom is "shared" among two or more variables. When more than one degree of freedom is thus distributed among variables, the κ and κ' matrices do not directly indicate which variables should be grouped into a set. There are at least two ways to determine this information. An example will follow the descriptions of these two methods.

C.2.1 Grouping Variables Using Vector Angles

The vectors that make up the process gain matrix **K** define a series of "directions" for the process. Finding equivalent inputs or measurements is the same thing as finding vectors of **K** that point in the same direction. It is a well-known result from linear algebra that the angle θ between two vectors a and b can be found by

$$\cos\theta = \frac{a^T b}{\|a\|_2 \|b\|_2}. \tag{4}$$

If two vectors lie in the same direction, the cosine of the angle between them is close to 1 or -1. Thus one way to find equivalent variables is to evaluate the angles between them.

C.2.2 Grouping Variables Using Inflated Gains

Another way to identify a set of equivalent variables is by artificially "inflating" the gains for one variable and recalculating the diagonals of the κ or κ' matrices. The other variables whose gain vectors point in the same direction will then appear to be nonessential, because their sensitivities will be relatively so much smaller compared to the inflated variable.

Example C-2: Identifying Equivalent Variables

Consider the gain matrix shown in Figure C.3. From the singular values, it should be clear that there are effectively only two degrees of freedom in this process. Because this is a simple example, it is easy to see from visual inspection of the gain matrix that the first and third column vectors are aligned, and that the second and fourth column vectors are aligned. How should manipulated variables be chosen from among the inputs for this process?

Form the pseudoinverse of **K**, also shown in Figure C.3. Then look at the diagonals of the product $\mathbf{K}^{+}\mathbf{K}$. They are 0.5126, 0.5264, 0.4873, and 0.4737. It is obvious that the two degrees of freedom are shared among the four manipulated variables. In fact, we may infer from these numbers that the variables should be grouped in sets of two, but it is not obvious which variables are equivalent.

Equivalency may be determined by checking the cosine of the angle between pairs of vectors in **K**. Figure C.3 also lists these cosines from Equation (4) above. From these cosines it is clear that vectors 1 and 3 are aligned (though in opposite directions), as are vectors 2 and 4 (in the same direction). The conclusion is that inputs 1 and 3 are equivalent, and inputs 2 and 4 are equivalent. The non-equivalent vectors are about 10-15 degrees away from being orthogonal, so the sets of inputs are not quite decoupled at steady state.

Equivalency may also be determined by inflating one of the columns (or rows) of **K** and recalculating κ' (or κ). This procedure is somewhat tedious manually, but it is easily programmed. For example, Figure C.3 shows the **K** matrix with the first column inflated

Original Gain Matrix (**K**): $\begin{bmatrix} .2752 & -.0167 & -.2588 & -.0218 \\ .0289 & .6211 & .0153 & .5854 \\ .7360 & -.2237 & -.7105 & -.1841 \\ .5788 & .5819 & -.5689 & .5768 \end{bmatrix}$

Singular Value Matrix (Σ): $\begin{bmatrix} 1.4411 & 0 & 0 & 0 \\ 0 & 1.1207 & 0 & 0 \\ 0 & 0 & 0.0276 & 0 \\ 0 & 0 & 0 & 0.0102 \end{bmatrix}$

Rank-2 Pseudoinverse of Σ: $\begin{bmatrix} 0.6939 & 0 & 0 & 0 \\ 0 & 0.8923 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$

Rank-2 Pseudoinverse of **K**: $\begin{bmatrix} .1558 & -.0758 & .4422 & .2529 \\ -.0553 & .4394 & -.2610 & .3337 \\ -.1534 & .0866 & -.4386 & -.2368 \\ -.0469 & .4139 & -.2318 & .3253 \end{bmatrix}$

Diagonals of $\mathbf{K}^+\mathbf{K}$: $[0.5126 \ 0.5264 \ 0.4873 \ 0.4737]$

Cosine of angle between vector pair

1-2:	0.216
1-3:	-0.999
1-4:	0.255
2-3:	-0.190
2-4:	0.999
3-4:	-0.229

K with first column inflated: $\begin{bmatrix} 2.752 & -.0167 & -.2588 & -.0218 \\ .289 & .6211 & .0153 & .5854 \\ 7.360 & -.2237 & -.7105 & -.1841 \\ 5.788 & .5819 & -.5689 & .5768 \end{bmatrix}$

Diagonals of new κ': $[0.9906 \ 0.5264 \ 0.0098 \ 0.4732]$

Figure C.3: Identifying Equivalent Variables by a factor of 10. Forming κ' using this new **K** gives diagonals of 0.9906, 0.5264, 0.0098, and 0.4732. Now the first input appears to account for nearly a whole degree of freedom, while the third one appears nonessential. Exaggerating the effect of the first input has diminished the comparative effect of the third input, but not the second or fourth inputs. We conclude that the first and input is equivalent to the third, but not to the second or fourth. Similar inflation of the other columns leads us to conclude that inputs 1 and 3 are equivalent, as are inputs 2 and 4.

(Of course, this method works only because the κ evaluation method is scale-sensitive--otherwise, the diagonals would not change when we introduced an artificial scaling factor. Although scale sensitivity is widely viewed as a drawback to an evaluation method, it may sometimes be used to advantage!)

In examples as simple as those cited here, examining the matrices of the singular value decomposition could have yielded the same information about nonessential variables, and perhaps even about variable equivalence. A systematic approach, though, is useful for more complicated systems.

C.3 An Analogy to the RGA

The Relative Gain Array introduced by Bristol (6) is among the most widely used tools for analyzing steady-state gain matrices. One of its limitations is that it requires *a priori* selection of the controlled and manipulated variables, because it can be calculated only for invertible gain matrices. The existence of the pseudoinverse suggests that an analogous tool might be available for the general case, in which **K** is not invertible.

The RGA may be simply calculated by $\mathbf{\Lambda} = \mathbf{K} \otimes (\mathbf{K}^{-1})^T$, where $\otimes$ denotes element-by-element multiplication. Each element of the RGA is the ratio of the open loop gain for an input/measurement pair to the gain with all other loops closed. In the case for which $\mathbf{K}^{-1}$ does not exist, though, the RGA is not defined. It is possible, though, to define an analogous measure, a "reduced-rank relative gain array" (RRGA) which uses the pseudoinverse of $\mathbf{K}$:

(5) $$\mathbf{\Lambda}^{*} = \mathbf{K} \otimes (\mathbf{K}^{+})^{T}$$

What is the significance of this measure? Just as the elements of $\mathbf{K}^{-1}$ are related to the "closed-loop" gains necessary for calculating the relative gains, the elements of $\mathbf{K}^{+}$ are related to the "least-squares" gains, i.e., the gains when the inputs and outputs are connected by the optimal least-squares controller described earlier. Thus the elements of Λ^{*} are ratios of open loop gains to closed-loop <u>least-squares</u> gains for processes with non-invertible gain matrices.

