

THE EFFECT OF ENERGY INTEGRATION
ON THE SWITCHABILITY OF CHEMICAL PROCESSES

by

Jeffrey Michael Howell
B.Sc.(Eng), A.C.G.I., M.Inst.M.C.

October 1984

A Thesis submitted
for the degree of Doctor of Philosophy of the University of London
and for the Diploma of Membership of the Imperial College

Department of Chemical Engineering and Chemical Technology
Imperial College of Science and Technology
London SW7

Contains

Microfiche

.

to Lin

ABSTRACT

Commercial pressures have encouraged the use of increasingly complex energy recovery schemes in chemical processes. Little research effort has been reported on the operability of processes with the high level of interaction associated with these schemes. This research project has directly attacked the problem of designing heat integrated chemical processes which are easily switched between widely spaced steady states. The work concentrates on switchability of networks of heat exchangers and distillation columns, chosen for their industrial importance and energy-intensive nature. Two aspects of switchability have been identified - feasibility and optimality. Theorems relating to feasibility of switches are presented for chemical processes in general and for some types of heat integrated networks in particular. From these it is apparent that for the systems under consideration, which have feasible steady states, switching feasibility is unlikely to be an issue. However, examples are given of switches which become infeasible as a direct result of heat integration.

The central part of the thesis considers optimal switching policies to derive an index of switchability, through which various heat integration schemes may be compared. Special models have been developed which are relatively simple while retaining the features considered important to switchability. The objective function is considered in detail as the basic index of switchability. Finally the development of the optimal control program and interface to an equation-oriented flowsheeting package is described. The package has been applied to three examples of energy integration; simple heat exchanger networks, distillation schemes with combined reboilers and condensers, and a complex heat-integrated distillation column for crude oil refining. A number of features have been identified which will influence the switchability of heat-integrated processes. Awareness of these features in the preliminary stages of design will help create intrinsically switchable networks.

Acknowledgements

I would like to express my deepest gratitude to Dr. John Perkins for his constant guidance and encouragement during the course of this work.

I thank British Petroleum, the Science and Engineering Research Council and the Imperial College Computer Centre for their financial and technical support.

I am indebted to Ken Morison, Mark Wong, Keith Hanson, Clive Beautyman and the SPEEDUP programming team for their valuable suggestions, criticisms and contributions to this research.

Finally I wish to thank Lin Howell for typing the manuscript.

TABLE OF CONTENTS

Abstract	3
Acknowledgement	4
Table of contents	5
List of figures and tables	9
CHAPTER 1 - INTRODUCTION	16
1.1 Introduction	16
1.2 Scope of the current work	18
1.3 Literature review	20
1.4 Significance of switchability	24
CHAPTER 2 - FEASIBILITY OF SWITCHES	26
2.1 How infeasible switches arise	27
2.2 Theorems on feasibility	34
2.2.1 A - General conditions for feasibility	35
2.2.2 B - Commutable switches	37
2.2.3 C - Global stability	37
2.2.4 F - Thermally coupled networks	38
2.3 A practical example	45
2.4 Summary	47
CHAPTER 3 - MEASURING SWITCHABILITY	48
3.1 Classification of dynamic performance indices	49
3.2 Direct profit estimate and marginal values	51
3.3 Quadratic off-specification cost function	54
3.4 Switch duration	57
3.5 Switchability and optimal control	60
3.6 Cost functions for switchability	61

CHAPTER 4 - MATHEMATICAL TECHNIQUES	65
4.1 Defining the optimal control problem	66
4.2 Solving the optimal control problem	68
4.3 Integration methods for mixed sets of equations	74
4.4 Interfacing with the minimisation package	79
4.4.1 Calculation of gradients	79
4.4.2 Limiting the linesearch step length	81
4.4.3 Scaling the decision variables	82
4.4.4 Accuracy of data	83
4.5 Convergence criteria and solution accuracy	84
4.6 Simulation using SPEEDUP	87
4.7 Problem definition through SPEEVOP	90
4.8 Implemented optimal control algorithm	94
4.9 Testing the optimal control package	97
4.9.1 Problem 1	97
4.9.2 Problem 2	98
4.9.3 Use of analytic jacobians	102
CHAPTER 5 - MODELLING OF HEAT EXCHANGERS	104
5.1 Requirements	105
5.2 Model classification	106
5.3 Applicability of previous work	110
5.4 Development of the new model	113
5.5 Evaluation of the new model	117
5.6 Utility exchangers	122
CHAPTER 6 - MODELLING OF DISTILLATION COLUMNS	123
6.1 Full order models	124
6.2 Extensions to the full order CHE model	127
6.3 Simplified full order models	129
6.4 Reduced order models	132
6.5 Node and section column model	135
6.6 COLUMN - Column section procedure	139
6.7 BUBBLE - Bubble point procedure	144
6.8 Evaluation of the new model	145

CHAPTER 7 - HEAT EXCHANGER NETWORK CASE STUDY	150
7.1 Network description	150
7.2 Determining the switches and objective	153
7.3 Closed loop control	156
7.4 Optimal control	163
7.5 Perfect control	169
7.6 Results and discussion	170
CHAPTER 8 - HEAT INTEGRATED DISTILLATION SYSTEM CASE STUDY	180
8.1 System description	180
8.1.1 LENMOA	182
8.1.2 LENMOB	182
8.1.3 LENMOC	182
8.2 Determining the switches and objective	183
8.3 Optimal control	186
8.4 Results and discussion	194
CHAPTER 9 - CRUDE OIL DISTILLATION SYSTEM CASE STUDY	200
9.1 System Description	200
9.2 The main column model	204
9.3 The heat exchanger networks	207
9.4 Switches and objective function	209
9.5 Optimal control	212
9.6 Results and discussion	219
CHAPTER 10 - CONCLUSIONS	224
10.1 Conclusions	224
10.2 Recommendation for future work	226

References	227
Table of commonly used symbols	235

APPENDICES

A	Atmospheric flash problem details	237
B	Heat exchanger model evaluation problem	239
C	Distillation column model evaluation problem	241
D	Heat exchanger case study details	245
E	Distillation system case study details	248
F	Crude oil refining case study details	254
G	Program listings	Microfiche (8 frames)

LIST OF FIGURES AND TABLES

CHAPTER 1 - none

CHAPTER 2

2.1	Exothermic reactor	27
2.2	Phase plane behaviour of (2.1)	28
2.3	Heat balance diagram	28
2.4	Exothermic reactor with heat recovery	29
2.5	Atmospheric flash process	30
2.6	Effect of a momentary increase in coolant flow	32
2.7	Effect of a momentary decrease in coolant flow	32
2.8	Neon liquefaction process	33
2.9	Homogeneous unit within a network	39

CHAPTER 3

3.1	Endpoint error of a first order system	63
-----	--	----

CHAPTER 4

4.1	Classification of single step methods	76
4.2	Line search for function minimisation	81
4.3	Effect of objective function error	83
4.4	VF02AD termination test in one variable	86
4.5	Logical data flow between SPEEDUP and SPEEVOP	91
4.6	Overlay structure of SPEEVOP	96
4.7	Optimal control policy for problem 1	98
4.8	Solutions to problem 1 by method	98
4.9	Nominal solution to problem 2	101
4.10	Effect of gradient scaling in problem 2	101
4.11	Progress of the solution for small scaling factors	101
4.12	Progress of the solution for large scaling factors	102
4.13	Progress of the solution with scaled Hessian	102
4.14	Comparison of Jacobian estimation methods	103

CHAPTER 5

5.1	Cascaded double convection model	107
5.2	Modified double convection model	108
5.3	Heat exchanger model with inverse response	111
5.4	Heat exchanger model with multiple solutions	112
5.5	Solutions of the model in (5.4)	112
5.6	First 4 eigenvectors of the cascade model	115
5.7	Proposed heat exchanger model	116
5.8	Dynamic model comparison	117
5.9	Response of model I to step temperature change	119
5.10	Response of model II to step temperature change	120
5.11	Response of model III to step temperature change	121

CHAPTER 6

6.1	Simple distillation tray	126
6.2	Lumped section binary column	133
6.3	Dahlquist column model	134
6.4	Simple binary node and section column	137
6.5	Conditions in a column section	139
6.6	Comparison of column solution methods	142
6.7	Model CHE25 flow diagram	146
6.8	Eigenvalues of model CHE25	147
6.9	Column section sizes for CLTEST	149
6.10	Steady state compositions and eigenvalues of major dynamic modes compared	149
6.11	Run times in CDC 6600 cpu seconds	149

CHAPTER 7

7.1	Network HEX7	151
7.2	Network HEX8	151
7.3	Inequality control constraint function	157
7.4	Graph of function (7.4) near origin	158
7.5	HEX7 closed loop control A-B	159
7.6	HEX7 closed loop control B-A	160
7.7	HEX7 closed loop control A-C	161
7.8	HEX7 closed loop control C-A	162
7.9	HEX8 closed loop control A-B	165
7.10	HEX8 closed loop control B-A	166
7.11	HEX8 closed loop control A-C	167
7.12	HEX8 closed loop control C-A	168
7.13	Closed loop switch costs	170
7.14	Optimal setpoint control switch costs	170
7.15	Perfect control switch costs	171
7.16	Steady state utility costs	171
7.17	Heater-cooler combination	174
7.18	Bypass control of a single exchanger	176
7.19	Steady state temperatures vs bypass control	178
7.20	Effect of a step change in u from 0.2 to 0.8	178

CHAPTER 8

8.1	Distillation networks	181
8.2	Effect of product purity on operating cost	185
8.3	LENMOA optimal control A-B	188
8.4	LENMOA optimal control B-A	189
8.5	LENMOB optimal control A-B	190
8.6	LENMOB optimal control B-A	191
8.7	LENMOC optimal control A-C	192
8.8	LENMOC optimal control C-A	193
8.9	Steady state and dynamic costs	194
8.10	System with a right-half-plane zero	197

CHAPTER 9

9.1	Main column schematic	202
9.2	Heat recovery networks	203
9.3	Boiling points of n-paraffins at 2 bars	205
9.4	Optimal control policies for switch A-B	215
9.5	Optimal target deviation for switch A-B	216
9.6	Optimal control policies for switch B-A	217
9.7	Optimal target deviation for switch B-A	218
9.8	Relative costs of networks 1-3	219

CHAPTER 10 - none

APPENDIX A

A.1	Process flow diagram	237
A.2	Steady state solutions	238
A.3	Vapour fraction function $\phi(T_4)$	238

APPENDIX B

B.1	Process flow diagram	239
B.2	Steady state operating conditions	239
B.3	Network of units for model 1	240
B.4	Eigenvectors of model 1	240

APPENDIX C

C.1	CHE25 unit interconnections	241
C.2	Tray parameters	242
C.3	Condenser unit	242
C.4	Reboiler unit	242
C.5	Control unit parameters	243
C.6	Column conditions for $x_F = 0.5$	243
C.7	Column conditions for all feeds	243
C.8	Composition eigenvectors of the two models	244

APPENDIX D

D.1	Nominal steady state design problem	245
D.2	Initial and target operating conditions	245
D.3	Network HEX7	246
D.4	Steady state controls for HEX7	246
D.5	Miscellaneous data for HEX7	246
D.6	Network HEX8	247
D.7	Steady state controls for HEX8	247
D.8	Miscellaneous data for HEX8	247

APPENDIX E

E.1	Network LENMOA	248
E.2	LENMOA Physical data	248
E.3	LENMOA Optimal operating conditions	249
E.4	Network LENMOB	250
E.5	LENMOB Physical data	250
E.6	LENMOB Optimal operating conditions	251
E.7	Network LENMOC	252
E.8	LENMOC Physical data	252
E.9	LENMOC Optimal operating conditions	253

APPENDIX F

F.1	General arrangement of the system	254
F.2	Schematic of the column model	255
F.3	Column mechanical data	256
F.4	Miscellaneous data	256
F.5	Column steady states	257
F.6	Column operating limitations	258
F.7	Heat recovery network 1 - BP	259
F.8	Network 1 mechanical data	260
F.9	Network 1 steady states	260
F.10	Heat recovery network 2 - Hanson	261
F.11	Network 2 mechanical data	262
F.12	Network 2 steady states	262
F.13	Heat recovery network 3 - Minimal	263
F.14	Network 3 mechanical data	264
F.15	Network 3 Steady states	264

APPENDIX G - Not applicable

This page intentionally left blank

CHAPTER ONE

1.1 INTRODUCTION

Increasing raw material and energy costs have encouraged the design of chemical plants which are both product selective and energy efficient. Although developments in catalysis and unit operations have contributed much to selectivity, the major thrust towards energy conservation has come from waste heat recovery. The plant designer is now concerned with energy integration at a relatively early stage in the project, since the complex nature of modern integration schemes can influence many aspects of a design.

Because of the direct relationship of operating cost and energy consumption, it is not surprising that a substantial research effort has been put into the design of energy recovery systems, notably networks of heat exchangers for process fluids. Algorithms now exist to generate heat exchanger networks automatically (subject to some simplifications) which will maximise heat recovery while minimising some characteristic cost index. The problem of designing such networks to maximise an arbitrary objective which is truly representative of cost is still a subject of research, but the basic technique of Flower and Linnhoff (1980) has gained wide acceptance in industry.

Despite the success of these designs for heat recovery at some nominal operating point, the designs are not necessarily tolerant of disturbances in the operating conditions. These may be deliberate, such as changes in set-point, or unavoidable, like a catalyst degradation. The increased interaction between various stages in the plant via the heat exchange system might also cause hitherto unsuspected operating problems. If the design which offers the lowest nominal operating cost is unable to handle even small deviations in operating conditions, it will not be the best choice. The design objective is thus complicated by three other desirable attributes - flexibility, controllability and switchability.

Flexibility allows the process to operate at other steady states without sacrificing efficiency or throughput. A flexible process has sufficient degrees of freedom and range of control to cope with all foreseeable changes in the steady state operating environment. Flexibility of distillation columns has been studied by Maarleveld and Rjinsdorp (1970), and of heat exchange networks by Morari, Marselle and Rudd (1982), amongst others.

Controllability is a measure of how easily the process may be regulated to the desired operating state. It depends on how a deviation from the target state is sensed, and the effectiveness of allowable corrective control action. Such properties are independent of the control system itself, so may be used in the process design stages. Features believed to degrade controllability include pure time delays, right-half-plane zeroes in the transfer function matrix, and modelling error sensitivity. Classical control theory for linear systems provides a valuable background to these topics, providing that the disturbances are small enough for the linear assumption to hold. Morari (1982) proposes an index of controllability which penalises pure time delays between the controls and the outputs of a system. Wong (1984) has investigated the subject of controllability in detail with particular emphasis on heat integrated systems.

Switchability describes how readily a process switches between different operating states, and is the subject of this thesis. It is concerned with the dynamic behaviour of systems which experience major disturbances in operating conditions, often exceeding the limits which have been set for flexibility. Under these conditions, the range of control action will be constrained by physical and safety factors. Non-linear features of the system response will also be emphasised, so any analysis based on linear models will be inappropriate. This is the essential difference between controllability and switchability. It is clear that the study of switchability must exploit techniques which do not dilute the underlying causes of the problem, namely, control constraints and non-linear response.

1.2 SCOPE OF THE CURRENT WORK

Before the problem can be discussed, the scope of the study must be defined, particularly the meaning of the terms 'switch' and 'switchability' in the current context.

Take for example an air separation plant. The products, namely oxygen, nitrogen and inert gases, are well defined, but the demand will vary from hour to hour. The cost of electricity, the prime raw material, will also vary according to price schedules which reflect the industrial and domestic power demand pattern. The most efficient operating state at any instant will be a function of these and other factors, such as available manpower and general condition of the plant, and can usually be calculated by routine optimisation programs. The operating policy actually chosen will not necessarily attempt to follow this exactly because the shift of operating state is not instantaneous, and generally causes some loss of production. Instead, when a long term trend has been identified, the decision to 'switch' the plant must be made, and a control policy implemented to accomplish this. Short-term fluctuations in operating conditions are controlled about the new set-points by conventional control schemes. A switch is therefore a transition between two nominal steady operating states, subject to any physical and safety constraints on the process.