Bristol (6) first noted that each column and row sum in the RGA is equal to 1. It can be easily shown that the column and row sums of the reduced-rank RGA are equal to the diagonal elements of $\mathbf{K}^{+}\mathbf{K}$ and $\mathbf{K}\mathbf{K}^{+}$, respectively. (For Bristol's RGA, this fact is trivial: for a square full-rank matrix, the pseudoinverse is simply the normal inverse, and both $\mathbf{K}^{+}\mathbf{K}$ and $\mathbf{K}\mathbf{K}^{+}$ are the identity matrix, whose diagonals are all 1's. These 1's indicate that each input and each measurement are essential if every degree of freedom is to be controlled.)

<u>Example C-3: The Reduced-Rank Relative Gain Array</u>

Figure C.4 shows a gain matrix and its RGA. (Row or column sums may not be exactly 1.0 due to rounding.) Two things are readily apparent from the RGA:

- The large numbers indicate a near-singular gain matrix; the system does not have three effective degrees of freedom.

Original Gain Matrix (**K**): $\begin{bmatrix} -0.51 & -0.89 & -0.13 \\ -0.59 & -3.56 & -0.83 \\ 1.52 & 1.82 & 0.15 \end{bmatrix}$

Bristol's RGA for **K**: $\begin{bmatrix} 27.9 & -58.4 & 31.6 \\ -3.4 & 24.1 & -19.7 \\ -23.5 & 35.3 & -10.8 \end{bmatrix}$

Rank-2 RRGA for **K**: $\begin{bmatrix} 0.09 & 0.005 & -0.005 \\ -0.24 & 1.08 & 0.14 \\ 1.10 & -0.15 & -0.03 \end{bmatrix}$

RGA of new 2x2 system: $\begin{bmatrix} -0.24 & 1.24 \\ 1.24 & -0.24 \end{bmatrix}$

Figure C.4: The Reduced-Rank Relative Gain Array

- It is usually poor policy to pair variables whose RGA element is negative, but here it is impossible to create a 3x3 pairing that does not involve at least one negative RGA element.

The singular values of the gain matrix are 4.38, 1.1, and 0.004. It appears that there are two effective degrees of freedom, so we may proceed by constructing the rank-2 pseudoinverse of **K** by the procedure outlined earlier. We may then form $\mathbf{\Lambda}_2^*$, the reduced-rank RGA with 2 degrees of freedom. This matrix is also shown in Figure C.4. The column and row sums suggest that the third input and first measurement are not essential to the two primary degrees of freedom, and should be omitted from a 2x2 control system. Figure C.4 also shows the RGA of the system that remains after the first row and third column are deleted from **K**; note that this matrix is very similar to the elements that would remain in $\mathbf{\Lambda}_2^*$ if its first row and third column were deleted! This same result may be observed in other matrices (for example, the matrix in Figure C.2): the elements of the "essential" rows and columns of $\mathbf{\Lambda}^*$ are approximately the same as the elements of the RGA formed from the reduced system formed by deleting nonessential variables from the gain matrix.

Example C-4: RRGA for Triangular Gain Matrices

It is well-known in the literature [see Jensen (29) and McAvoy (54), for example] that the RGA for a triangular plant gain matrix is simply the identity matrix. Figure C.5 shows an example of a gain matrix for which the RGA recommendation may be misleading. The RGA for this matrix shows the only possible variable pairing for which the two measurements could be independently controlled, but it is obvious from the gain

Gain Matrix (**K**): $\begin{bmatrix} 1 & 100 \\ 0 & 1 \end{bmatrix}$

RGA for **K**: $\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$

Rank-1 RRGA for **K**: $\begin{bmatrix} .0001 & .9998 \\ .0000 & .0001 \end{bmatrix}$

Figure C.5: RRGA for Triangular Gain Matrices

matrix that there is really only one effective degree of freedom in the process. The control will be almost certainly be very poor unless the second input is used to control the first measurement, while the second measurement remains uncontrolled. This observation is not at all obvious from the RGA, though, which seems to indicate "noninteracting" control when pairings are assigned along the main diagonal.

A singular value decomposition indicates the potential problem: the singular values for this gain matrix are 100 and .01; the condition number is 10,000! Omitting the weak degree of freedom leads to the rank-1 RRGA shown in Figure C-5, which suggests a more rational variable pairing.

Appendix D: Eastman Challenge Plant Dynamic Responses

The four candidate control systems for the Eastman Challenge Plant were introduced in Chapter 7. Although it is difficult to find a basis for numerical comparison of these strategies, it may be useful to compare their dynamic responses to several disturbances. Downs and Vogel (15) recommend four disturbances as a basis for comparison:

- stream 4 step change in A/C ratio (disturbance 1)
- step change in reactor cooling water inlet temperature (disturbance 4)
- stream 4 random variation in A, B, and C compositions (disturbance 8)
- random change in condenser cooling water inlet temperature, combined with sticking condenser coolant valve (disturbances 12 and 15).

The reactor pressure is the variable most sensitive to many of the plant disturbances. As the pressure varies, so do the other critical plant variables, and the reactor pressure is the variable most likely to violate the equipment constraints. Therefore, the reactor pressure is a convenient indicator of the ease with which disturbances are rejected.

Figures D.1 through D.16 show the dynamic response of the reactor pressure to each of the disturbances listed above for the four candidate control strategies. Each disturbance is introduced to the plant at a time of 1 hour.

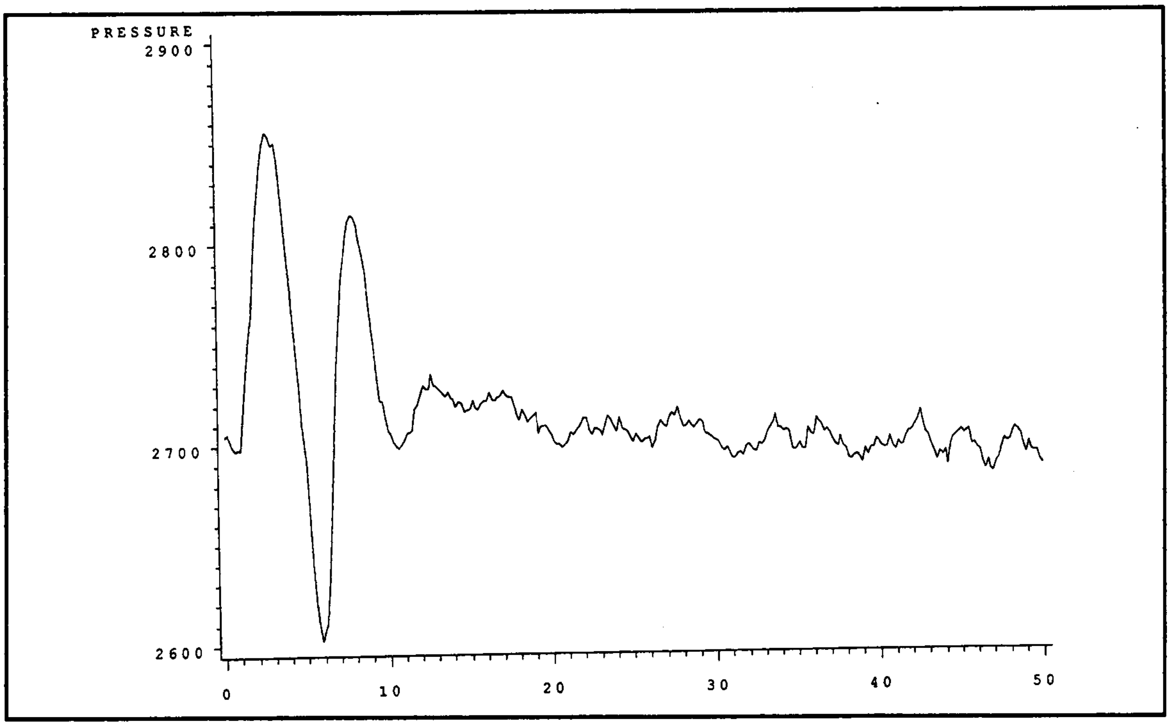

Figure D.1: Strategy 1 Response to Disturbance 1

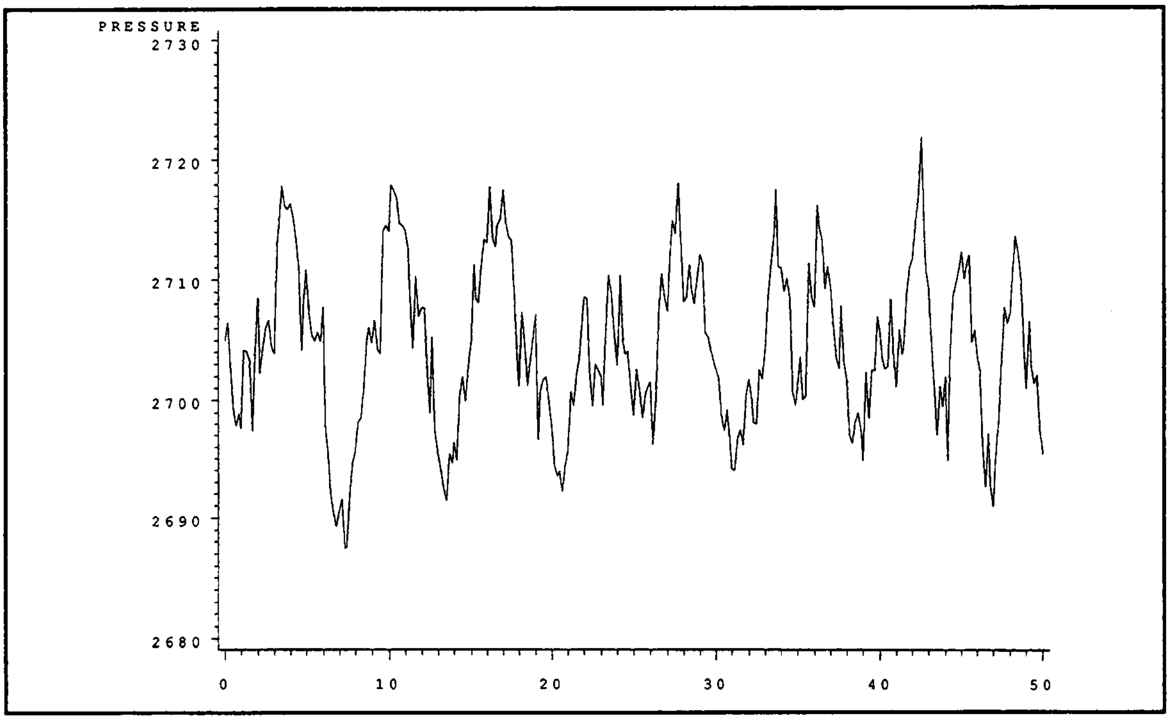

Figure D.2: Strategy 1 Response to Disturbance 4

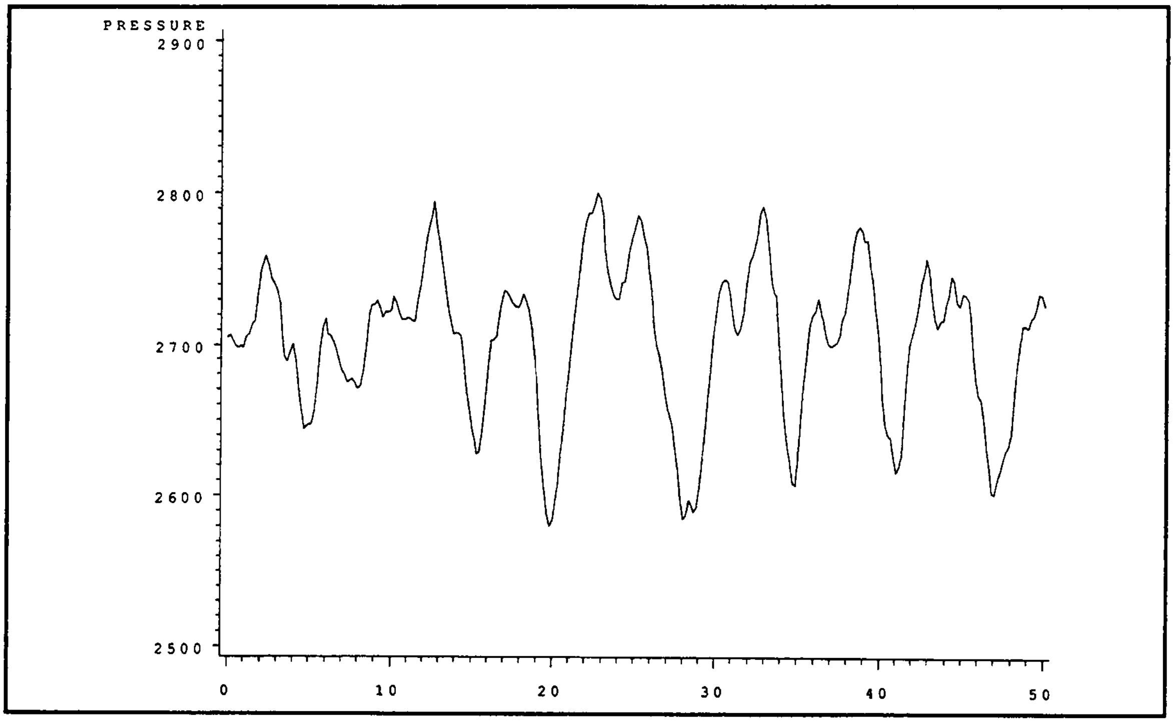

Figure D.3: Strategy 1 Response to Disturbance 8

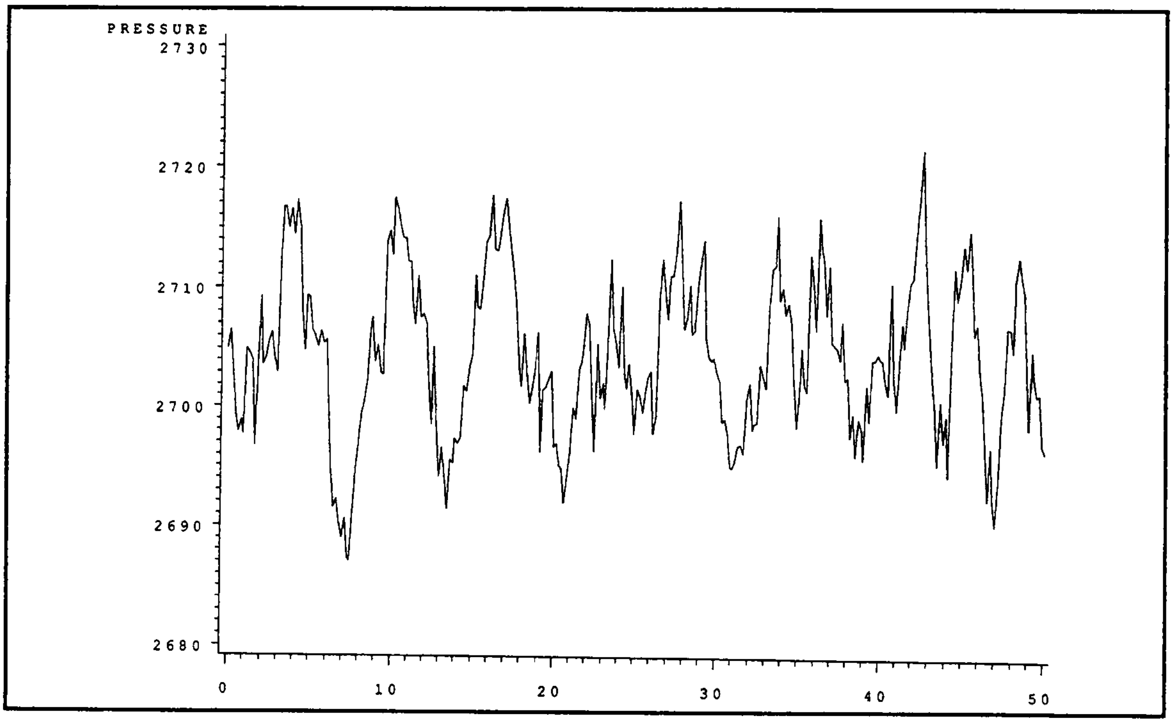

Figure D.4: Strategy 1 Response to Disturbance 12/15

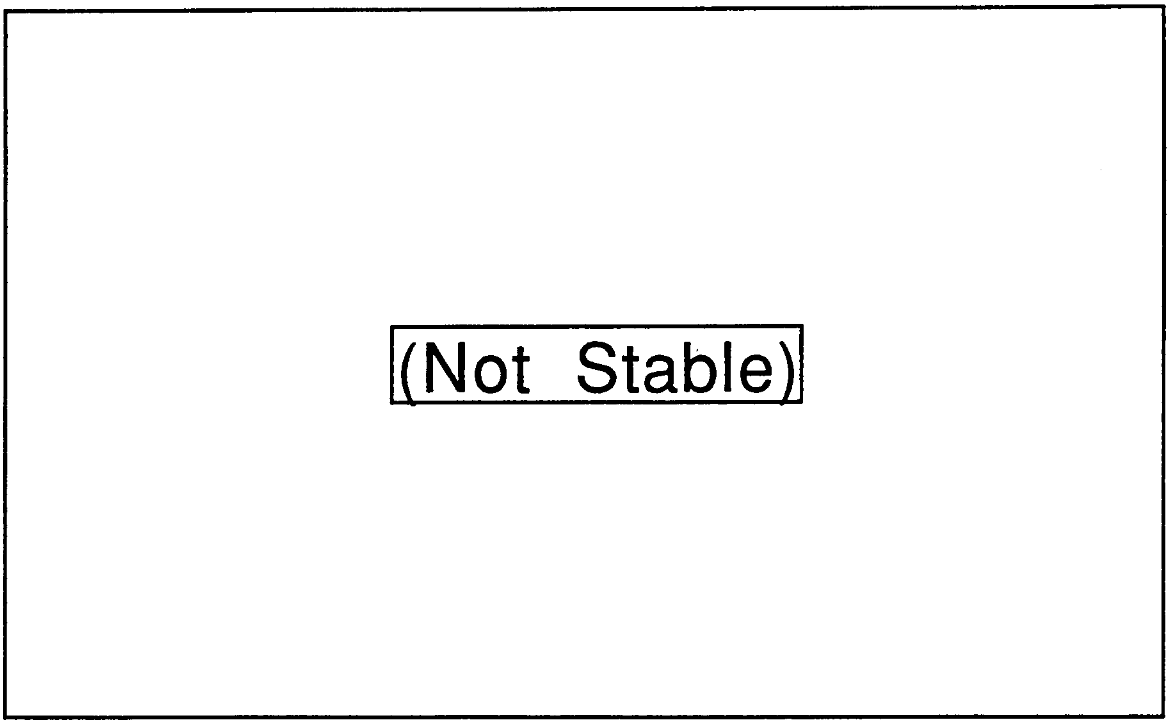