Switchability is an index of difficulty of achieving a given switch (or family of switches), which is determined by the process itself. Since the control policy is arbitrary, the control system is irrelevant - the study considers only open-loop properties wherever possible. The economic cost of a switch is such an index. If a switch is infeasible the cost is infinite, whereas perfect switchability implies that it is free, i.e. uses no more utilities etc. than normal operation.

For any realistic chemical process there will be many control policies

which will achieve a given switch, so which one should be used to compare switchability? As we are considering process networks independently of the associated control system, we can assume that the control policy is 'optimal' - that is, it will give the lowest cost switch possible subject to the constraints imposed by the system itself. This optimal policy may be deduced by the mathematical techniques described in subsequent chapters.

It has been assumed throughout that the switch occurs between states which are feasible steady states. It is perfectly possible to consider switching from unsteady states that a system could be driven to by abnormal conditions; but it is felt that this is an unnecessary complication of limited practical value.

1.3 LITERATURE REVIEW

It appears that no previously published work directly addresses the problem of switching heat integrated systems. Several relevant aspects have been tackled in recent years, and the work falls broadly into two categories - the constraint control approach and the global optimisation approach. Two examples of the former and three of the latter have been found.

Constraint control covers any scheme which is partitioned in some way into a sequence of steady states. The switch is accomplished by hopping from one state to the next, and relies for its success on choosing the correct sequence. Each state is chosen according to the active constraints on control or process variables, and the optimal dynamic policy is assumed to pass through these states. To obtain a true measure of switchability we must consider schemes which do not pass through intermediate states, so this approach does not appear suitable.

Global optimisation includes optimal control and optimal design of process plant. Constraints, where present, are readily incorporated into the optimising package, so the active constraints do not need to be determined a priori. Unlike the former approach, transient behaviour is also optimised. Although this can satisfy all the requirements for a study of switchability, it is computationally very much more expensive.

The work of Arkun and Stephanopoulos (1980) is in some ways close to the present study, in that it identifies both the feasibility and optimality aspects of switching. Following the established principles of constraint control, the task of switching a single process unit is treated as a constrained optimisation problem to find a minimum cost route for the steady state which joins the initial and target states. The switch is achieved by moving the set points of all controllers along this route, and making the appropriate assignments for manipulated variables according to the active constraints.

Because the optimum will usually lie at a vertex in the manifold of feasible steady states, and that each edge corresponds to a different combination of constraints, the switch route is along those edges which connect the initial and target states. The vertices along this route form the sequence of steady states which make up the switch. To determine the optimal sequence, all sequences which connect the terminal states are screened on physical and safety grounds. Of the remainder, starting at the initial state, each step is chosen to minimise the cost objective, based on the local gradients. Thus a feasible and efficient (but not necessarily optimal) switching policy can be constructed.

Throughout the paper, the authors assume that the actual state will follow the set points closely. This highlights more drawbacks to this method when applied to switchability:

1. Constraints may be violated during transients.
2. Transient costs are equated to target steady states.
3. Transient duration is not quantitatively compared.

Furthermore, the optimisation strategy does not consider routes which might lead to very cheap terminal sequences, although they are initially more expensive. Neither does it consider schemes which drop more than one constraint at a time. For these reasons the technique does not really address the problem of switchability, and will not be considered further.

Nishitani and Westerberg (1982) looked specifically at process start-up of a distillation system. They identified the major problem as bringing the column from cold to a state of total reflux, subject to limited utility availability. This was solved, for both a single isolated column, and a pair of columns which share utilities, using a transportation problem type of formulation.

The resultant optimal policies show limiting control action for much of

the start up period, reducing the start-up policy to a simple sequence of intermediate states. This paper also includes a brief discussion of column modelling and explains that a column may be switched from a state of total reflux and an arbitrary composition profile to full production. This is described as 'a typical control problem ... (for which) an attractive solution is obtained without solving the variational problem'. No indication is given of how attractive the solution is, or whether it will apply to more complex systems.

Pollard and Sargent (1970) use numerical optimisation techniques to find the optimal startup and shutdown policies for a simplified binary distillation column model. A control policy in which the controls are always at the bounds is optimised by adjusting the switching times and the result is used as a starting point for a comprehensive optimal control algorithm. The results show that there are two distinct phases to each transition, one in which all controls are extreme, and the other with an intermediate level on one or more controls. Some extreme control action is necessary throughout the switch, lending support to the constraint control theories since they were derived by an independent route.

Sullivan (1978) developed the techniques further to assist with finding a minimum cost feed changeover policy for a crude oil separation system. The optimal control policy was found and a simple practical closed loop policy was derived from it. He found that the switching operation could be split into two stages, one to establish the new flow rates and the other to switch to the required compositions, without degrading the total switch cost. While this also supports the sequential switching theories, the numerical results show virtually no evidence of extreme control action. These results must therefore be treated cautiously, since they are not directly compatible with other work in this area.

Hitherto we have only considered switching existing designs, but the underlying purpose of this research is to develop some method for use at the design stage, so that the plant is intrinsically switchable.

This can take two forms, either a heuristic guide, or, continuing the previous theme, a dynamic component of a structural optimisation package. There are no references to heuristics to promote switchability, although such things are commonplace in control and steady-state design. This is probably because the subject has not previously been researched, rather than any fundamental theoretical problem.

However, Nishida and Ichikawa (1975) chose to tackle the problem at the highest level, namely as a problem of design of a dynamic process. This is the only approach which refers to a system of several process units with variable structure. The optimal control policy is embedded in the decision variables for the entire design, which include structure parameters and equipment sizes.

Results are presented for two simple examples, but it is clear that the computational requirements severely limit its application. In particular, it is not obvious how to extend the algorithm to cover a range of operating conditions, although the batch processes which were studied are well suited to it. The use of structure parameters to describe the network topology has been around since the late sixties, but most attempts to apply them to system synthesis have foundered on the combinatorial difficulties of including all possible structures in the initial formulation. For this work, the topology will be determined largely by the steady state requirements, and the only structural decisions will relate to bypasses and stream splits in the heat recovery scheme. Under such conditions the relatively small number of structure parameters can be included easily in the total optimal control set. The technique therefore has some value in the current study, but is unlikely to form part of any practical method for designing switchable processes.

1.4 SIGNIFICANCE OF SWITCHABILITY

There are several reasons why switchability has been neglected in the synthesis of chemical plants. Firstly, traditional processes do not in general exhibit switching problems. Any deficiencies in the controllers are usually smoothed over by the inherent open-loop stability of the process, and the system is left to 'settle down' without further intervention. Secondly, it is assumed that the controllers will have sufficient proportional bandwidth to cope with all likely disturbances. The assumption is often justified because the control elements are habitually oversized. Finally, the duration of a switch that could halt production is very small compared with normal steady-state operation.

Switchability should be recognised by the designer of modern plants for the following reasons, which all stem from desired increases in plant efficiency. To begin with, the designer must be aware of the cost of switching, as even small savings can represent a large proportion of a narrow profit margin. Next, extensive heat integration introduces many new constraints and interactions within a complex plant. A formerly stable plant can easily become open-loop unstable due to heat integration, rendering traditional start-up methods unworkable. Examples of this are given in the next chapter.

Lastly, the practice of oversizing to ensure controllability is expensive, and does not automatically give sufficient control range for an efficient switch. For example, it is useless to increase the area of a heat exchanger if the desired cold stream outlet temperature exceeds the hot stream inlet temperature.

It should be noted that neither steady state analysis nor classical control theory may predict any of these problems. Both dynamic simulation and optimal control studies are impractical at the initial stages of a plant design. A simple, direct method is needed to recognise designs which are prone to poor switchability and eliminate them at an early stage. The purpose of this work is to investigate the

possibility of developing such a method for a restricted class of plants, for which heat recovery is an important feature. This work concentrates on networks of heat exchangers and distillation columns for several reasons:

1. The models are relatively simple.
2. The processes are industrially important.
3. The large scale of many separation processes make the comparatively small savings from switchability improvements very significant.

The study covers two areas - feasibility and optimality, of a plant with respect to switching. The section on feasibility is fairly general, but its main purpose is to show that switch feasibility is not a major issue for distillation systems. The determination of optimal switching policy is directly related to finding a quantitative representation of 'switchability' for such systems.

CHAPTER TWO

FEASIBILITY OF SWITCHES

As mentioned in the previous chapter, the cost of an infeasible switch is essentially infinite. Any design procedure in which decisions are made on the basis of switchability must be able to identify designs which possess infeasible switches. This chapter describes how such switches can arise in practical process networks, and offers theorems with which designs can be screened for infeasible switches.

2.1 HOW INFEASIBLE SWITCHES ARISE

The concept of infeasible switches is best illustrated by examples. The first example may be constructed from the classical exothermic reactor system of Bilous and Amundsen (1955), shown below.




Figure 2.1 Exothermic reactor

An exothermic reaction takes place in a well-mixed tank at a rate dependent on the temperature, which is regulated by a heating/cooling coil. In this case the feed conditions are fixed. Bilous and Amundsen considered the relationship between the concentration and temperature within the reactor given different starting conditions. The phase plane plot (2.2) shows how each initial state leads to one of two stable steady states, according to which side of the separatrix abc it lies.

The steady states (e,b,f) and separatrix (a,b,c) are functions of the control, which in this instance is the heating/cooling rate. The behaviour of these states is readily apparent if the heat production and consumption are plotted against temperature, as suggested by Van Heerden (1953), figure 2.3.




Figure 2.2 Phase plane behaviour of 2.1




Figure 2.3 Heat balance diagram

The three steady states correspond with the intersections of the two lines. The effect of additional heating shifts the dashed line to the right, and vice versa. The effect will be limited by the bounds on the heating and cooling rates. As the heating rate is increased, the unstable steady state moves downwards to meet the stable steady state, and they merge at h , negating one another. The heating rate at which this occurs is the minimum required to guarantee that a switch from e to f is feasible. Similarly the dashed line must lie on the left of g if a switch from f to e is to be feasible.

In other words, the heating duty must be sufficient to push the reactor temperature to an 'ignition' point and the cooling duty must be able to 'quench' the reactor, to ensure switchability. These duties will generally exceed those required for normal regulatory control. Van Heerden identifies switchability as a significant issue when he states 'the most characteristic feature is the necessity of an ignition by external heating before a steady state can be established'.

A closely related situation arises when the reactor product stream preheats the feed stream in an external heat exchanger. The general arrangement is shown in figure 2.4.




Figure 2.4 Exothermic reactor with heat recovery

Shinskey (1979) p262 compares the control problems of such a system with the single well-mixed reactor. The similarity is easily seen if the two intersecting curves of Van Heerden's heat balance graph are identified within the system. The S-shaped curve for heat generation arises from the nonlinear dependence of reaction rate on temperature as before. The straight line for heat loss is now due to the difference between feed and product temperature, which is directly related to the reactor outlet temperature. Thus the same form of graph results and similar conclusions can be drawn.

The system may be regulated as suggested by Shinskey, with bypasses and trim heaters, but special arrangements are still required to achieve 'ignition' and 'quench'. Note that the switchability problems are entirely due to energy integration - if independent heaters and coolers of sufficient capacity for steady state operation were employed the system would be unconditionally switchable.

The previous example contains a highly non-linear relationship between reaction rate and temperature, and it is plausible that this is the underlying cause of the multiplicity of steady states. One may suggest that a non-reacting system will possess a unique steady state. The following example proves this conjecture to be false.




Figure 2.5 Atmospheric Flash Process

A hot, pressurised fluid is cooled below its normal boiling point and flashed at atmospheric pressure to remove dissolved gases. It is further subcooled to avoid cavitation in the pump before it is reheated by the incoming feed. No heat is created or transferred except by simple conduction processes. Details of the model are given in appendix A.

The normal operating condition of this system (S_1) has a flash vapour fraction around 2%, maintained by controlling the cooling water flow. However, the equations describing the system have two other steady state solutions, with 54% (S_2) and 100% (S_3) vapourisation. The former is an unstable state, which must lie on the separatrix between the other two. This instability is indicated by the positive sign of the system eigenvalue at S_2 and verified by dynamic simulation. A unit perturbation (positive or negative) in the cooling water flowrate is sufficient to drive the operating state towards S_1 or S_3 respectively, as shown in figures 2.6 and 2.7. The latter state exists when the feed is not being cooled by product, so it flashes entirely to vapour.

Consider a switch from S_1 to S_3 . If the cooling water is shut off completely, it can be shown that S_3 becomes a unique steady state, so the switch is feasible. The reverse switch is not. At S_3 , there is virtually no flow through the cooling heat exchanger, and the control has negligible effect. Again, the problem is caused by the integration of heat transfer processes. One solution is described in section 2.3 as an example of the application of theorems developed later in this chapter.




Figure 2.6 Effect of momentary increase in coolant flow




Figure 2.7 Effect of momentary decrease in coolant flow

ISF.T2 refers to T2 in Appendix A, p237

Having shown that both chemical reaction kinetics and temperature induced flow changes can cause infeasible switches, one may be tempted to assert that nonreacting constant flowrate systems have unique steady states. Even this is not sufficient, as this last example shows.




Figure 2.8 Neon Liquefaction Process

Neon gas at high pressure is cooled to near its liquefaction temperature by a return flow of lower pressure neon. An expansion valve provides adiabatic cooling to liquefy the gas. This system relies on the Joule-Thomson effect to provide cooling, and this only operates below a certain temperature, which for neon is around 220K. Above this temperature (known as the Joule-Thomson inversion temperature) adiabatic expansion actually 'heats' the gas. A little thought will show that, as given, the system can never be cooled from ambient temperature.

2.2 THEOREMS ON FEASIBILITY

The examples show that both limited control action and multiple steady states can introduce infeasible switches. Theorem A helps to measure the significance of each of these factors, and also to set up a framework for subsequent theorems.

Firstly, the conditions for existence and uniqueness of solutions for differential equations must be considered.

Lipschitz condition:

A function $\underline{f}(\underline{x}, t)$ satisfies the Lipschitz condition whenever the following inequality holds,

$$|\underline{f}(\underline{x}', t) - \underline{f}(\underline{x}, t)| < k|\underline{x}' - \underline{x}| \quad (2.1)$$

where k is a positive constant which may depend on the space X and $\underline{x}, \underline{x}' \in X$. The norm $|\underline{x}|$ is defined by

$$|\underline{x}| = \sup [|\underline{x}_i|] \quad (2.2)$$

Conditions for existence and uniqueness of solutions:

A system of equations

$$\frac{d\underline{x}}{dt} = \underline{f}(\underline{x}, t) \quad (2.3)$$

satisfy the Lipschitz condition and is continuous in a space X . Then a unique trajectory $\underline{x}(t)$ of the system passes through every point $\underline{x} \in X$.

These conditions are stated fully in Lefschetz (1957) p30 et seq., with precise definitions of the vector space X .

The majority of chemical engineering equations satisfy the Lipschitz condition and the cause of exceptions can often be traced back to imperfect modelling. For example, the equation for the flow F of liquid through an orifice under a pressure difference P is frequently represented by

$$F = kP^{\frac{1}{2}} \quad (2.4)$$

This does not meet the Lipschitz condition, since as $P \rightarrow 0$ the gradient dF/dP increases without limit. However, the equation is only accurate for turbulent flow, and at very low flowrates the flow is laminar, hence it is proportional to P .