Figure D.5: Strategy 2b Response to Disturbance 1

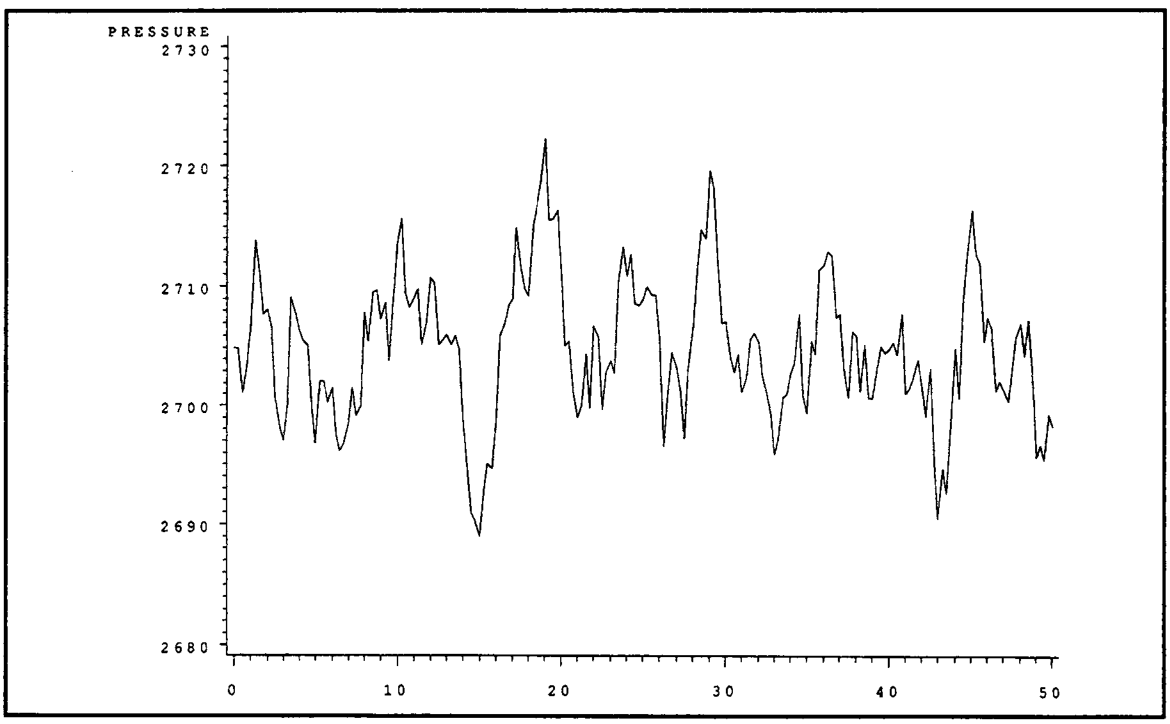

Figure D.6: Strategy 2b Response to Disturbance 4

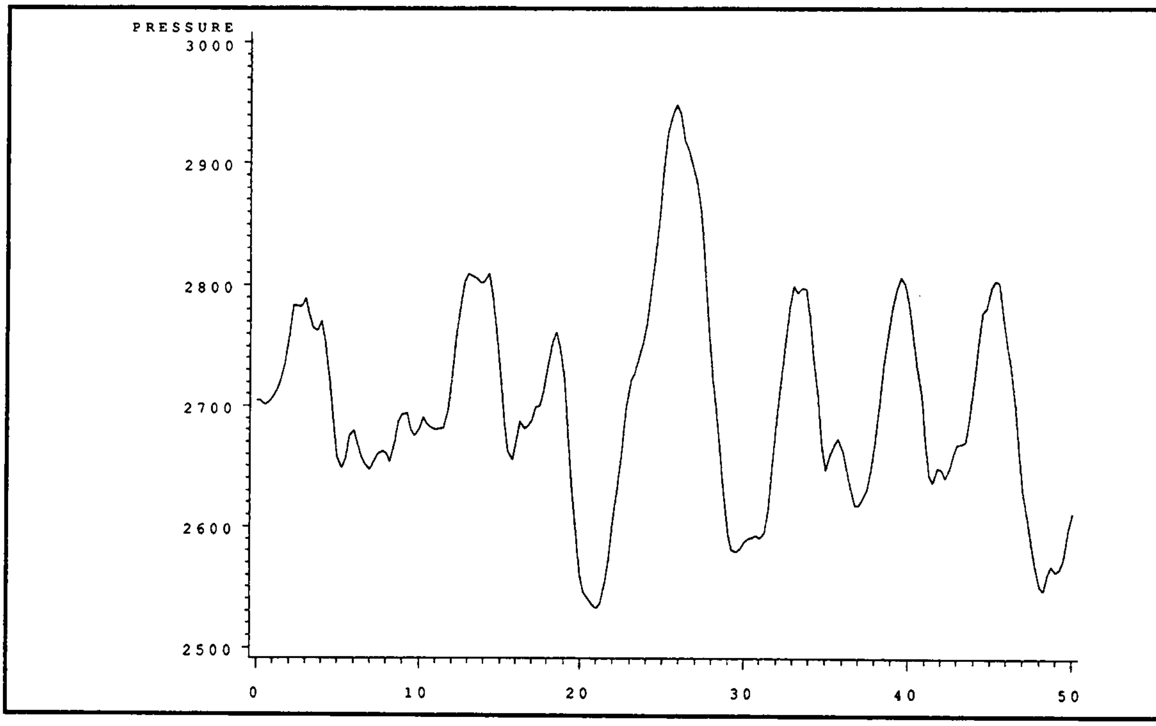

Figure D.7: Strategy 2b Response to Disturbance 8

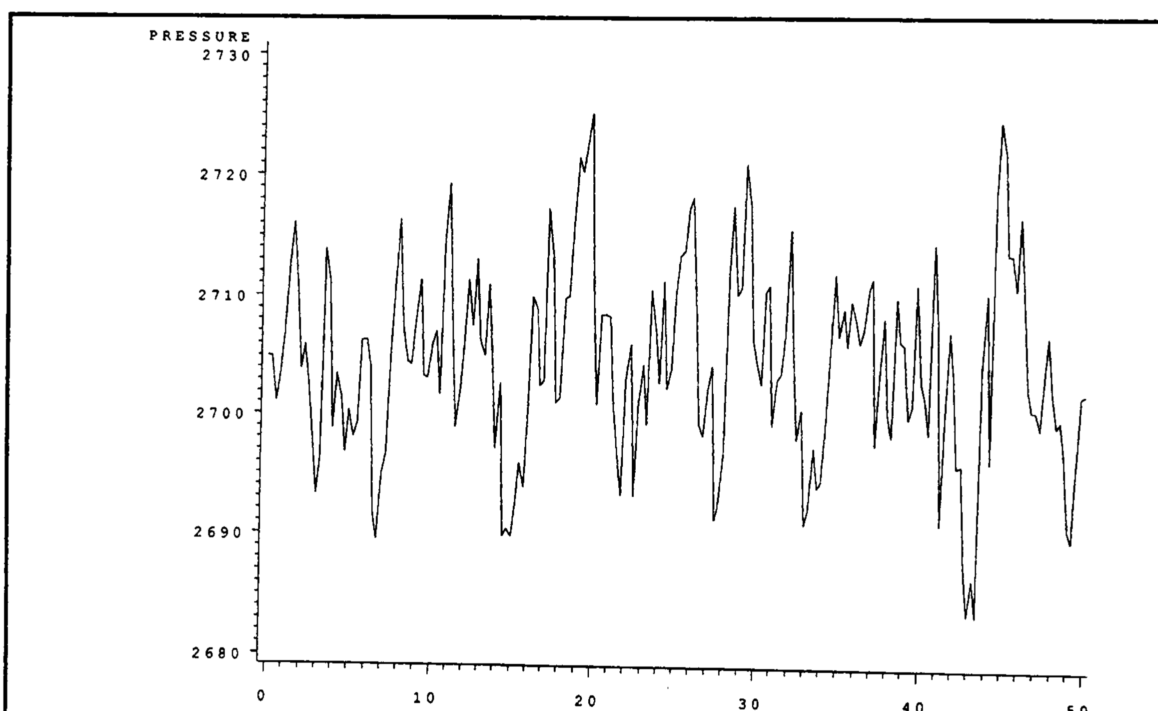

Figure D.8: Strategy 2b Response to Disturbance 12/15

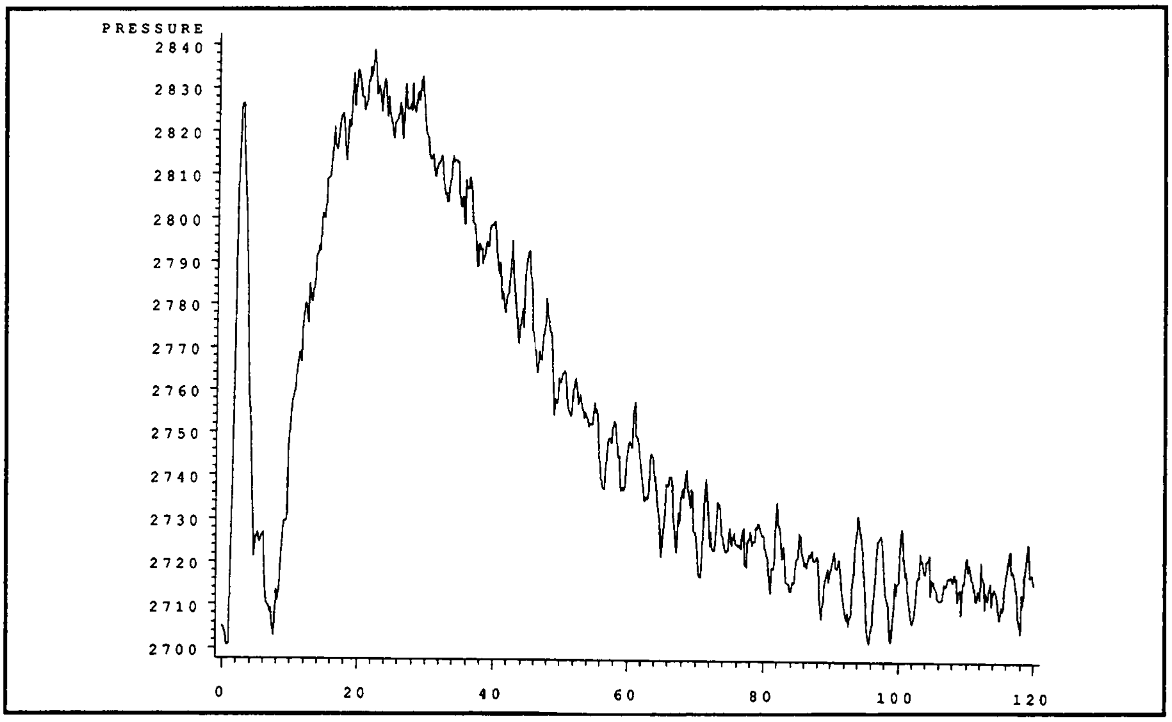