2.2.1 Theorem A: (General conditions for feasibility)

A system is defined by a set of ordinary differential equations which satisfy the Lipschitz condition,

$$\dot{\underline{x}} = \underline{f}(\underline{x}, \underline{u}) \quad (2.5)$$

and there exists a steady state $\underline{x}_{ss}, \underline{u}_{ss}$ for which

$$0 = \underline{f}(\underline{x}_{ss}, \underline{u}_{ss}) \quad (2.6)$$

where

$$\underline{x} \in \mathbb{R}^n, \underline{f} \in \mathbb{R}^n, \underline{u} \in U \in \mathbb{R}^r$$

Associated with this steady state is an escape space X_E and a capture space X_C defined thus:-

$$\underline{x}(t) \in X_E \text{ if } \underline{x}(0) = \underline{x}_{ss} \text{ for some constant } \underline{u} \in U \quad (2.7)$$

$$\underline{x}(0) \in X_C \text{ if } \underline{x}(t) \rightarrow \underline{x}_{ss} \text{ as } t \rightarrow \infty \text{ for } \underline{u} = \underline{u}_{ss} \quad (2.8)$$

Then a switch between two steady states $\underline{x}_{ss,1}$ and $\underline{x}_{ss,N}$ is feasible if there exists a finite sequence of states $\underline{x}_{ss,i}$ for which

$$X_{E,i} \cap X_{C,i+1} \neq 0, \quad i=1,2,\dots,N-1 \quad (2.9)$$

Proof:-

From the definition of $X_{E,1}$ there is a feasible constant control $\underline{u}$ which will drive $\underline{x}$ to any state within $X_{E,1}$. Select a state within the non-null intersection of $X_{E,1}$ with $X_{C,2}$ and apply the appropriate constant control to reach it. Then by definition of $X_{C,2}$ the state can reach $\underline{x}_2$ from here with steady state constant control action. This process can be repeated so the system passes through states $\underline{x}_3$, $\underline{x}_4 \dots$ to $\underline{x}_N$.

Discussion

The capture space X_C is the domain of the state space within which the system is asymptotically stable, so any state within X_C will converge to $\underline{x}_{ss}$. The escape space X_E is the reachable domain using a step change in the controls. It is clear that if the escape space overlaps the capture space of another state, a switch to the latter is feasible. Theorem A is a formal extension of that observation. The theorem is very general, and applies to systems with limit cycles, multiple steady states, and so on. It is now straightforward to show that the escape space is determined by the range of control action, and the capture space is dependent only on the topology of the n dimensional state space.

Escape space:- Assume that the system has a unique dynamic solution under constant control action $\underline{u}$ from a state $\underline{x}(t)$. Then every trajectory which starts from $\underline{x}_{ss}$ will sweep out part of X_E . The union of such trajectories for each feasible control $\underline{u} \in U$ is X_E , so any change in U will directly affect X_E .

Capture space:- This is constructed using a constant control $\underline{u}_{ss}$, so X_C is independent of the feasible control space U , but will depend on

the asymptotic values of the solutions of the system equations from different initial states. These will be determined by the topology of the n -dimensional state space, which identifies such features as limit cycles and multiple steady states. These are well understood for 1- and 2-dimensional systems, see Lefschetz p170 et seq, but unfortunately the theory has not been developed for higher dimensions.

Two further theorems can now be stated and proved.

2.2.2 Theorem B (Commutable switches)

A system which can switch reversibly between a reference state and all members of a set of steady states, can also switch reversibly between any members of this set.

Proof:-

If necessary, the switch may be accomplished by switching to the reference state and then away from it. In the framework of theorem A, the reference state becomes part of the sequence of intermediate steady states.

This theorem is useful for verifying that a system can switch between operating states if it can be started up and shut down from each state.

2.2.3 Theorem C (Global stability)

A system which is globally asymptotically stable at a steady state ($\underline{x}_{ss}$) is always switchable to that state.

Proof:-

$X_C = R^n$ by definition of global asymptotic stability. Therefore

$$X_E \cap X_C = X_E \neq 0 \text{ for all } X_E \in R^n \quad (2.10)$$

This satisfies the conditions of theorem A, so the switch is feasible.

Many classes of chemical processes are known to be stable - either by experience or through formal proof. Theorem C can show that within the range of validity of the describing equations, feasibility is not a problem. This fact is now used to show that heat integration will not cause infeasible switches in general if certain conditions are satisfied.

2.2.4 Theorem F (Thermally coupled networks)

A time-invariant network of splitters, mixers and heat exchangers is always thermally switchable if the following conditions are satisfied.

- A) All flowrates may be constrained to their constant target values.
- B) Enthalpy is only a function of temperature.
- C) Each element of the network is thermally homogeneous. (i.e. the temperature at all points within each element is the same)

Proof:

The proof consists of establishing the form of the describing equations and proving global stability by a Lyapunov method. Theorem C proves switchability of such a system.




Figure 2.9 Homogeneous unit within network

Consider a unit k within a network of n units which are thermally coupled to one another and to their environment by convection, conduction and radiation. The convective terms are represented by constant flowrates F_{ik} (condition A) so the holdup H_k will also be constant. The condition that each unit is thermally homogeneous (C) implies that the specific enthalpy of the product streams h_k is the same as the bulk specific enthalpy. The net radiative and conductive heat flux Q_{ik} is a function only of the relative temperature $T_k - T_i$. These flows between two typical network elements are shown in figure 2.9 above. Flows to the environment are represented by subscript 0. The enthalpy balance on unit k is (2.11),

$$H_k \frac{dh_k}{dt} = \sum_{i=0}^n [F_{ik} h_i + Q_{ik} - F_{ki} h_k] \quad (2.11)$$

The definition of the switch implies the existence of the target steady state $\underline{h}_{ss}$, so the equation can be rewritten in terms of the perturbation from this state, $\underline{x}$,

$$\begin{aligned} \dot{x}_k &= \sum_{i=0}^n \left[\frac{F_{ik}}{H_i} x_i + y_{ik}(x_i, x_k) - \frac{F_{ki}}{H_k} x_k \right] \\ &= f_k(x_1 \dots x_n) \end{aligned} \quad (2.12)$$

$$\text{where } x_i = H_i(h_i - h_{i,ss}) \quad (2.13)$$

$$y_{ik} = Q_{ik} - Q_{ik,ss} = -y_{ki} \quad (2.14)$$

The elements of the system jacobian are (2.15, 2.16),

$$\frac{\partial f_k}{\partial x_i} = \frac{F_{ik}}{H_i} + \frac{\partial y_{ik}}{\partial x_i}, \quad i \neq k \quad (2.15)$$

$$\frac{\partial f_k}{\partial x_k} = \sum_{i=0}^n \left[-\frac{F_{ki}}{H_k} + \frac{\partial y_{ik}}{\partial x_k} \right] \quad (2.16)$$

The flow terms are always positive by definition. The radiation and conduction terms can also be shown to be positive, as follows. The specific enthalpy is a function only of temperature (condition B) and from the thermodynamic definition of enthalpy the gradient of enthalpy versus temperature must be positive. Next, the rate of heat transfer between two media by conduction and radiation alone will always increase with the temperature difference. This is confirmed by experimental evidence. Thus the change of heat flux will always oppose the change of specific enthalpy of the media, and the gradient $\partial y_{ik}/\partial x_i$ will always be positive. We can now sum the terms in the columns of the jacobian matrix (2.17),

$$\begin{aligned} \sum_k \frac{\partial f_k}{\partial x_i} &= \sum_{\substack{k \neq 0 \\ k \neq i}} \left[\frac{F_{ik}}{H_i} + \frac{\partial y_{ik}}{\partial x_i} \right] \\ &\quad - \sum_{k \neq 0} \left[\frac{F_{ik}}{H_i} + \frac{\partial y_{ik}}{\partial x_i} \right] \\ &\quad - \left[\frac{F_{i0}}{H_i} + \frac{\partial y_{i0}}{\partial x_i} \right] \end{aligned} \quad (2.17)$$

The first summation, due to the off-diagonal terms (2.15) exactly cancels the diagonal terms (2.16) except for the heat flux leaving the network. Hence, the summation can either be zero or negative,

depending on whether the unit has a suitable connection to the environment (2.19) or not (2.18).

$$\sum_k \frac{\partial f_k}{\partial x_i} = 0, i=1,2,\dots,n \quad (2.18)$$

$$\sum_k \frac{\partial f_k}{\partial x_i} < 0, i=1,2,\dots,n \quad (2.19)$$

We now use a stability proof due to Rosenbrock (1963) which depends essentially on the following properties.

- (i) There is a flow of some conserved quantity between units (in this case, heat).
- (ii) The direction of flow reduces the driving force.
- (iii) The flow increases monotonically with driving force.

Rosenbrock demonstrated that binary CMO^{*}distillation columns and isolated heat exchangers possessed these properties. To extend the proof to heat integrated networks we must consider the associated rigorous mathematical conditions:

R.A) The partial derivatives $\partial f_k / \partial x_i$ exist and are continuous over all i and k . (Lipschitz condition)

R.B) For all i ,

$$\sum_k \frac{\partial f_k}{\partial x_i} + \psi_i = 0, \quad \psi_i \geq 0 \quad (2.20)$$

R.C) For all $i, i \neq k, \frac{\partial f_k}{\partial x_i} \geq 0 \quad (2.21)$

R.D) The equations are linked in such a way that every variable for which $\psi_i = 0$ appears in an equation which is dependent on a variable for which $\psi_i \neq 0$. (This condition is stated in detail by Rosenbrock).

* See p.129

Theorem (Rosenbrock)

If conditions R.A, R.B, R.C and R.D are satisfied within a closed space and the solution $\underline{x}$ never leaves this space, then the solution is asymptotically stable.

The proof of this theorem uses a Lyapunov function of the form;

$$L = \sum_i |f_i| \quad (2.22)$$

For the case of heat integrated networks, conditions R.A, R.B and R.C are verified by inspection of the equations and inequalities derived above. Condition R.D can be shown to be satisfied by considering the network itself. Since the network has at least one output, the variable describing it will have a $\Psi > 0$. As every sub-unit of the network can indirectly affect the outputs, there is the necessary structural link between the outputs and sub-units. The closed space for the solution vector is defined by the maximum and minimum temperatures present in the network at the initial state. The temperatures within the network at all subsequent times will be bounded by these since, as previously stated, the heat flows will always be in the direction which reduces the driving force. As the four conditions will be met at all temperatures, the system will be asymptotically stable.

Discussion:

Before the value of this theorem can be judged, its applicability to real processes must be considered. The conditions will be considered in order of importance in chemical engineering.

Condition B renders the theorem inapplicable to systems with chemical reactions involving the use or production of heat. Separation processes where enthalpy depends strongly on composition are also excluded. However, in many important processes, the enthalpy of a stream is virtually unaffected by stream properties other than temperature. Take for example the separation of hydrocarbon mixtures. The specific enthalpy of all saturated hydrocarbons from C_4 to C_{20} as vapour at 298K and 1 bar lies within 12% of the average (API, 1981). The enthalpy is thus almost independent of composition for the mixtures encountered in a distillation system.

Condition C will not normally be satisfied directly for common process equipment such as heat exchangers, which present a temperature gradient, but can be satisfied indirectly by using alternative models for the equipment. Consider for example a heat exchanger in which the process fluid is nominally in plug flow. As shown in chapter 5, a cascade of well-mixed tanks with suitable thermal coupling may be used to model such a unit. As the theorem places no limit on the interconnection or number of such tanks, the theorem is applicable to any plug flow heat exchanger. Similar reasoning may be possible for other inhomogeneous process units.

There are two aspects to condition A, relating to flows between physical units and within physical units. Flows between units will normally be easily constrained to the target values in single phase systems. Multiphase systems, like the flash drum example given earlier, can impose physical restrictions on the flows. As demonstrated, these may lead to infeasible switches.

Flows within physical units must also be constrained if the unit has been partitioned so the model meets condition C. Again, in single phase systems, the flow will normally distribute itself evenly within the unit, so the internal flows are simple ratios of the external flows. If the external flow can be regulated, then one can be reasonably confident that the internal flows are also regulated.

This reasoning is invalid for multiphase units, e.g. reboilers and condensers. The internal flows are determined by the geometry of the heat transfer surfaces and phase distribution within the units. This can change with temperature, thus violating condition A, and lead to abnormal temperature/heat transfer relationships. In physical terms, the heat transfer surface may 'boil dry', with consequent drop in heat transfer coefficient compared to nucleate film boiling. The work of Farber and Scoriah (1948) clearly illustrates this phenomenon. The theorem can only be applied to networks containing multiphase units when the scale of the phase distribution is small enough for the flow distribution assumption to hold. An example is a reboiler operating at relatively low temperature difference between boiling liquid and heated surface. Under these conditions the rapid circulation of liquid and small bubble size will ensure thermal homogeneity.

In conclusion, although there are several important restrictions, theorem F can be applied to a wide range of heat integration schemes. Exceptions will usually be recognised by the nature of the processes involved, as demonstrated in the examples. However the theorem only defines sufficient conditions for feasibility, so a system which fails to meet all of them is not necessarily unswitchable.

The theorem has an interesting corollary, namely:

For a process which satisfies the conditions A-C above, its ability to start-up and shut-down will be unaffected by any heat integration scheme.

2.3 PRACTICAL EXAMPLE

The flash drum problem (second example, section 2.1) will be used to illustrate the application of theorems A and F. As previously shown, the basic design has two stable steady states, S_1 (vapour fraction 0.02) and S_3 (vapour fraction unity), and it is not possible to switch from S_3 to S_1 .

Theorem A tells us that it is sufficient to find an intermediate state which lies within the escape space of S_3 and the capture space of S_1 . Firstly, we must identify the capture space of S_1 . Since we have a single state variable in the system, the capture space will be determined by the sign of the eigenvalue of the system in the vicinity of each steady state. The intermediate (unstable) steady state S_2 has a positive eigenvalue so the state will always move away from it, towards either S_1 or S_3 . The capture space of S_1 therefore comprises all flash feed temperatures below S_2 .

Next, the escape space of S_3 must be examined for a steady state which lies in the capture space of S_1 . The escape space of S_3 is virtually nil, because the cooling water exchanger has very little flow in it on the process side. Although the process stream temperature may be controlled, it has negligible effect on the flash feed temperature.

It is clear that the process must be modified to move the state from S_3 to some intermediate state below S_2 . Theorem F suggests that if all flowrates were constrained to some intermediate steady state value, the system could always be switched to this intermediate steady state. The third step is to find a means to regulate the split fraction independently of the temperature into the flash drum, since this is the point at which the flowrates are not controlled. One method is to fit a stop-valve in the vapour line from the flash drum. When this valve is closed, all feed will be directed through the cooling exchanger. A suitable intermediate steady state, which is known to be stable from theorem F, can be found by a simple heat balance.

The final strategy for the switch now becomes;

- (i) Close vapour line stop valve.
- (ii) Set intermediate cooling water rate.
- (iii) Wait until flash feed temperature falls below S_2 .
- (iv) Open vapour line stop valve.
- (v) Set target cooling water rate.
- (vi) Flash feed temperature will now approach S_1 .

Although this may not be optimal, the process is now suitable for further development using an optimal control package. The switch feasibility has been assured for the price of a single low-cost valve.

2.4 SUMMARY

Feasibility of switching is undoubtedly an important issue in certain types of process, but it rarely presents an insurmountable problem. Where feasibility is known to cause difficulties, the engineering profession has devoted special efforts to solving it, as typified by the extensive research on stability of exothermic reactors (Van Heerden; Bilous and Amundsen; and others).

For heat integration in separation processes, only cryogenic systems are likely to show switchability problems, because of the nature of the prime source of energy. In normal petrochemical separations, it has been proven that if certain assumptions can be made about the enthalpy and heat transfer mechanisms of the streams, switches will always be feasible. In other fields, where these assumptions are invalid, theorems A B and C may be used to prove switchability, or suggest simple process modifications to assure it.

CHAPTER THREE

MEASURING SWITCHABILITY

The previous chapter has shown that feasibility of switches is not usually an important practical consideration in the design of heat integrated distillation systems. We should therefore consider some quantity which is representative of the switchability of each alternative design. In this chapter the dynamic performance indices reported by other workers in related areas are classified and discussed. An index of switchability suitable for this work is then developed and qualitatively justified.

3.1 CLASSIFICATION OF DYNAMIC PERFORMANCE INDICES

All indices of dynamic performance of chemical process may be regarded as approximations to a rigorous economic index. Such an index would be constructed by considering the costs of feedstock and energy, the value of the product streams, depreciation and maintenance of the plant, and so forth. These terms would be summed over the period of interest, and the best system or strategy would be the one maximising the resulting net profit. While this is unquestionably the most attractive index from a commercial standpoint, it is very inconvenient for engineering purposes, not least because much of the data is unobtainable.