Figure D.9: Strategy 3 Response to Disturbance 1

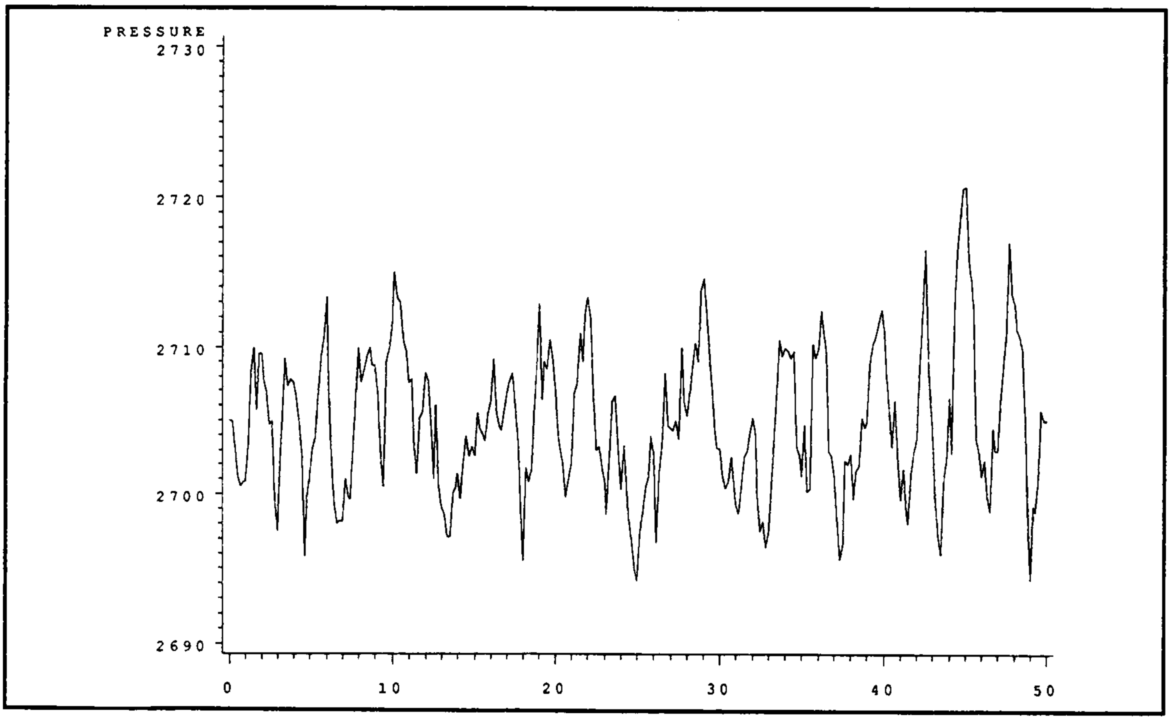

Figure D.10: Strategy 3 Response to Disturbance 4

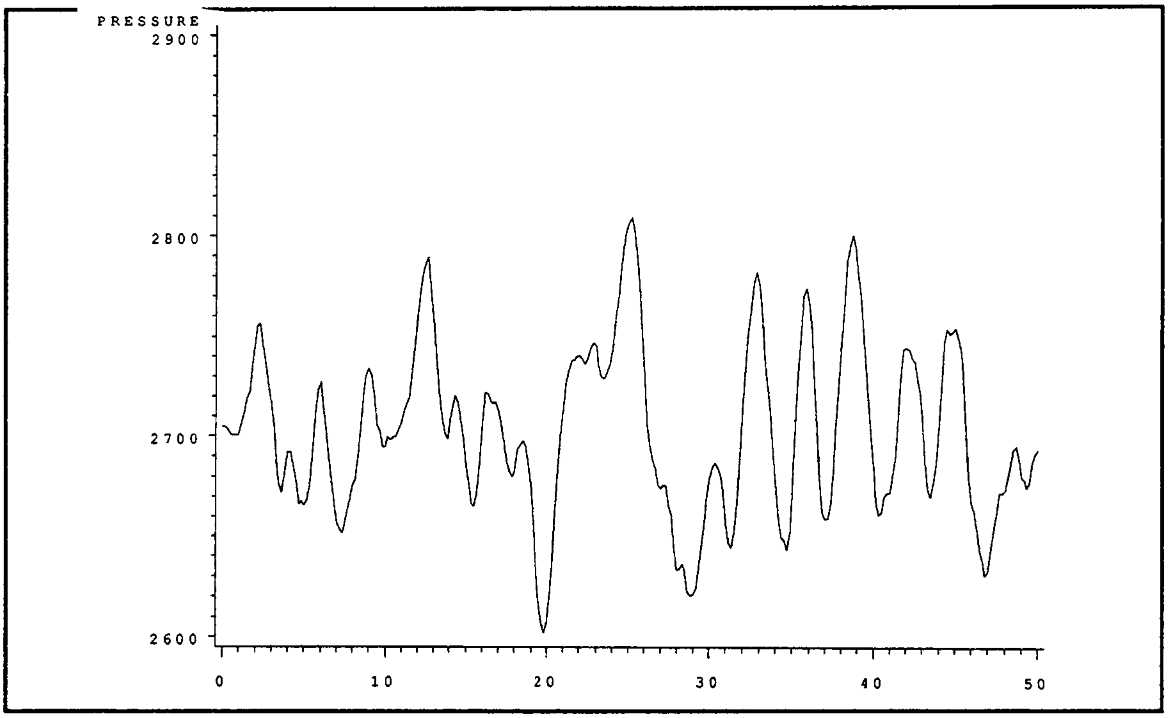

Figure D.11: Strategy 3 Response to Disturbance 8

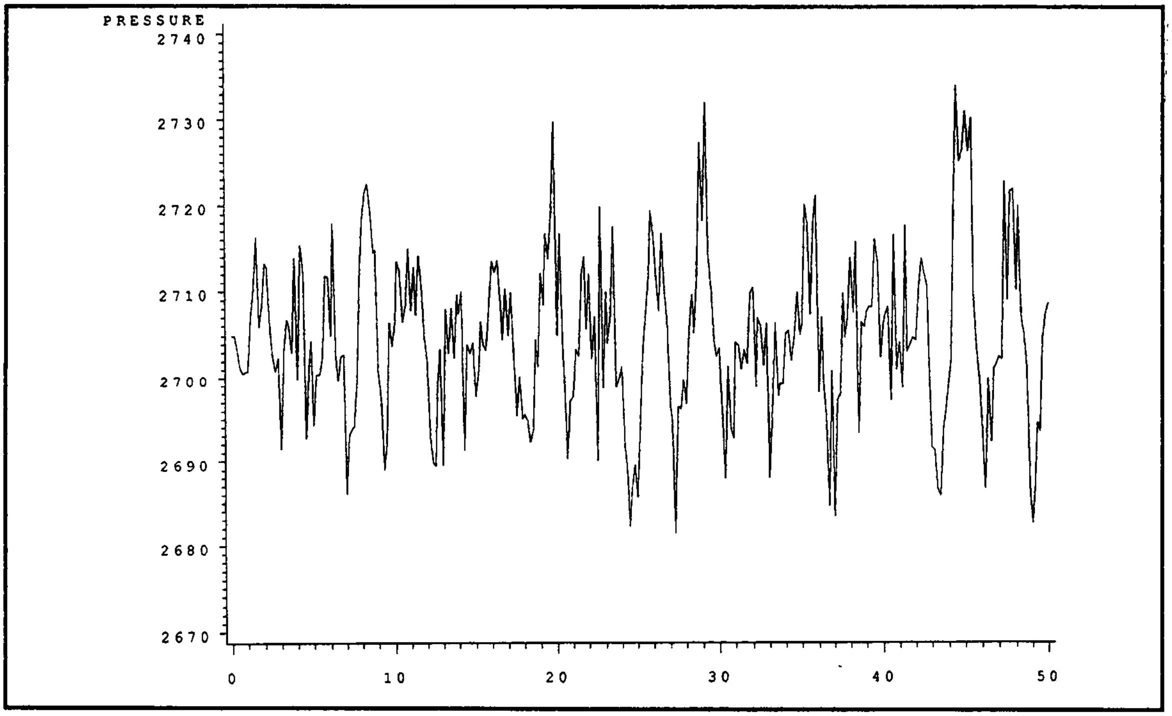

Figure D.12: Strategy 3 Response to Disturbance 12/15

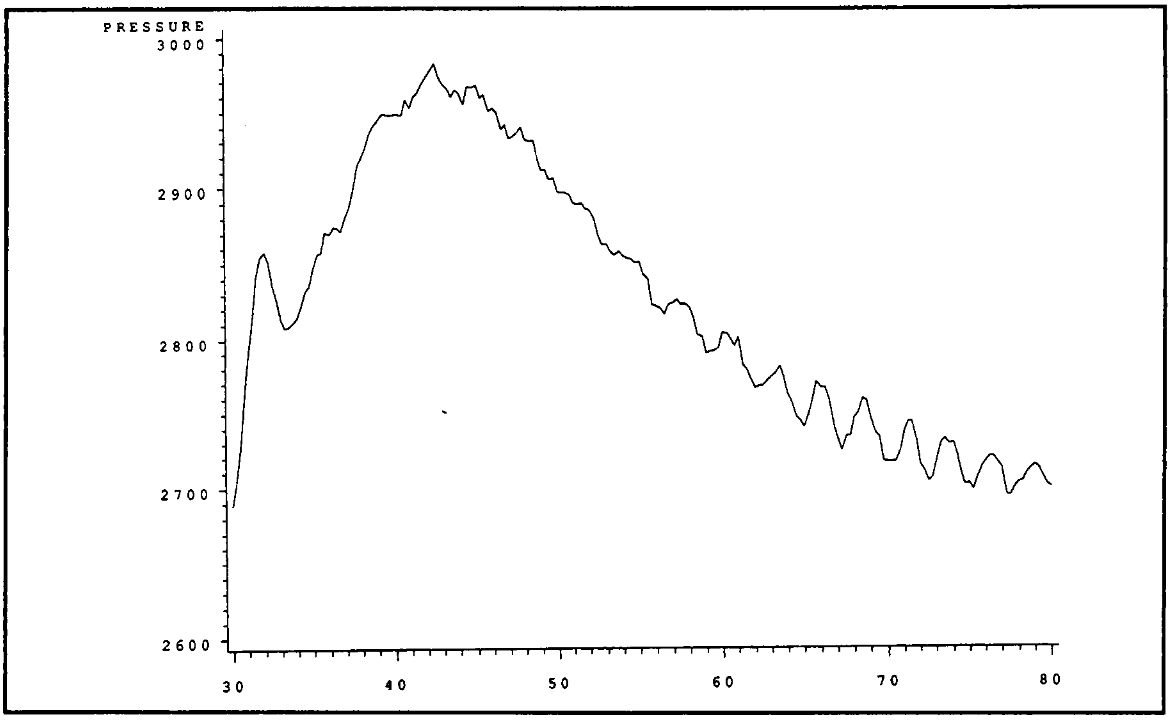

Figure D.13: Strategy 4 Response to Disturbance 1