Although dynamic performance can be estimated from transfer functions of linearised models (Morari, 1982; Wong, 1984) as mentioned in chapter 1, such estimates have little basis on economic grounds, and will not be considered further. The essential feature of the economic index is that it considers the behaviour of the system at every point along a switching state trajectory, and an index of switchability should preserve this.

The dynamic performance indices, or objective functions, considered for switchability are all integrals of some function of the system state and control levels which represents an 'instantaneous cost'. There are three terms in this function, each contributing some approximation to the 'instantaneous cost':

- (i) Direct profit estimate,
- (ii) Offspec product penalty,
- (iii) Switch duration.

Each of these terms is considered in detail in subsequent sections.

The overall objective function is then given by (3.1) - (3.3) below;

$$J = \int_{t_1}^{t_2} L(\underline{x}, \underline{u}, t) \cdot dt \quad (3.1)$$

$$\text{subject to } \underline{x} = \underline{x}_1, \text{ for } t < t_1 \quad (3.2)$$

$$\underline{x} = \underline{x}_2, \text{ for } t > t_2 \quad (3.3)$$

where J is the objective function of the switch from $\underline{x}_1$ to $\underline{x}_2$ and L is the instantaneous cost function obtained by summation of the component terms (i) to (iii) above. The construction of a suitable function L for switchability is the subject of section 3.6.

It is essential for this work that the value of L at the target steady state, $L(\underline{x}_2, \underline{u}_2)$, is zero. This is necessary because of the way a switch has been defined, see chapter 1. A switch is a transition between two steady operating states. If these are coincident, namely we are considering a switch from one steady state to itself, the cost of the switch must be zero. The integral (3.1) nevertheless defines the cost for any arbitrary interval t_1 to t_2 . The only meaningful value of L which satisfies this requirement is $L(\underline{x}_2, \underline{u}_2) = 0$. By similar reasoning, the instantaneous cost L of any steady operating state must also be zero. This is referred to as the zero-steady-state-cost condition in subsequent sections.

3.2 DIRECT PROFIT ESTIMATE AND MARGINAL VALUES

The first of three types of instantaneous cost function is the direct profit estimate, which is implemented using the marginal values associated with the product specifications. It seeks to include an approximation to every term in the rigorous economic cost formula. These terms can represent product value, feed cost, energy cost and the change of inventory within the system. Depreciation and maintenance costs are not usually dependent on operating conditions so are best regarded as time proportional terms, see section 3.4.

For obvious commercial reasons the system will be operated at the control settings which maximise the overall profit. Since we are considering switches between steady states (see chapter 1) the steady state controls must maximise the instantaneous cost function. A rigorous mathematical model of the system will by definition exhibit a profit maximum at the operating state which coincides with the desired initial or target state. If the rigorous economic cost formula is linearised about this state it may be used directly as a direct profit estimate, whose optimum will coincide with the true optimum.

In practice the mathematical model will not be exact, but will contain many simplifications to facilitate efficient computation, see chapters 5 and 6. These simplifications will introduce small but significant errors in the steady state solutions, so that even if the linearised rigorous costs were known, the optimum operating states may not coincide with the desired end states of the switch. A cost function must be found which can compensate for both the lack of data and the model error.

The practical solution is to assume that the end states are optimal with respect to all control variables. It is then possible to write

$$\frac{dL}{du_i} = 0, \quad i = 1 \dots r \quad (3.6)$$

Now because the system is at the steady state, the singular solution of the dynamic equations will uniquely determine all states from the controls, and also the gradients of each state with respect to each control. Hence

$$\frac{dL}{du_i} = \frac{\partial L}{\partial u_i} + \sum_j \frac{\partial L}{\partial x_j} \frac{dx_j}{du_i} = 0 \quad (3.7)$$

The partial derivatives $\partial L/\partial x_j$ are the marginal values of each state, while the partial derivatives $\partial L/\partial u_i$ are the marginal values of each control. Some of these will be known - for example the value of on-spec product or the cost of fuel oil. Since there are r equations in 3.7, we may solve 3.7 for r unknown marginal values. The choice of suitable marginal values will often be a matter for engineering judgment, because the number of unknown marginal values will frequently exceed r . Where product purity is of interest, a logical choice is the composition of key component in each stream.

Taking the first-order terms of a Taylor-series expansion of the instantaneous cost gives the following direct profit estimate, equation 3.8,

$$L(\underline{x}, \underline{u}, t) = L(\underline{x}_{ss}, \underline{u}_{ss}) + \underline{m}^T \begin{bmatrix} \underline{x} - \underline{x}_{ss} \\ \underline{u} - \underline{u}_{ss} \end{bmatrix} \quad (3.8)$$

where $\underline{x}_{ss}$, $\underline{u}_{ss}$ pertain to the desired optimum state, and $\underline{m}$ is the vector of marginal values. Undefined terms of $\underline{m}$ must be set to zero.

This approach is used by Sullivan (1978) for a splitter column with three controls and three product streams. Numerical results suggest that this is a perfectly suitable method for distillation studies, but the main purpose of his investigation is to justify the next method to be considered, namely the minimal off-specification cost function.

There are several drawbacks to the marginal cost formulation. Although the need for extensive economic data is avoided, some cost information is still required. Secondly the choice of unknowns to solve 3.7 for is not necessarily obvious. Finally the marginal values are only a linear subset of the true economic cost function, so will only be strictly accurate near the reference state.

3.3 QUADRATIC - OFF SPECIFICATION COST FUNCTION

This cost function penalises any deviation from the initial or target state according to the square of the deviation. The general formula is

$$L = \frac{1}{2} (\underline{x} - \underline{x}_{ss})^T Q (\underline{x} - \underline{x}_{ss}) \quad (3.9)$$

where Q is a positive definite matrix.

In practice, Q is either the unit matrix or has terms on the leading diagonal only, because of the difficulty in estimating the off-diagonal terms.

The principal advantage of this function is that the target state automatically minimises the cost L . Although it does not have the same theoretical basis as the marginal value function it nevertheless has a satisfactory qualitative interpretation. Consider a binary column which requires a certain heat input to achieve a steady state optimal separation of a valuable light product from a waste heavy product. A rise in purity of light product will use more energy but at reduced efficiency, so will increase the overall operating cost. A fall in purity to save energy will cause more light product to be lost in the waste, thus also increasing overall operating cost. At the optimum these two effects are minimised (by definition), and this coincides with the quadratic minimum.

Sullivan demonstrated convincingly for his splitter example (three component separation) that the optimal control policy obtained from the quadratic function gives virtually the same cost as the marginal value profit function. The former also possesses superior numerical properties for optimisation algorithms, with a much sharper minimum. For his work, the reference setpoint was chosen to be the one which minimised L . This was based on selecting a product storage tank according to the instantaneous product quality.

The quadratic cost function can have an optimal linear state feedback solution when applied to a linear system. The problem must be defined as;

$$\begin{aligned} \text{minimise } J = & \frac{1}{2} \underline{x}^T(t_f) S_f \underline{x}(t_f) \\ & + \int_{t_0}^{t_f} \left(\frac{1}{2} \underline{x}^T Q \underline{x} + \frac{1}{2} \underline{u}^T R \underline{u} \right) dt \end{aligned} \quad (3.10)$$

$$\text{subject to } \dot{\underline{x}}(t) = A \underline{x}(t) + B \underline{u}(t) \quad (3.11)$$

$$\underline{x}(t_0) = \underline{x}_0 \quad (3.12)$$

The feed-back law is derived and discussed in e.g. Bryson and Ho (1975), and is given by

$$\underline{u}(t) = -R^{-1} B^T S(t) \underline{x}(t) \quad (3.13)$$

where $S(t)$ is the positive definite solution of the matrix Riccati equation

$$\dot{S} = Q + A^T S + S A - S B R^{-1} B^T S \quad (3.14)$$

$$S(t_f) = S_f \quad (3.15)$$

Often, the final time horizon is extended to infinity. The problem is still meaningful because the state error will tend to zero. In this case $S_f = 0$ and the stationary value of S , given by the solution of

$$0 = Q + A^T S + SA - SBR^{-1}B^T S, \quad (3.16)$$

can be used in 3.13 above.

A variant on this is used by Lenhoff and Morari (1982) as the dynamic performance index in a study of heat-integrated distillation columns. The results obtained are compatible with those of transfer function analyses of similar systems. However the choice of matrices Q and R , and initial perturbation from the steady state, $\underline{x}_0$, appear to have little theoretical basis.

It is tempting to use this linear state feedback solution to obtain an index of switchability directly. This is not recommended for two reasons. Firstly, the solution is only valid for linear systems. The well-known nonlinear behaviour of distillation columns over the wide operating range traversed during a switch would invalidate this. Heat exchangers also exhibit nonlinear properties. Secondly, the control deviation penalty matrix R has no meaning in many chemical engineering problems, since control action in itself often incurs no cost. We cannot simply ignore it (treat R as zero) because R^{-1} appears in (3.13 - 16).

3.4 SWITCH DURATION

Here the time taken by the system to transfer from the initial to the final state is regarded as the switch cost. The cost function L is simply unity. Unlike the previous cost function components, this is only meaningful if the initial and final states are achieved within the allotted switching period. For compatibility with the previous definitions of L we can write

$$L(\underline{x}, \underline{u}, t) = 1 \text{ if } \underline{x} \neq \underline{x}_1$$

$$\text{and } \underline{x} \neq \underline{x}_2$$

$$L(\underline{x}, \underline{u}, t) = 0 \text{ otherwise} \tag{3.17}$$

The justification for this cost function rests on the fate of off-specification product in many industrial separation processes. Where the off-spec products are recycled, the value of such products may be assumed equal to the value of the feestock. Let us suppose that normal production is only possible at either target state. During the switch, the process will be consuming energy as normal, but all products are being recycled, so it will not make anything of value. Furthermore, if the overall energy requirement is constant then the switch cost, given by the total energy consumption, will be proportional to switch duration.

A practical application of this cost function need only consider the system behaviour while $L = 1$. This is a direct consequence of the definition of state variables;

"The future behaviour of a system is fully defined by the current values of the state variables and the future values of the control variables".

Once the sytem has reached the final state, any control action other than final steady state control will incur a penalty, since it must

cause the state to deviate from the final value. By a similar argument, the start of the switch can coincide with the beginning of the allotted switching period. The problem now becomes a two-point boundary value problem, because the initial and final states are constrained. Such problems are difficult to solve numerically, and this has restricted the application of time-optimal objective functions in chemical engineering.

One method of avoiding these difficulties is to approximate the final time constraint with a penalty function which is quadratic in the final states. This is used successfully by Pollard (1967), reviewed in chapter 1. The switch duration index appears to have slight advantages over the economic index, although Pollard does not make any direct comparison. Sullivan (1978) proposes a more elaborate formulation which actually treats the final state error as an equality constraint, but does not appear to have used it himself.

The switch duration cost function can be applied strictly only to systems which are completely controllable to the target state. This is because the switch is not complete until $\underline{x}_2$ is attained exactly, equation (3.17). The conditions for the feasibility of a switch given in chapter 2 are insufficient because these only ensure that the system will tend asymptotically towards the target state.

A system is completely controllable to a given state if there exists a feasible control policy which will transfer the system to this state from the initial state in a finite time. Since the system will always move into a region which is close to the target state (asymptotic stability) we need only consider the behaviour within a small region surrounding the target state. Within this region a linearised model of the system is often appropriate.

A linear system described by equations (3.11-12) is completely controllable if the rank of the controllability matrix G is equal to the dimension of the state space, n , where

$$G = [B | AB | A^2B \dots A^{n-1}B] \quad (3.18)$$

The proof of this can be found along with a thorough discussion of controllability in non-linear and time-varying systems in Athans and Falb (1966) p.197 et seq.

This does not mean that incompletely controllable systems cannot be analysed by switch duration comparisons, but due allowance must be made for the subspace of the uncontrollable states. The use of quadratic end-point penalties which effectively relax the condition (3.17) on $\underline{x}_2$ is possible. Another method is to introduce dummy controls orthogonal to the span of G . These will only affect the system in the previously uncontrollable state subspace. Finally, if the uncontrollable subspace only contains isolated states, these states may be omitted from the equation (3.17) entirely.

3.5 SWITCHABILITY AND OPTIMAL CONTROL

So far we have discussed methods of measuring the cost of a switching policy without defining the control action required to achieve the switch. As mentioned in chapter 1, the switchability of the process is the optimal cost - that is, the lowest cost out of every control policy which achieves the switch.

It is possible to define the switching policy by one method and measure the cost by another, for example apply closed loop controls and use a quadratic penalty on the product specification. This is not considered important for this work, because it contradicts the definition given above. If the switching cost has not been minimised, the control policy is incorrect.

Lenhoff and Morari (1982) consider only the use of linear state feedback control, which is optimised as outlined in section 3.3. They argue that because a real system will be operated with feedback controllers, these must be incorporated in any model for dynamic performance estimation. Although this is compatible with the requirement for optimality, the use of such simple controllers severely constrains the allowable control action. The authors make no mention of how their proposed SISO proportional-only controllers would be developed into a practical multivariable 2 or 3 term control scheme.

The use of feedback controllers to define the switching trajectory is considered most unrealistic for switchability studies, and will only be used when there are over-riding numerical reasons. Apart from introducing extra equations into the system model, the choice of control system will invariably affect overall performance. The optimal control schemes used by Pollard (1967) and Sullivan (1978) avoid the need to introduce controllers in the network model. The optimal control package itself becomes the controller, and the objective function will be minimised when the process variables are regulated near their target values.

3.6 COST FUNCTIONS FOR SWITCHABILITY

It is clear from the foregoing review that although the optimal open-loop switching trajectory is an essential ingredient of a general switchability study, each of the cost functions has its own advantages and disadvantages.

The marginal value cost function performs poorly in numerical optimal control packages but it alone measures the cost of energy to accomplish the switch. Any switchability study on energy integration should consider this cost.

The off-spec product cost function cannot strictly be equated to the true economic cost, although Sullivan (1978) shows that it can give very similar results.

The minimum time cost function appears to fall between the other two. By suitable formulation it will guarantee that the initial and target state constraints are met within well-defined limits.

The ideal cost function would combine the best features mentioned above without compromise. The following formula is proposed;

- (i) A term to represent the direct cost of control action. This will penalise switching policies which sink large amounts of energy through a system to achieve rapid state variations.
- (ii) A term to account for off-specification product. This will encourage the target product setpoints to be established rapidly.
- (iii) A term to represent final state error when the switch control action is complete.

Terms (ii) and (iii) will both be quadratic. Since we are mainly interested in the states which affect the product streams, it is sufficient to use the same term for both. During the switch it will represent offspec product, and afterwards it will represent final state error. This assumes that the system is asymptotically stable, which will be true for the systems considered (see chapter 2). This method will also handle incompletely controllable systems.

Hence,

$$J = \int_{t_1}^{t_2} L(\underline{x}, \underline{u}_{opt}) dt + \int_{t_2}^{t_f} L(\underline{x}, \underline{u}_2) dt \quad (3.19)$$

where

$$L(\underline{x}, \underline{u}) = (\underline{x} - \underline{x}_2)^T Q (\underline{x} - \underline{x}_2) + \underline{m}^T \left[\frac{\underline{x} - \underline{x}_2}{\underline{u} - \underline{u}_2} \right] \quad (3.20)$$

The first integral will be the penalty for offspec product and the operating costs during the switch, and the second integral indicates how close we get to the target state $\underline{x}_2$. There is no need to integrate to infinite time to evaluate this. Two or three system time constants will give sufficient accuracy for practical purposes, as shown in the following example.

For a one variable system, let

$$\dot{x} = -x \tag{3.21}$$

$$x(t_2) = 1 \tag{3.22}$$

$$J(t_F) = \int_{t_2}^{t_F} x^2 \, dt \tag{3.23}$$

then

$$J(t_F) = \frac{1}{2} (1 - \exp 2(t_2 - t_F)) \tag{3.24}$$

Table 3.1 below shows that the terminal error converges very rapidly to the true error as the time horizon t_F increases.

Time $t_F - t_2$	Approach, $J(t_F)/J(\infty)$
0	0
1	0.86
2	0.98
3	0.998
∞	1

Table 3.1 End point error of a 1st order system

It only remains to decide on the quadratic weighting matrix Q . If this is too small the optimiser will ignore the setpoint constraints completely and simply minimise operating costs regardless of product quality. The opposite situation will give a policy which may be highly uneconomical because it endeavours to hold the product specification throughout the switch. To maintain the balance, the effect of a cutback in operating cost on steady state product quality must be computed. The following method is employed:

- (i) Decide on an acceptable setpoint deviation Δx_i for each product stream specification, x_i .
- (ii) Starting with the steady state control $\underline{u}_{ss}$, compute the product specification for progressively lower operating costs, L .
- (iii) Interpolate the results of step (ii) to find the steady state control policy $(\underline{u}_{ss} - \Delta \underline{u})$ for which

$$\|\underline{x}(\underline{u}_{ss}) - \underline{x}(\underline{u}_{ss} - \Delta \underline{u})\| = \|\Delta \underline{x}\| \quad (3.25)$$

where $\|\underline{x}\|$ is the Euclidean norm $(\sum x_i^2)^{\frac{1}{2}}$

- (iv) Select Q as the diagonal matrix, with equal elements q where a variable is a product stream, or 0 otherwise such that

$$\frac{\Delta L}{\Delta u} + \Delta \underline{x}^T Q \Delta \underline{x} = 0 \quad (3.26)$$

This ensures that the savings in operating cost are just balanced by the loss of product quality at the bottom end of the specification range. Although it is not suggested that the resulting Q will be optimal it will be least be of the correct magnitude.

CHAPTER FOUR

MATHEMATICAL TECHNIQUES

In the previous chapter the optimal switch cost was shown to be a potentially useful index of switchability to compare different heat integration schemes. Finding the optimal cost is in general a far from trivial task. In this chapter the development of an optimal control program suitable for this work is described. Starting with the fundamental theories of optimal control a number of general techniques for solving the optimal control problem are reviewed, and the use of a gradient method is justified. Before a practical implementation can be produced, the problem must be redefined with discrete control intervals. Techniques for the numerical solution of differential equation sets and constrained minimisation are also discussed in context.

The resulting algorithm has been linked to a general chemical engineering flowsheeting package called SPEEDUP so that different networks can be studied without lengthy preparation. SPEEDUP is used to define the initial and target states for each network, and the intermediate files containing the detailed problem description are used directly by SPEEVOP. This is a comprehensive evolutionary optimal control package written specially for this work incorporating the optimal control program described herein. The operation of the overall package is verified with two simple test problems described at the end of the chapter.

4.1 DEFINING THE OPTIMAL CONTROL PROBLEM

Before solution methods can be discussed the problem must be formally defined.

Consider a system described by the following equations

$$\dot{\underline{x}}(t) + f(\underline{x}(t), \underline{u}(t), \underline{v}) = 0 \quad (4.1)$$

$$\underline{x}(t_0) = \underline{x}_0 \quad (4.2)$$

where t is the independent variable, usually time,

$\underline{x}(t) \in \mathbb{R}^n$ is the vector of dependent variables,

$\underline{u}(t) \in \mathbb{R}^r$ is the vector of control variables,

$\underline{v} \in \mathbb{R}^s$ is the vector of design variables.

The objective function which is to be minimised is

$$J = J(\underline{x}(t_f), \underline{v}) \quad (4.3)$$

and the entire system is constrained by

$$u_i^{\min} \leq u_i(t) \leq u_i^{\max} \quad (4.4a)$$

$$v_i^{\min} \leq v_i \leq v_i^{\max} \quad (4.4b)$$

$$g_i^{\min} \leq g_i(\underline{u}(t), \underline{v}) \leq g_i^{\max} \quad (4.5)$$

where $\underline{g} \in \mathbb{R}^q$ is the vector of non linear control constraints. The objective function in (4.3) is compatible with equation (3.19) if an extra variable x^0 is used to represent the integral term. Hence

$$J = x^0(t_f), \quad 0 = x^0(t_0) \quad (4.6)$$

$$\dot{x}^0 = L(\underline{x}(t), \underline{u}(t), \underline{v}) \quad (4.7)$$

The optimal control problem is to choose an admissible set of control and design variables to minimise the objective function, J , subject to conditions 4.1 - 4.7 where applicable. The definition may be extended to cover path constraints, end-point constraints and variable final time t_f , but these extensions will not be required for this work.

Matrix A in (4.1) need not be of full-rank. If it is the negative unit matrix we obtain the classical formulation used until very recently in all publications on optimal control. A more flexible approach allows the system to contain algebraic and differential equations by setting the appropriate rows of A to zero. It is often convenient to solve a mixed system simultaneously, rather than solving the algebraic subsystem as part of the calculation procedure to evaluate the state derivatives with respect to time.

4.2 SOLVING THE OPTIMAL CONTROL PROBLEM

The solution of the optimal control problem as defined in the previous section requires sophisticated mathematical techniques. A fundamental feature of these is the adjoint system, defined as follows;

$$\dot{\underline{p}}^T(t)A - \underline{p}^T(t) f_x(t) = 0 \quad (4.8)$$

$$\underline{p}^T(t_f)A - J_x^T(t_f) = 0 \quad (4.9)$$

where $\underline{p} \in \mathbb{R}^n$ is the vector of adjoint variables. These expressions are a simplified form of equations derived by Morison (1984) using a variational approach. Many properties of the adjoint system for the classical differential systems of equations also apply to the system with algebraic equations and variables. Since this extension of the theory to mixed systems (of differential and algebraic equations) is quite recent, the references in the remainder of this section only consider pure differential systems.

The development of the necessary conditions at the solution of the problem is well known. It can be shown, (Bryson and Ho (1975)) that the gradient of the objective to the instantaneous value of the controls is;

$$J_u = \underline{p}^T f_u \quad (4.10)$$

thus in the absence of any constraints on $\underline{u}$, along the optimal solution;

$$0 = \underline{p}^T f_u \quad (4.11)$$

The optimal unconstrained control problem is then defined by the three sets of equations (4.1-2), (4.8-9) and (4.11). Control constraints and singular arcs, where 4.11 does not define $\underline{u}$, will modify the equation set, but the overall equation structure remains similar.

Bryson and Ho (1975) propose three classifications for the various numerical methods used to solve the two-point boundary value problem contained in these equations. These are,

- (i) Neighbouring extremals (shooting methods).
- (ii) Gradient methods.
- (iii) Quasilinearisation methods.

The methods differ only in the order in which the subsystems of equations are solved.

Method 1 solves the full set of equations simultaneously from boundary condition (4.2). It attempts to 'hit' the adjoint boundary condition (4.9) by trial and error in the original adjoint estimates at t_0 . The advantage of this method is that every trajectory is optimal. Unfortunately, because of the minus sign in (4.8) the adjoint system will almost always be unstable, so the final value of $\underline{p}$ will be very sensitive to the initial value. This will hinder convergence to $\underline{p}(t_f)$.

Method 2 solves (4.1-2) to obtain the state trajectory, and (4.8-9) for the corresponding adjoint trajectory. (4.10) is then used to adjust the controls to minimise J through a conventional minimisation package. The adjoint system is solved from t_f backwards so the instability problem is avoided. The success of this method rests largely on the minimisation package, which may be first or second order.

Method 3 estimates a state and adjoint trajectory, then solves (4.11) to find the control policy. The state and adjoint trajectories are then modified through (4.1-2) and (4.8-9) and the process repeated until the estimate is suitably close to the solution trajectory. As with method 1, it is somewhat sensitive to the initial guess.

For the large constrained systems encountered in chemical engineering problems the gradient methods (2) have gained universal acceptance.

Sullivan (1978) reviews the development of gradient methods and quotes fifteen favourable reports since 1963, in addition to his own very successful experience. They are preferred because they will handle control constraints naturally within the minimisation package, and are relatively robust. However, these advantages are easily lost by careless implementation or a poor choice of minimisation package.

The minimisation sub-problem can be stated as follows;

choose $\underline{y} \in \mathbb{R}^m$ to minimise $J(\underline{y})$

$$\text{subject to } y_i^{\min} \leq y_i \leq y_i^{\max}, i=1, \dots, m \quad (4.12)$$

$$g_i^{\min} \leq g_i(\underline{y}) \leq g_i^{\max}, i=1, \dots, m \quad (4.13)$$

In practice, the space of the decision variables $\underline{y}$ has a finite dimension m , whereas the control policy should be optimised over all time between t_0 and t_f . The most direct strategy is to discretise the interval t_0 and t_f into a finite number k of sub-intervals for each control within $\underline{u}(t)$. The start time of each sub-interval, and the value of the control during the sub-interval, form a pair of decision variables within $\underline{y}$. As the problem is defined in section 4.1, this will give the dimension of $\underline{y}$ as

$$m = r(2k - 1) + s \quad (4.14)$$

A problem with 8 subintervals of 4 control variables will require 60 decision variables. This is a considerable number even for the packages available. Sullivan adopted an effective compromise which uses a common set of sub-intervals for all controls. This only needs

$$m = (r + 1)k - 1 + s \quad (4.15)$$

The revised requirement is 39 decision variables, which saves typically 58 percent of the original computer storage, since the workspace is related to the square of the problem size for a second

order minimisation method.

The decision variable vector $\underline{y}$ now appears thus:

$$\underline{y} = (\underline{u}_1, \underline{u}_2, \dots, \underline{u}_k, t_2, \dots, t_k, \underline{v})$$

To solve the minimisation problem efficiently, the gradients of the objective and the constraints with respect to this vector must be computed. Morison (1984) has derived the following equations for the mixed differential-algebraic system.

$$\frac{dJ}{du_1(t_j)} = J_{u_1} - \int_{t_{j-1}}^{t_j} \underline{p}^T(t) f_u(t) dt \quad (4.16)$$

$$\frac{dJ}{dt_j} = \underline{p}^T(t) A (\underline{\dot{x}}(t_{j+}) - \underline{\dot{x}}(t_{j-})) \quad (4.17)$$

$$\frac{dJ}{dv_1} = J_{v_1} - \int_{t_0}^{t_f} \underline{p}^T(t) f_v(t) . dt \quad (4.18)$$

The theory on which the minimisation algorithm is based requires the gradients to be continuous in $\underline{y}$. This will be met if the derivatives of J and f are continuously differentiable.

Constraint information falls into two groups; linear, which applies to the bounds on decision variables 4.12, and nonlinear, which covers the remaining inequalities 4.13. The linear constraints are trivial. The minimisation package will customarily only handle simple inequality constraints, so each control constraint becomes;

$$c_{2(j+(i-1)k)-1} = u_{1j} - u_1^{\min} \geq 0 \quad (4.19)$$

$$c_{2(j+(i-1)k)} = u_1^{\max} - u_{1j} \geq 0 \quad (4.20)$$

The subinterval start times will also be constrained so that the subintervals retain finite duration.

$$c_{j+2kr} = t_j - t_{j-1} > 0 \quad (4.21)$$

Design variable constraints are similar to the control constraints.

$$c_{2(kr+i)+k-1} = v_i - v_i^{\min} \geq 0 \quad (4.22)$$

$$c_{2(kr+i)+k} = v_i^{\max} - v_i \geq 0 \quad (4.23)$$

The gradients for all linear constraints are simply +1 with respect to the appropriate decision variables.

The search for methods which will effectively handle general path constraints remains an important research topic. No attempt will be made here to incorporate such methods in the optimal control program. Instead, all path constraints will be treated as functions of the decision variables alone. The constraints of (4.5) then map directly into (4.13).

One drawback with the above is that the constraint values need to be determined explicitly. This is equivalent to the classical ordinary differential equation problem formulation which requires all algebraic equations to be eliminated from the system. It would be consistent with the present mixed d.a.e. formulation if the constraints were specified as part of the main equation set. This may also be more efficient since the constrained values will often appear naturally in the system. An example is the reflux flow in a distillation column, which must be constrained to positive values.

Let us augment the state vector with the constraint vector g, so x $\in \mathbb{R}^{n+q}$ in (4.2). Then the constraint values seen by the minimisation package are

$$c_{2(kr+s+j+(i-1)k)+k-1} = x_{n+i}(t_j) - x_{n+i}^{\min} \geq 0 \quad (4.24)$$

$$C_{2(kr+s+j+(i-1)k)+k} = x_{n+1}^{\max} - x_{n+i}(t_j) \geq 0 \quad (4.25)$$

The gradients with respect to the controls may be found by differentiating equation 4.1 then (since we are only interested in the time - independent variables) ignoring $\dot{x}$ and solving for the control derivatives at any point within the sub-interval j . Hence;

$$\frac{\partial x}{\partial u}_{t_j} = - (f_x)^{-1} f_u \quad (4.26)$$

$$\frac{\partial x}{\partial v}_{t_j} = - (f_x)^{-1} f_v \quad (4.27)$$

from which the gradients of (4.24, 4.25) may be extracted.

The conceptual solution algorithm can now be stated.

1. Define an initial piecewise constant control policy and series of switching times.
2. Solve equations (4.1-2) from t_0 to t_f by forward integration.
3. Solve equations (4.8-9) from t_f to t_0 by reverse integration along the path obtained at step 2.
4. Solve equations (4.16-18) to obtain the control and switch time gradients.
5. Compute constraint magnitudes and gradients from (4.19-4.27).
6. Pack gradient and constraint information into a suitable form and call the minimisation package.
7. If the minimum has been found, stop.
8. Unpack the updated decision vector into the control policy and switch times. Repeat from step 2.

The practical implementation will differ slightly to maximise computational efficiency (section 4.8).

4.3 INTEGRATION METHODS FOR MIXED SETS EQUATIONS

Steps 2 and 3 of the conceptual algorithm require the solution of mixed sets of differential and algebraic equations. This solution process, commonly called 'integration', is in practice the most time consuming part of the algorithm, so the solution method should be carefully chosen. The systems of equations encountered in chemical engineering are frequently both sparse and stiff. These terms and their significance will now be discussed.

The sparsity of an equation set normally refers to the density of non-identically-zero elements in the jacobian matrix. Sparse matrix methods exist which can solve such equation sets without using appreciably more computer storage than is required to store the original matrix. Because they operate on a much smaller set of data, they are generally faster than the equivalent full matrix implementation. The integrator should use a sparse method where matrix operations are required.

The stiffness ratio of a set of ordinary differential equations is the ratio of the maximum to minimum magnitude eigenvalues of the jacobian matrix. For a mixed system of differential and algebraic equations one must consider the eigenvalues of the equivalent canonical form, see Sinovec et al (1981). Alternatively the algebraic subsystem may be eliminated, see Wong (1984). Stiffness ratios of over 100 are not unusual in chemical engineering, and are widely recognised as a source of difficulty in solving differential equations. The stiffness of around 10^8 in the full distillation column model of chapter 6 was the main reason that a full dynamic simulation was omitted. The integrator should be able to handle moderately stiff systems, with stiffness up to 1000, such as may be constructed from the models proposed in chapters 5 and 6.

Cameron (1981) investigated the various methods available for solving this class of system, and gives an in-depth discussion of the above topics. The main points of this are summarized below.

The methods may be classified as explicit or implicit. An explicit method computes the solution at a given time by considering only the past behaviour of the solution. Although capable of high accuracy, they are generally unsatisfactory for stiff nonlinear problems because small errors in the solution are not corrected. Instead, the solution will diverge from the true solution with time, and the method is said to be unstable. To minimise this error a small time step must be used, and this is computationally inefficient. An implicit method expresses the solution at a given time as a set of algebraic equations which incorporate the time derivative information. This set can be solved using standard iterative methods, so that any desired accuracy may be obtained, subject to limitations in the formulation of the time derivatives. Stable solutions are thus attainable with relatively large time steps. An additional advantage for mixed differential and algebraic systems is that the algebraic subset of the system will always be consistent.