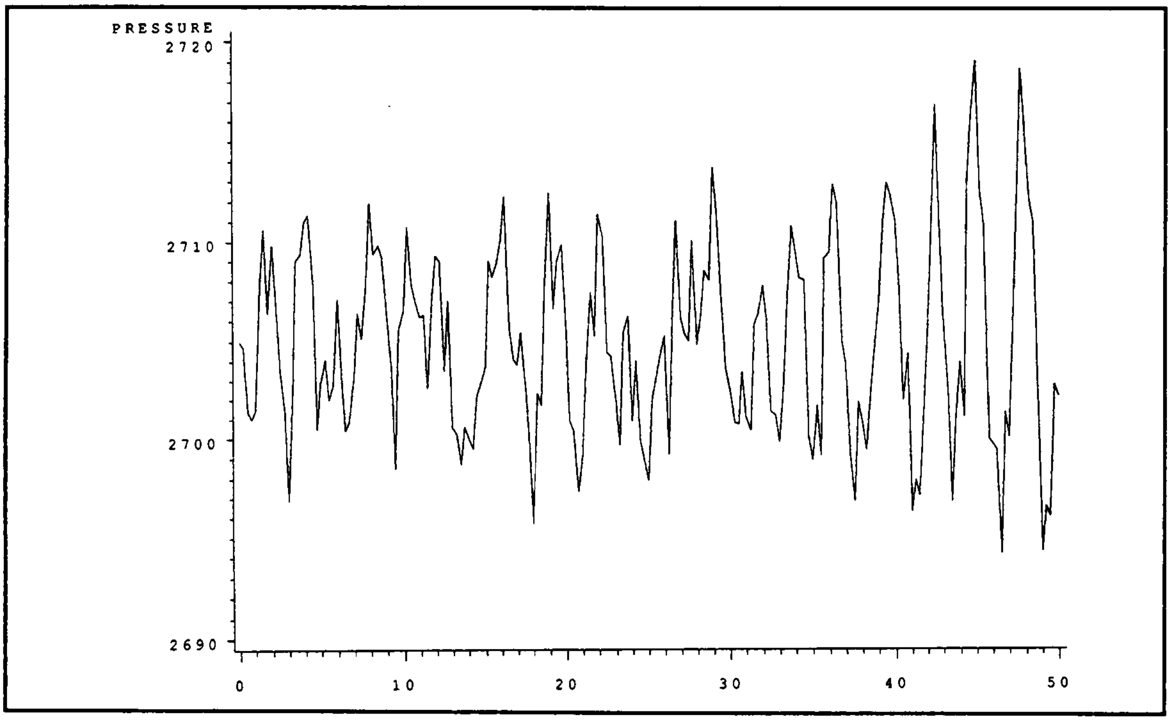

Figure D.14: Strategy 4 Response to Disturbance 4

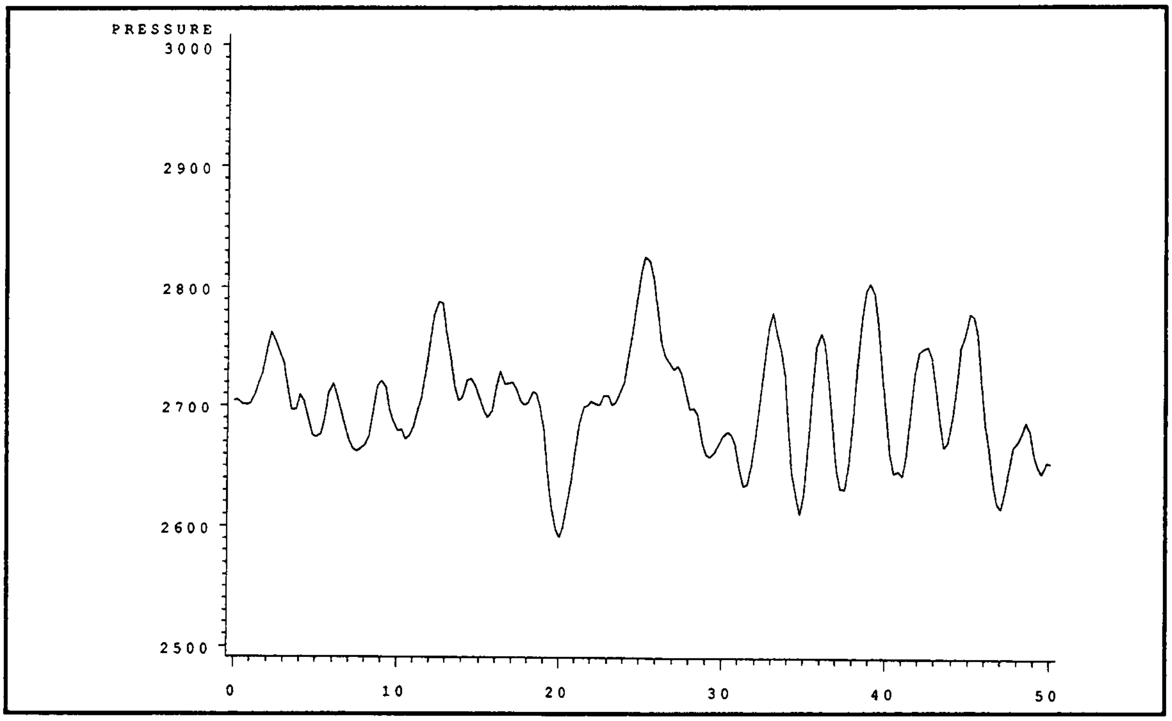

Figure D.15: Strategy 4 Response to Disturbance 8

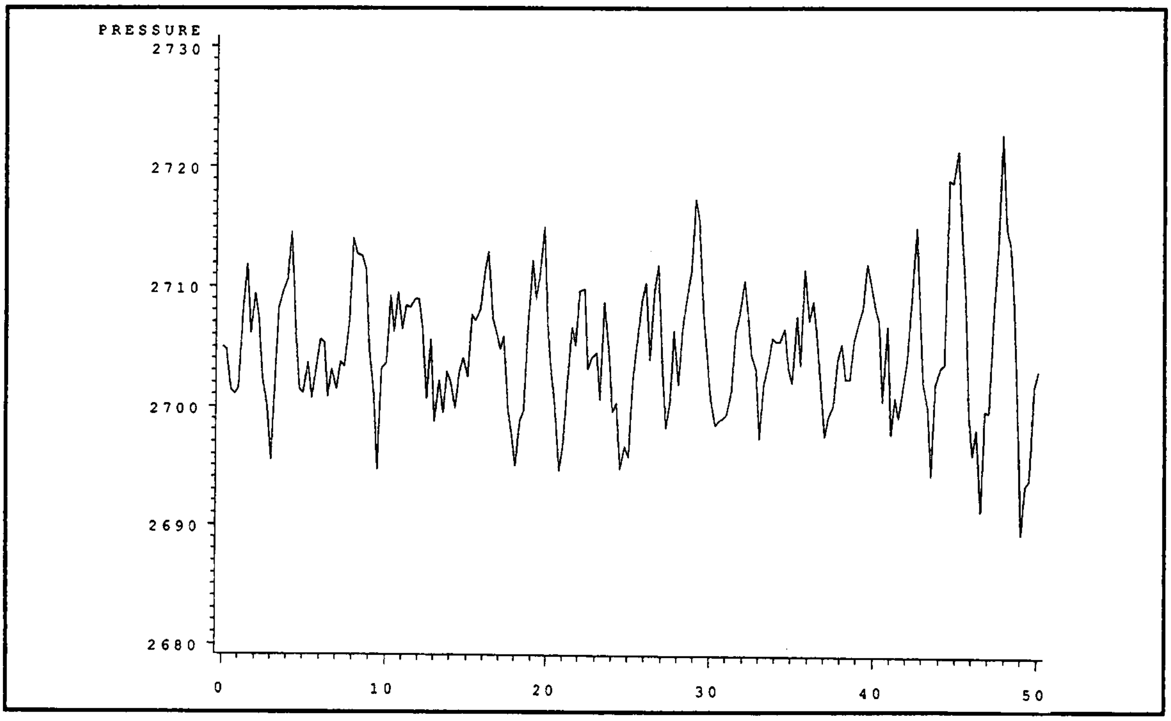

Figure D.16: Strategy 4 Response to Disturbance 12/15

VITA

Erin Percell was born October 27, 1967, in Biloxi, Mississippi. He attended public schools in Missouri and Tennessee before graduating as valedictorian of North Side High School in Jackson, Tennessee, in June, 1985. In August, 1985, he enrolled at Tennessee Technological University in Cookeville, Tennessee, from which he received his B.S. in Chemical Engineering in June, 1989. His professional experience includes a one year co-op assignment with Milliken & Company in Inman, South Carolina, and two years working in the Engineering Research Division of the Eastman Chemical Company in Kingsport, Tennessee, from 1989 until 1991. He entered the graduate program at the University of Tennessee, Knoxville, in August, 1991, and received his doctoral degree in August, 1996. Upon completing his degree, he will be employed at Exxon's Baton Rouge Chemical Plant in Baton Rouge, Louisiana.