Another (compatible) classification is the order of the method. An n -th order method is defined as one in which the Taylor expansion of the solution up to and including the term in h^n is given exactly. An n -th order method must use at least n points to solve each time step h . High order methods can achieve a given accuracy with longer time steps than a low order method, and the advantage greatly outweighs the extra computation involved. Unfortunately the stability properties deteriorate as the order increases, so in practice orders greater than 5 are rarely used. This restriction applies to both implicit and explicit methods.

Sullivan (1978) used a single step implicit method with a fixed step length to solve his systems of ordinary differential equations for distillation column models. Two forms, the trapezoidal rule and the Liniger-Willoughby rule, were applied with remarkable success. A

third form, the backward Euler method, was not tried. The solution of a set of ordinary differential equations is defined through:-

$$\underline{x}(t + h) = \underline{x}(t) + h((1 - e) \underline{\dot{x}}(t + h) + e \underline{\dot{x}}(t)) \quad (4.28)$$

where h is the time step length and e is a parameter between 0 and 1, dependent on the method (Table 4.1).

Method	e	order
Trapezoidal	0.5	2
Liniger Willoughby	variable	1
Backward Euler	0	1

Table 4.1 Classification of single step methods

The Liniger-Willoughby parameter e is dependent on the step size and system eigenvalues, as computed from (4.29).

$$e = \frac{1}{h|\lambda_0|} - \frac{1}{\exp(h|\lambda_0|) - 1} \quad (4.29)$$

where λ_0 is the real part of the eigenvalue of the jacobian matrix such that,

$$|\lambda_j - \lambda_0| \leq b \text{ for all } h|\lambda_j| > 1, \text{ for all } j \quad (4.30)$$

It can be seen that both trapezoidal and Backward Euler rules are special cases of the Liniger-Willoughby rule. As employed by Sullivan, a fixed number of iterations were used to converge (4.28), and no error check was included. For the relatively small number of cases studied, it was possible to select the step size and method by trial and error for each case.

Rather than develop these methods for mixed systems and incorporate suitable error correcting algorithms, it was decided to use an

established integration package. Gears backward differential formula (Gear, 1971) was chosen mainly because a suitable implementation was readily available. The method is a variable-order variable step length prediction and correction algorithm based on the following approximation to the true solution;

$$h_j \dot{\underline{x}}_j = \sum_{i=0}^k a_{ij} \underline{x}_{j-i} \quad (4.31)$$

where $\underline{x}_j$ is the solution after step j , h_j is the length of the j -th step and a_{ij} is a constant which ensures that the approximation is accurate to k -th order.

Substituting this into (4.2) we obtain

$$A \sum_{i=0}^k a_{ij} \underline{x}_{j-i} + h_j f(\underline{x}_j, \underline{u}_j, t) = 0 \quad (4.32)$$

This is a standard nonlinear equation in $\underline{x}_j$ which can be solved by a Newton or quasi-Newton method. The number of iterations required to reach convergence is dependent on the initial guess of $\underline{x}_j$. A suitable starting value is given by

$$\underline{x}_j = \underline{x}_{j-1} + h_j \dot{\underline{x}}_{j-1} \quad (4.33)$$

where $\dot{\underline{x}}_{j-1}$ is computed from (4.31).

The practical implementation contains several refinements. The most important of these is the automatic order and step size control. The accuracy of the prediction (4.33) can be estimated by the subsequent behaviour of the equation solver on (4.32). If the prediction is good, both stepsize and order may increase. Conversely a poor prediction causes reductions of stepsize and order. The exact

conditions for stepsize and order changing can have a marked influence on the overall performance of the integrator. Another refinement is to hold the past history ($\underline{x}_{j-1} \dots \underline{x}_{j-k}$) as approximations to the time derivatives of $\underline{x}$ (the Nordsieck vector) at the current step. This makes the coefficients a_{ij} independent of h , and can simplify the coding. The code used for this work is due to Morison (1984), who provides a detailed description of the Gear algorithm and associated issues.

Examination of equations (4.8) and (4.24-27) will show that the solution of (4.1, 4.2) is required to evaluate the jacobian or the constraint violations. The jacobian is also required for the solution of equation (4.32). Generating the code for the jacobian by hand is considered unrealistic due to the size and number of cases it is intended to study, and no algebraic manipulation software is presently available to generate it automatically. The technique used is therefore to compute a finite difference approximation at every point required. This is a lengthy computing task, so precautions have been taken to accomplish it as efficiently as possible.

The first of these is to store and manipulate the jacobian in sparse form. The sparse linear equation solving routine MA28 from Harwell (Duff, 1979) is used whenever a matrix equation of the type $A\underline{x} = \underline{b}$ is solved. Secondly, the finite difference approximation is computed with groups of variables selected by the CPR algorithm (Curtis, Powell and Reid, 1974). The number of groups is typically one sixth of the number of equations, thus substantially reducing the computation time over a full finite difference process. Finally, Gears method itself is able to determine when the jacobian is no longer a sufficiently close approximation to the real system to ensure rapid convergence - the jacobian will only be updated when necessary.

4.4 INTERFACING WITH THE MINIMISATION PACKAGE

The performance of the optimal control program is critically dependent on the choice of, and interface to, the minimisation package at step 6 of the conceptual algorithm. For this work, a modified version of the Harwell code VF02AD is used. This is a standard non-linear optimisation routine which implements the Han and Powell quadratic programming algorithm (Powell, 1978). It has proved the most robust and efficient of all such packages tested and found suitable for optimal control work (Morison, 1984). This section is intended to illustrate some important aspects of using such a package in this study.

The minimisation problem as presented to VF02AD is defined in section 4.1, equations 4.12-13. For full details of the algorithm used VF02AD, the reader should consult Powell (1978). The following discussion will only consider aspects of the algorithm which directly affect the interface with the optimal control program.

4.4.1 Calculation of gradients

The basic form of VF02AD requires both the function and gradient information at every iteration. It is relatively expensive to compute the gradients of the objective function with respect to the decision variables, especially as they are not used at every iteration. In order to save computational effort by skipping the gradient evaluation, the VF02AD algorithm must be understood. The following explanation applies to unconstrained problems only. The presence of active constraints will modify the gradient vector and Hessian matrix so that the solutions always satisfy the constraints.

The objective function about a current solution $\underline{y}$ is represented by the first and second order terms of a Taylor expansion;

$$J(\underline{y} + \underline{d}) = J(\underline{y}) + \underline{d}^T \underline{w}(\underline{y}) + \frac{1}{2} \underline{d}^T H \underline{d} \quad (4.34)$$

where H is the Hessian matrix d^2J/dy^2 , $\underline{w}$ is the gradient vector dJ/dy and $\underline{d}$ is the perturbation in the solution from $\underline{y}$. The value of $\underline{d}$ which minimises $J(\underline{y} + \underline{d})$ is given by the Newton formula,

$$\underline{d} = -H^{-1} \underline{w} \quad (4.35)$$

The vector $\underline{d}$ is computed at each iteration and becomes a search direction for the next stage of the algorithm. This is a one-dimensional line search to minimise $J(\underline{y} + \alpha \underline{d})$, where α is a positive step length parameter not greater than one. As originally implemented, the line search starts with $\alpha=1$ and progressively reduces it until a reduction in J is obtained which is suitably close to the reduction predicted by (4.34). When the line search is complete the gradient values at the new solution are used to revise the Hessian as explained in Powell (1978). Precautions are taken to ensure it remains positive definite after the revision.

Powell explains this line search as a means 'to force convergence from poor starting approximations'. Another interpretation is that it enables the algorithm to work when the objective function is not adequately represented at the current solution by (4.34). Experience with the optimal control package has shown that many problems fall into this category, so the potential savings are large. This contrasts with Powell's own test problems for which 'on nearly every iteration the step length α has the value one'.

Figure 4.2 shows a hypothetical line search. Because the Hessian is positive definite, the predicted step will cause a reduction of the objective. This is easily seen from 4.34, in which the term $\frac{1}{2} \underline{d}^T H \underline{d}$ is always greater than zero. Although the main optimal control program cannot apply the same tests for step acceptance, we can be certain

that a step will never be accepted if the objective has not decreased. The optimal control program will therefore only compute the gradients if the corresponding objective function is less than all its predecessors.

Constrained optimisation problems are only amenable to this trick when all constraints are satisfied exactly. This applies unconditionally for linearly constrained problems after the first iteration. Non-linear constraints will only be satisfied as linearised approximations, so iterations which cause the objective to increase may still be accepted by the algorithm if they move the solution from an infeasible to a feasible point. The gradients must therefore always be evaluated if nonlinear constraints are present.




Figure 4.2 Line search for a function minimisation

4.4.2 Limiting the linesearch step length

Although the strategy given above is valuable in reducing the number of gradient evaluations, it is still necessary to compute the objective function at every trial step. The initial step length d depends only on the Hessian and current gradients, so if the curvature estimate of the revised Hessian remains too small, the sequence il-

lustrated above will be repeated at every iteration. This effect was observed on many occasions during SPEEVOP development.

The code has been modified to remember the last successful step length factor α . At each iteration, the first trial α is not allowed to exceed this by more than a predetermined factor. After the initial iteration this effectively limits the number of unsuccessful steps to one. For details of this modification, see Morison (1984).

4.4.3 Scaling the decision variables

Equation 4.34 is scale-independent; that is, the step-length as computed by 4.35 is not affected if all decision variables are transformed by multiplying the decision vector by a constant diagonal matrix. However the initial value of the elements of H is not known, so the first step of the minimisation algorithm is performed with a nominal approximation to H , often the unit matrix. This results in a steepest descent step causing one half unit change in the objective function, as can be seen by substituting I for H in 4.35, 4.34. This is generally a good strategy, as steepest descent methods are reported to perform well at points distant from the solution (Bryson and Ho, 1975).

Scaling is important in practice because subsequent progress of the algorithm depends on how quickly the Hessian can be updated to approximate the actual system Hessian. If the true Hessian has diagonal elements of 1000, several updates will be required before this can be represented in the approximation to H in the algorithm. Similar problems occur if the system Hessian has very small elements. Furthermore, very large and small numbers can cause numerical problems in the computation, so should be avoided. The effect of poor scaling is clearly illustrated in section 4.10. Some measure of protection can be provided in the algorithm by setting the initial estimate of H to be the unit matrix multiplied by some norm of $\underline{w}$. This has been incorporated in the current version of VF02AD.

4.4.4 Accuracy of data:

The accuracy of computing the gradients and objective function is important in this work because both are found by iterative integration schemes. (Sullivan (1978) avoided this problem by using an 'open-ended' implementation for his integration scheme, with a fixed number of corrector steps and fixed step length.) Since the number of iterations and time steps are variable, the result can only be guaranteed monotonic to within twice the specified accuracy. The optimisation process will be making relatively small changes in the objective function towards the end of an optimal control run. These changes could be less than the required optimisation accuracy, and will be accepted or rejected incorrectly if the errors are unfavourable, as shown in fig 4.3. The objective function should therefore be computed to at least twice the accuracy required in the optimizer. The gradient accuracy is also important since the change in the gradient is used in the Hessian update.




Figure 4.3 Effect of objective function error

4.5 CONVERGENCE CRITERIA AND SOLUTION ACCURACY

The iterative minimisation algorithm will produce a series of solutions of monotonically decreasing objective functions. After each iteration some tests must be applied to check whether the specified accuracy has been achieved, or some condition has arisen which prevents the algorithm making further iterations. The choice of termination criteria must be based on a set of assumptions about the behaviour of the objective and the decision variables. No set of termination criteria is suitable for all optimisation problems and all methods - the following discussion applies only to the algorithm and problems previously outlined.

We are primarily interested in obtaining a good estimate of the objective function because this will become our index of switchability. The most important assumption is that the objective function is smooth with respect to the decision variables, i.e. it is twice continuously differentiable. If the objective and its derivatives could be computed exactly the models used would satisfy this condition because their solutions are continuous. The iterative nature of the solution methods introduces discontinuities as explained in the previous section. The assumption will hold if a sufficiently high accuracy is specified for the objective and gradient determination.

Secondly we must assume that the quadratic approximation (4.34) is truly representative of the objective function. Superlinear convergence is the most convenient indicator of this. Powell (1978) describes how this property arises in the minimisation algorithm.

Thirdly the solution obtained is assumed to represent the global minimum of the problem, although we know that any solution from VF02AD will only be locally optimal. For this work, a rapid shift to the new target state and good regulation of the product specifications is a sign that the solution is globally optimal. By careful choice of

initial control trajectory, we can avoid unsound solutions to the switching problem.

There are two aspects to the test for convergence;

- (i) Has the desired accuracy been achieved?
- (ii) How accurate is the current solution?

The second aspect is of importance here as the problem complexity will sometimes prevent the routine from terminating by satisfaction of the applicable criteria in a reasonable number of iterations.

The convergence criteria may be stated in non-mathematical terms as follows (Murray, 1978);

- (i) The sequence of decision vectors $\underline{y}_k$ is converging.
- (ii) The sequence of objective values J_k is converging.
- (iii) The gradients dJ/dy are small.
- (iv) The Hessian is positive definite.
- (v) The decision vector satisfies all feasibility constraints.

Conditions (iii) - (v) are relatively easy to establish subject to the aforementioned assumptions. The remainder, (i) and (ii), cannot be checked rigorously because the exact values of $\underline{y}$ and J at the optimal solution are necessarily unknown. There are a number of methods for estimating whether the specified accuracy has been reached in the vicinity of the current solution.

The VF02AD code uses a simple test based on the projected change in the objective neglecting the curvature of the objective function surface. It is thus a first order approximation of the anticipated decrease in the objective using the current value of the reduced gradient and the previous step d. The reduced gradient is the gradient of the objective function with respect to the decision variables subject to any active constraints. The algorithm terminates when the projected decrease is less than the required accuracy. The test is illustrated in figure 4.4 below for a one-variable optimisation.




Figure 4.4 VF02AD termination test in one variable

Additional causes of termination in VF02AD relate to the provision of gradients and function values which are not compatible with previous iterations, or render the problem badly posed. These abnormal termination conditions are easily recognised and appropriate corrective action taken.

4.6 SIMULATION USING SPEEDUP

The process networks which are the subject of this study can be described by a mixed set of differential and algebraic equations, formed by linking together the appropriate mathematical models (see chapters five and six). The resulting equation set is quite large, containing over one hundred equations in some cases, and to code these by hand is both tedious and prone to error. In view of the relatively large number of distinct networks to be investigated, it was decided at an early stage to use an existing flowsheeting package to convert the modular description of the process networks into equation sets. The same package could also be used to find steady-state and dynamic solutions.

The SPEED-UP package developed at Imperial College was chosen for this purpose. SPEED-UP is an acronym for "Simulation Package for the Evaluation and Evolutionary Design of Unsteady Processes", and a general description may be found in Perkins and Sargent (1982). This equation-oriented package will take a process network described by the unit operations and associated interconnections, and solve the steady-state or dynamic problem for any consistent set of unknown variables within the network. This means that (for example) the operating conditions may be found to meet a certain product specification, as well as normal simulation. The user has complete flexibility in modelling the unit operations with both equations and procedures, so there is no difficulty in incorporating the models described later.

SPEED-UP operates in two phases. The translation phase takes the modular process description, and generates a series of files containing the set of Fortran equations and the associated numerical data. The solution phase uses the information of these files to solve the original steady state or dynamic problem. The user can specify the solution method to be used, and a number of full and sparse matrix non-linear equation solvers are available.

When setting up a SPEED-UP problem, some care must be exercised to ensure convergence of the equation solver. In general, the user will be able to guess the initial values of many of the variables, or place relatively tight bounds on them using information derived from simpler models. The rate of convergence of many examples was found to be very dependent on the starting point. This is not surprising, as the systems are highly nonlinear.

The equation solver used for all steady state simulations in this work is a quasi Newton method due to Broyden, and implemented by Paloschi (1982). The advantage of this class of methods over the standard Newton methods is that derivatives of the equations with respect to the unknown variables do not need to be supplied at each iteration. Instead, they are estimated from the function values generated in successive iterations. The initial Jacobian matrix (the matrix of these derivatives) is estimated by finite differences (see section 4.3).

A further feature of SPEED-UP is the ability to recognise a dependent block structure in the unknown variables. This is used to decompose the problem into a series of subproblems which may be solved sequentially. Detailed discussions on these topics are given in Bogle (1983) and Paloschi (1982).

Dynamic simulation within SPEED-UP uses Gear's BDF method, which is discussed in section 4.3. The initialisation difficulties encountered with steady state problems are avoided because the initial state is generally specified as a perturbed solution of a compatible steady state.

For the purpose of this research, the SPEED-UP steady state simulation is used to define the initial and final states of a switch. Certain key variables associated with each product stream are fixed, and the equations are solved to find the appropriate steady state control policy at each operating point. This policy becomes $\underline{u}_2$ in equation (3.19) for each switch to that steady state.

If a single input - single output closed loop control is easily derived, an approximation to the optimal control policy may be found by dynamic simulation. The main drawback is that the controllers will contribute extra equations which increase the overall complexity of the system. These equations may have to be eliminated from the system before the open-loop optimal control policy can be found.

The solution from both steady-state and dynamic simulations of the system is readily accessible to other software packages on the set of files generated by the SPEED-UP translator and updated by the equation solver. The SPEEVOP package written specially for this work, is described in section 4.8. SPEEVOP will construct an optimal control problem from the problem definition and solutions within these files.

4.7 PROBLEM DEFINITION THROUGH SPEEVOP

Given the system of equations which described the process network and their steady state solution at the terminal states of the switch, it should be possible to set up the optimal control problem. In practice it is necessary to provide additional information to overcome restrictions in both the original flowsheeting package and the optimal control algorithm. A comprehensive optimal control package, SPEEVOP (SPEEDUP VF02AD OPTimiser) has been developed to redefine the problem as initially presented by SPEEDUP and implement the algorithm given earlier.

The problem is defined through three routes:

1. Executable code - system equations.
2. SPEEDUP data - variable types, initial solution etc.
3. Optimisation data - all information not specified in 1 or 2.

The system equations are a set of FORTRAN routines generated by the SPEEDUP translator. They are compiled and loaded as an integral part of the optimisation package. This means that each system to be investigated must be loaded separately. The SPEEDUP data is contained in a series of files which form a database for subsequent optimisation runs. These files will only store information relating to the original SPEEDUP simulation. All other information specific to the optimisation is given directly by the user to SPEEVOP.

The logical flow of information between SPEEDUP, SPEEVOP and the set of files is shown in figure 4.5 overleaf.




Figure 4.5 Logical data flow : SPEEVOP and SPEEDUP

The user dialogue with SPEEVOP is designed so that the same commands can be used in a remote or interactive manner. Each command consists of a keyword followed by a series of parameters, such that wherever possible, each line of input is independent of the preceding lines or the SPEEVOP response. This prevents the package entering undesirable modes of operation after a simple typographical error. The following keywords are recognised by SPEEVOP.

A) Keywords effecting the initial control trajectory and states.

- NEW - Clears COMMON file.
- PRESET - Presets any variable or group of variables.
- SAVE - Saves part or all of COMMON file.
- USE - Copies another COMMON file into part or all of current COMMON file.

B) Keywords to specify the optimal control problem.

CONTROL - Allows user to specify control variables to optimiser
DESIGN - As above, but for design variables.
ELIMINATE - Specifies equations and variables to be omitted from
the system equations.
LIMIT - Specifies algebraic variables which are to be limited
as control constraints.

C) Keywords which affect the optimisation algorithm directly.

FINT - Set final time t_f .
OPTIMISE - Starts optimal control algorithm.
SCALE - Specifies scaling parameters for minimisation routine.
SWITCH - Sets initial control interval times.

D) Keywords to control the output format.

OUTPUT - Set frequency and destination of output dump.
RESULT - Specify variables to be output.
TINT - Fix minimum simulation time between dumps to COMMON file.
UPDATE - Set COMMON file update frequency.

E) Miscellaneous.

AUDIT - Debugging aid.
END - Terminate SPEEVOP run.
* - Treat line as a comment.

The problem definition is constructed primarily in the COMMON file, which remains compatible with the file of the same name created by SPEEDUP. Where the COMMON file has no facility to store certain information, it is held in a series of tables in the SPEEVOP program. When the problem definition is complete the OPTIMISE command may be issued to begin an optimal control run. Once this terminates the program is ready to accept further commands, for example to print

selected groups of variables, or to reoptimise with a different set of conditions.

The optimisation run itself will dominate the total time spent in the package, so every effort has been made to reduce the problem to the simplest form. All problem reduction is carried out as the first stage of the optimal control run, before the actual optimal control program begins. The most important part of this is to eliminate all redundant equations and variables in the original SPEEDUP problem description from the system of equations (4.1) used by the optimal control routines. The savings can be quite considerable, especially through careful use of the ELIMINATE keyword. It is not generally possible to omit the equations from the SPEEDUP system, since they are required to define certain constraints such as product purity which become extra degrees of freedom in the open-loop optimisation. The reduced sets of equations and unknowns are packed into contiguous groups through indirection lists, so the optimal control routines only manipulate the minimum vectors needed to describe the problem as in (4.1) et seq. The sparse matrix information for the various derivatives is derived from the original SPEEDUP occurrence matrix and similarly reduced. The algorithm to accomplish the problem reduction is one of straightforward list manipulation - details are contained in the program listings (Appendix G - microfiche).

4.8 IMPLEMENTED OPTIMAL CONTROL PROGRAM

The optimal control program forms the heart of the SPEEVOP package, and as such warrants a detailed description of the individual steps within it. The original implementation of Morison has been substantially modified for this work. The conceptual algorithm of section 4.2 is adapted for computational efficiency as follows:

1. Initial control trajectory.

The initial switching times are defined through SPEEVOP and the control levels are extracted from the COMMON file. They are scaled and packed for VF02AD.

2. Forward integration.

Equations (4.1-2) are solved with the current control policy using Gear's method, to a tolerance of one fifth of the optimisation tolerance. The state trajectory at every time step is stored on a scratch file in Nordsieck vector form. This is a vector of the state time derivatives which allows accurate interpolation formulae to be used. The full SPEEDUP data structure is written to the COMMON file at predetermined time intervals.

3. Descent step test.

The value of $J(t_f)$ for the current control policy is checked against the best obtained by the program so far. If it has decreased, and there are no active nonlinear constraints, then the minimisation routine is likely to accept the step, so gradients must be computed. Otherwise step 4 in the algorithm will be skipped.

4. Backward (Adjoint) integration.

Equations (4.8-9) are solved about the state trajectory determined at step 2 in the algorithm. Gear's method is used to find the adjoint trajectory with a tolerance of five times the optimisation tolerance. After each control sub-interval, equations (4.16-18) are solved for the gradients. Gradient integrations

are explicit so trapezoidal rule is used. By finding the adjoint vector and gradients concurrently it is unnecessary to save the adjoint trajectory for the entire period t_0-t_f .

5. Linear Constraints and gradients.

Equations (4.19-23) are solved for the constraint values and the gradients are set up by writing 1 and -1 to the appropriate elements of the constraint gradient array.

6. Nonlinear constraints and gradients.

The algebraic subset of (4.1) is determined by examination of the occurrence matrix of the jacobian. For each time interval the state trajectory at the end of the interval is used to solve (4.24-25) for the constraint values. The same point in the state trajectory is used to find the jacobian matrices to solve (4.26-27) for the gradients.

7. Nonlinear minimisation.

The minimisation program VF02AD is called with the values and gradients previously calculated, packed and scaled by steps 4-6 of the algorithm. If the minimum has been detected, or some other condition arises which inhibits further progress, the appropriate flag value will be sensed on return and the optimal control program will terminate. Otherwise, a new decision vector is set up, and the program continues from step 2 of the algorithm.

On termination, the optimal control program returns control to SPEEVOP. The latest version of the optimal control policy and associated state trajectory remains in the COMMON file. This becomes a starting point for subsequent optimal control runs. The policy can evolve under close control of the user, who may fine-tune certain variables to accelerate convergence between runs. The overlay structure is shown in figure 4.6. This mirrors the logical dependence of the various stages given above.




Figure 4.6 Overlay structure of SPEEVOP

4.9 TESTING THE OPTIMAL CONTROL PACKAGE

The total size of the SPEEVOP package code is around 10000 lines - a program of this size is vulnerable to coding errors during development and maintenance. In addition, changes of operating system, and revisions of the SPEEDUP flowsheeting package can give rise to obscure faults in formerly working routines. To verify SPEEVOP both during development and following any changes, two simple problems have been set up. These problems use all the facilities of SPEEVOP, so can be used to check every part of the package.

4.9.1 Problem 1

This 3-variable problem has a known continuous solution, so it provides an opportunity to judge the penalty of discretising the control action. The problem is to minimise the total square deviation error and control action over the period of interest for the system below.

$$\left. \begin{aligned} \dot{x}_1 &= u \\ \dot{x}_2 &= x_1 \\ \dot{x}_3 &= x_1^2 + x_2^2 + u^2 \\ t_0 &= 0, \quad t_f = 1 \\ x_1(t_0) &= x_2(t_0) = 1, \quad x_3(t_0) = 0 \\ J &= x_3(t_f) \end{aligned} \right\} \quad (4.36)$$

SPEEVOP obtained an optimal solution to this problem with 5 time intervals in 15 iterations. This solution is shown, together with the continuous optimal solution, in figure 4.7. The continuous solution was obtained by the quasi-linearisation method of section 4.2. This

could be successfully applied because the control u is well defined by the time derivative of the objective function, $\dot{x}_3$. The mean difference of the two solutions is relatively small, and the objective function error is hardly significant, even with only 3 control intervals, table 4.8.




Figure 4.7 Optimal Control policy for problem 1

Solution	no. of intervals	$J(t_f)$
Quasi linearisation		2.710
SPEEVOP	5	2.716
SPEEVOP	3	2.725

Table 4.8 Solutions to problem 1 by method

4.9.2 Problem 2

This problem is superficially similar to problem 1, but is solved for a different objective and with nonlinear control constraints which are active at the solution. The problem is to minimise the terminal error on the two differential variables for the following system;

$$\dot{x}_1 = u + v$$

$$\dot{x}_2 = x_1$$

$$x_3 = x_1^2 + x_2^2$$

$$x_4 = (u^2 + v^2)^{\frac{1}{2}}$$

$$t_0 = 0, t_f = 1$$

$$x_1(t_0) = x_2(t_0) = 1$$

$$x_4 \leq 3$$

$$J = x_3(t_f)$$

(4.37)

The problem has one design variable v and one control variable u which are mutually constrained through the bound on x_4 . This restriction limits the total control action so as to prevent both x_1 and x_2 reaching 0 within the time allowed. This constraint may equally be written as

$$x_4^2 = u^2 + v^2 \tag{4.38}$$

but this does not define x_4 uniquely, since it may be either positive or negative. The expression in (4.37) carries the implicit assumption that x_4 will always be positive.

SPEEVOP will solve this problem when posed with two control intervals in twelve iterations. Scaling has a significant effect on the number of iterations for the reasons mentioned in section 4.4. The scaling factor is defined in this work as the constant by which a specified gradient is multiplied before presentation to the minimisation package. The gradient is increased by reducing the span of the decision variables, as shown overleaf (4.39).

$$\begin{aligned}
 y_1 &= u_1/s_1 \\
 y_2 &= u_2/s_1 \\
 y_3 &= t_1/s_2 \\
 y_4 &= v/s_3
 \end{aligned}
 \tag{4.39}$$

The nominal solution obtained with unit scaling and a specified accuracy of 0.001 is shown in table 4.9. The effect of scaling the gradients on the number of iterations to reach this solution is given in table 4.10, and the solution progress is plotted in figures 4.11 - 4.13. The small curvature of the objective function surface in this problem was evident by the variation of the decision variables about the nominal solution for different scaling factors. Relative variations up to 0.04 were noted although the objective function lay well within the specified accuracy. This illustrates the advantage of quadratic programming algorithms over steepest descent methods near the solution, as noted in many comparative reports eg. Bryson and Ho (1975).

The effect of prescaling the unit matrix for the initial value of the Hessian is also shown in figures 4.12 and 4.13. The prescaling was disabled for the results in table 4.10. The principal advantage is to make the algorithm more robust to gross scaling errors, by limiting the step length through an artificially high initial objective function surface curvature.

Summarising, this is a deceptively tough problem, but SPEEVOP was able to solve it even when badly scaled. The worst effects of bad scaling were poor progress of the algorithm; in practice these could easily be recognised and corrected.

Variable	Value at optimum solution
$u(t_0)$	-2.888
$u(t_1)$	+2.888
v	-0.816
t_1	0.5804
J	0.509

Table 4.9 Nominal solution to problem 2

s_1	s_2	s_3	Iterations
0.02	0.02	0.02	30
0.2	0.2	0.2	17
1	1	1	12
10	10	5	12
50	50	10	20

Table 4.10 Effect of gradient scaling in problem 2




Figure 4.11 Progress of the solution for small scaling factors




Figure 4.12 Progress of the solution for large scaling factors




Figure 4.13 Progress of the solution with scaled hessian

4.9.3 Use of analytic jacobians

Initially it was proposed to use a jacobian matrix derived analytically for all parts of SPEEVOP. The proposal was dropped for reasons given earlier (section 4.3), but it is interesting to assess the overall effect of this decision on the performance of SPEEVOP. Sullivan (1978) developed a binary column model for which he derived

the analytic jacobian. This model was rewritten in SPEEDUP format and submitted to SPEVOP with and without the analytic jacobian. The results are shown in table 4.12.

Jacobian	Analytic	Finite differences
Iterations	29	27
Objective value	-56.47	-56.48
Run time	606.2	680.0
(CDC 6600 CP seconds)		

Table 4.12 Comparison of jacobian estimation methods

It is perhaps a little surprising that the finite difference method has outperformed the analytic method except in CP time. This is probably a chance result, in that the truncation errors have accelerated the convergence at some point in the program (see section 4.4). The significant fact is that the average time per iteration is only slightly greater - certainly there is insufficient difference to justify the trouble of setting up analytic jacobians for each case study.

CHAPTER FIVE

MODELLING OF HEAT EXCHANGERS

In this chapter the mathematical model of a liquid-liquid heat exchanger suitable for optimal control studies is developed. This model retains the features which are considered important to switchability of more complex models while using far fewer equations. The responses of the model to typical operating disturbances are shown to compare favourably with the responses obtained from a complex multi-section model. Finally a model is presented for exchangers which involve a utility stream (steam or cooling water).

5.1 REQUIREMENTS

To perform optimal control studies on a network containing heat exchanger units, a suitable mathematical model of a heat exchanger is required. This model must be as simple as possible, to facilitate the large number of calculations required by the optimal control algorithm, while retaining the relevant features of the real system. These are;

1. The steady state solution must be accurate at both extremes of the switch.
2. The dynamic response must be qualitatively correct for every disturbance or control. This implies that the dominant eigenvalues should be similar, to give the response the correct shape.

These requirements are not actually met by any model reported in the literature which is sufficiently simple for optimal control. There is however a considerable body of information to help develop a suitable model.

Before proceeding with the literature survey, it is helpful to define the scope of this work. The following assumptions apply:-

1. The heat capacity flowrate is constant for each stream.
2. Composition lags through the heat exchanger are negligible.
3. Heat transfer rate is a function of the physical construction of the exchanger and temperature difference alone.
4. All heat capacity is confined to the working fluid.

Assumptions 1 and 3 have been used in almost all previous work. They are retained here as a logical first step in the simplification process. The second assumption is justified because, if necessary, composition lags may be lumped into the associated distillation units. Similarly the thermal capacity of the exchanger material may be lumped with the capacity of fluid.

5.2 MODEL CLASSIFICATION

Many investigations have been reported on the modelling of heat exchangers; Mozely (1956), Campbell (1958), Matsubishi (1960), Koppel (1962), Stermole et al (1963) and Finlay et al (1965) present various aspects of dynamic models. The general approach of these investigators is to compare rigorous mathematical or experimental results with proposed models, to determine a suitable level of simplification. Three distinct classifications emerge from these studies:-

- a) Distributed parameter models.
- b) Cascaded double convection models.
- c) Modified double convection models.

In the first class, the heat flow is expressed by partial differential equations, in which the distance along the exchanger is the independent variable. The steady state solution of these equations for plug flow of parallel fluid streams gives the classical log-mean temperature difference formula. Koppel, Finlay and Stermole solve the equation for the response of exit temperature to flowrate, while Mozely, Campbell and Matsubishi consider only the effects of inlet temperature change.

All research work on concentric pipe exchangers suggests that distributed parameter models are the most accurate over the full range of disturbances. The transfer function solutions, which are presented for different cases in every publication given above, can become very complex. However the distributed parameter model is only a cascaded double convection model in which the number of stages is infinite - the principal difference is in the method of solution. Therefore, any conclusions which apply to the latter can be extended to the former, so there is no need to discuss it separately.

The second class of models treat the exchanger as a collection of perfectly stirred thermally coupled tanks. Heat transfer between any two tanks is a function of their relative bulk temperatures. All types of multipass configurations can be represented by suitable networks of tanks. Figure 5.1 shows the network for a simple counter-flow arrangement.




Figure 5.1 Cascaded double convection model

$$H_{1i} \frac{dT_{1i}}{dt} = F_1(T_{1,i-1} - T_{1i}) - Q_i \quad (5.1)$$

$$H_{2i} \frac{dT_{2i}}{dt} = F_2(T_{2,i+1} - T_{2i}) + Q_i \quad (5.2)$$

$$Q_i = UA_i(T_{1i} - T_{2i}) \quad (5.3)$$

where H is the thermal holdup,

F is the heat capacity flowrate,

T_{mn} is the temperature of stream m at n ,

UA is the heat transfer coefficient, and

Q is the heat transfer rate.

This model receives favourable references from all who consider it. Mozely (1956) reports very good agreement with experiment for a 5 section model, and fair results from even a single section. Campbell (1958) states that seldom will more than 5 sections be needed for any

practical exchanger model. Matsubishi (1960) claims good agreement to inlet temperature variations between experiment, distributed parameter, and two- and four-section models. Finlay et al (1965) find the single section model is reasonably close to experiment for flow variations, and a 6 section model even better, particularly at lower frequencies.

As the number of sections is increased, while holding the total capacity constant, the cascade model will approach the distributed parameter model in all respects. Finlay et al suggest that discrepancies between the distributed parameter model and experiment are due to axial dispersion effects, so it appears that a high-order cascade is the closest to the 'real' system.

The third class confines each stream to a single tank, but computes the heat transfer rate by assuming a steady state temperature profile within each tank along a common wall. Thus the driving force is based on some mean temperature difference between inlet and outlet temperatures of both tanks. The mean fluid temperatures may be calculated at any points along the exchanger, but conventionally either the midpoint or the exit temperature is used. See figure 5.2.




Figure 5.2 Modified double convection model

$$H_i \frac{dT_{im}}{dt} = F_i(T_{i1} - T_{i2}) \pm Q, i=1,2 \quad (5.4)$$

$$Q = UA T_{LM}(T_{11} - T_{22}, T_{12} - T_{21}) \quad (5.5)$$

$$T_{im} = e T_{i2} + (1-e)T_{i1}, i=1,2 \quad (5.6)$$

Where T_{LM} is the log mean temperature difference,

T_m is the mean temperature,

e is the position along the length of the exchanger ($0 < e < 1$) where T_m is calculated.

Mozely provides convincing results for this model using $e = \frac{1}{2}$ and an arithmetic mean driving force for heat transfer. Phase and amplitude response show very good agreement with experiment up to a phase shift of 160° where the phase response diverges sharply, at which frequency the amplitude is relatively small.

Stermole et al reach similar conclusions for the case where $e = 1$, also using an arithmetic mean driving force. This model is actually a better representation of their experimental results than the 'rigorous' distributed parameter model up to the frequency of first resonance, when the time delays in the exchanger reinforce the sinusoidal disturbance (This feature is probably due to axial dispersion in the experimental apparatus).

Since the inlet and outlet temperature differences are of similar magnitude in these cases, the use of log mean temperature differences is not necessary. Nevertheless it would be essential if accurate steady states were required, as in the present work.

5.3 APPLICABILITY OF PREVIOUS WORK

None of these classes of model is suitable for the present work. Distributed parameter models are excessively complex, especially for numerical solution. The cascaded double convection models have a maximum thermal efficiency, E , set by the number of stages; n . In the worst case, where the heat capacity flowrates are equal in pure counter-current flow, it can be shown that

$$E \leq \frac{n}{n+1} \quad (5.7)$$

$$\text{where } E = \frac{\max_j (|T_{j,IN} - T_{j,OUT}|)}{|T_{2,IN} - T_{1,IN}|}, \quad j=1,2 \quad (5.8)$$

The efficiency of heat exchangers in the networks under consideration is relatively high, so the model will require more stages to satisfy the steady state conditions than are needed for dynamic accuracy. As two extra equations will be introduced per stage this is an undesirable solution.

This only leaves the modified double convection model, which has two peculiarities. Firstly, unless the accumulation term is computed at the exit alone ($e = 1$), the model will artificially introduce right-half-plane zeroes into the transfer response. These are manifest as inverse responses to certain disturbances, and the resultant model will not be qualitatively correct. Consider for example the system shown in fig 5.3, with holdups at the midpoint of each stream ($e=\frac{1}{2}$).




Figure 5.3 Heat exchanger model with inverse response

$$G(s) = \frac{(s + 4.828)(s - 0.828)}{(s + 6)(s + 2)} \quad (5.9)$$

This has one RHP zero on the real axis, although no other model class would have any poles or zeroes in the RHP for a similar unit. It is, of course, possible to approximate a pure time delay using transfer functions containing right half plane zeroes. The well-known Pade formulae are of this type;

$$G(s) = e^{-ts} \approx \frac{1 - ts/2}{1 + ts/2} \quad (5.10)$$

Because these are derived from a Taylor series expansion, they will only be accurate over a small frequency range. The higher frequency terms in a step change (such as would be applied through the optimal control package) will be subject to phase distortion. It is this which gives rise to the observed inverse response.

Secondly, the model will not handle the wide range of inlet temperatures which the optimal control algorithm might impose. The assumption of steady state temperature profiles only works while the temperature difference has the same sign at both ends of the exchanger. A large inlet temperature transient could invalidate this, and the equations would become insoluble. We cannot simply switch to an alternative formula for the duration of the transient, because the

model may still reach an alternative (incorrect) solution using this formula.

This is illustrated by the following example, Figure 5.4, in which the model uses the arithmetic mean when it cannot employ the log mean formula. Different initial conditions give two steady state solutions;




Figure 5.4 Heat exchanger model with multiple solutions

	T_{12}	$T_{12} - T_{21}$	T_{22}	$T_{12} - T_{22}$
Arithmetic mean	80.0	80.0	120.0	-20.0
Logarithmic mean	83.832	83.832	97.009	2.991

Table 5.5 Solutions of model in figure 5.4

The test for crossover has selected the appropriate formula in both cases, but one is obviously wrong, since crossover cannot occur at steady state.

The restrictions on the classical models are unacceptable, so a new heat exchange model has been developed specially for this work.

5.4 DEVELOPMENT OF THE NEW MODEL

The derivation of and justification for the log mean temperature difference as a driving force is completely accepted for all steady state chemical engineering applications. Differences between actual and ideal units are routinely corrected by published factors, e.g. TEMA (1978). It is therefore reasonable to base our model on this driving force if it is compatible with the dynamic requirements.

The model is to be incorporated in a set of equations to be solved by a general simulation package. This requires a variant which is unconditionally well behaved, as the variables will take on a wide spectrum of values as the solution converges. The difficulty with using the log mean formula directly is that it cannot be evaluated under certain conditions. A summary of this problem is given by Paloschi et al (1983). The most satisfactory solution is to implement it as a procedure, so that the appropriate equations can be selected and numerical difficulties are avoided.

The log mean temperature difference equation and associated steady state heat balances may be solved analytically to give the outlet temperatures in terms of the inlet temperatures, specific heat flows and heat transfer coefficient. For a counterflow unit with constant specific heat;

$$T_{12} = \frac{T_{11}(1-F_1/F_2) + T_{21}(\exp(x)-1)}{\exp(x) - F_1/F_2} \quad (5.11)$$

$$T_{22} = T_{21} + (T_{11}-T_{12}) F_1/F_2 \quad (5.12)$$

where T_{11} , T_{21} are the inlet temperatures,

T_{12} , T_{22} are the outlet temperatures,

F_1 , F_2 are the heat capacity flowrates

and x is a factor which indicates the temperature profile in the unit.

$$x = \frac{UA}{F_1} - \frac{UA}{F_2}, \quad (5.13)$$

where UA is the heat transfer coefficient - heat transfer area product.

If either F value is very small, there is a risk of division by zero. If the F's are equal, the log mean will be undefined, so the arithmetic mean formula should be used instead. Finally, if the heat transfer coefficient is very large, the exponent of x will exceed the numeric range of the machine. Thus we need 5 separate solution methods if we are to cope with all possible cases.

The profile parameter x can be used as an indicator to select the appropriate solution. If x is small, use the arithmetic mean formula for T_{12} ;

$$T_{12} = \frac{T_{11}((2/UA + 1/F_2) F_1 - 1) + 2 \cdot T_{21}}{(2/UA + 1/F_2) F_1 + 1} \quad (5.14)$$

For large x, indicating total recovery from one stream, neglect the temperature difference at that outlet;

$$T_{12} = T_{21} \quad \text{for } F_1 \leq F_2 \quad (5.15)$$

The remaining two cases cover either or both flows tending to zero. As x will be infinite in such cases, the algorithm must trap them before attempting to find x. The outlet temperature of a zero flow stream is equated to the inlet temperature of the other stream in all these cases.

A procedure called SSHEX (see Appendix G) has been written incorporating these equations and tests. The input parameters are the terms T_{11} , T_{21} , F_1 , F_2 and UA above, and the outputs are T_{12} and T_{22} .

The foregoing model is ideal for steady state calculation but cannot be used directly for dynamic models. To determine the dominant dynamic features, the eigenvalues and eigenvectors of the 4-stage cascaded double convection model were found. They are illustrated in the manner of Levy et al (1969). The vertical axis is the real part of the eigenvector, the horizontal axis indicates the position of the corresponding state variable along the exchanger in figure 5.6 below.




Figure 5.6 First 4 eigenvectors of cascade model

The slowest eigenvalue, almost a third of any other, is symmetrical, and represents the heat capacity of the complete exchanger. This will be strongly excited by all simple disturbances entering the exchanger and will appear as a first order lag on the output temperatures. Superimposed on this will be the other modes, which represent the exchange of heat between individual sections of the exchanger model. Since we are only interested in the principal time constant, these fast modes may be safely ignored. The final model now appears as a steady state log-mean driving force heat exchanger with a first order lag on each output temperature.




Figure 5.7 Proposed heat exchange model

$$T_{12} = f(F_1, F_2, T_{11}, T_{21}, UA) \quad (5.16)$$

$$T_{22} = f(F_2, F_1, T_{21}, T_{11}, UA) \quad (5.17)$$

$$\tau_1 \frac{dT_{13}}{dt} = T_{12} - T_{13} \quad (5.18)$$

$$\tau_2 \frac{dT_{23}}{dt} = T_{22} - T_{23} \quad (5.19)$$

where τ_1 , τ_2 are the effective time constants on each stream. For units with relatively small temperature changes in each stream, these time constants will be comparable to the residence time of the streams. However, in the high efficiency units considered in this work, where thermal coupling is tight, the time constants will be closer to the total residence time;

$$\tau_1 = \tau_2 = \frac{H_1 + H_2}{F_1 + F_2} \quad (5.20)$$

5.5 EVALUATION OF THE NEW MODEL

Numerical studies were performed to compare the new model with the 4 section cascade double convection and also with a modified double convection model. Several different heat exchangers were considered over a range of operating conditions, and in every case the new model satisfied all requirements laid out at the beginning of this chapter. The case described in Appendix B is considered to be the most difficult, as it demonstrates high heat recovery with transient temperature crossover.

The 4 section cascade model (I) and candidate model (II) are as previously described. The modified double convection model (III) uses the log mean formula to compute the driving force while the temperature differences are suitable, otherwise the heat transfer rate is set to zero. For reasons discussed previously, $e = 1$. Table 5.8 shows the results obtained for the three models. In line with previous work, model I is used as the 'true' representation of a system with finite axial dispersion.

Model number	I	II	III
Graph name	4 section cascade	LMTD+CSTR	LMTD Double C
Principal eigenvalue **	-2.69	-2.36 *	-2.04
Next eigenvalue **	-7.28	-	-11.1
Accuracy of response	-	good	poor
Computation time			
CDC 6600 CP seconds	20	9	73

* repeated

** eigenvalues of state matrix from linearisation of initial state

Table 5.8 Dynamic model comparison

The dynamic response for each system is shown in figures 5.9-11

The general agreement of principal eigenvalue for each model is misleading, as the dynamic simulation graphs show. The modified DC model shows the change from transient to logmean formula very clearly, and this is also the cause of the long execution time. The integration routine had to reduce the steplength by several orders of magnitude to pass the discontinuity without loss of accuracy. (The run time for the same model without a discontinuity was only 15 cp seconds) In contrast, the new model satisfies all requirements for the optimal control package, and has been used in subsequent work.





Figure 5.9 Response of model I to step temperature change





Figure 5.10 Response of model II to step temperature change





Figure 5.11 Response of model III to step temperature change

5.6 UTILITY EXCHANGERS

A practical heat exchange network may contain units which use steam and cooling water to heat or cool process streams. Such units which have a process stream on one side only are collectively known as utility exchangers, and differ from normal inter-stream units in two respects. Firstly, the control action is not free - utilities represent energy consumption by the network, and any switchability study must take this into account. Secondly the model need not be so elaborate, since these units are only concerned with one stream.

The model employed in this work treats the utility exchanger as a well-mixed tank, into which the heat content of the utility (steam or water) is completely dissipated. This agrees with the practical operation of a steam-heated unit, where the latent heat of the steam, which provides virtually all the heat, is completely transferred to the process stream. While not strictly true for cold utilities, it is nevertheless adequate, since cooling water is substantially cheaper than steam. The cost of utilities in unit time is therefore proportional to the utility flowrate.

CHAPTER SIX

MODELLING OF DISTILLATION COLUMNS

Suitable dynamic models of distillation columns must be incorporated in the mathematical description of the heat integrated network. This chapter states the fundamental equations and simplifying assumptions which are currently used to model distillation processes, then progressively develops the simpler models reported in the literature. A review of significant publications in both full and reduced order models is included. A new model is proposed which has significant advantages over previous models for the study of switchability. Comparative data for a binary column shows the new model preserves all the important features of the full-order rigorous model which was tested.

6.1 FULL ORDER MODELS

These models are derived by writing material and energy balances for each stage; they can yield accurate representations of dynamic and steady state features, including non-linear effects. A number of assumptions are commonly made, as follows:-

1. Negligible heat losses from column.
2. Negligible heat capacity of column material.
3. Negligible vapour holdup.
4. Liquid and vapour leaving tray are in thermal equilibrium.
5. Liquid on tray is perfectly mixed.
6. Murphree efficiency determines vapour composition.

The equations which can describe a tray (see figure 6.1) are;

Component mass balance

$$\frac{d}{dt}(H_n x_{in}) = L_{n-1} \cdot x_{i,n-1} - L_n \cdot x_{in} - V_n \cdot y_{in} + V_{n+1} \cdot y_{i,n+1} \quad (6.1)$$

Overall mass balance

$$\frac{d}{dt} H_n = L_{n-1} - L_n - V_n + V_{n+1} \quad (6.2)$$

Enthalpy balance

$$\frac{d}{dt}(H_n h_n^L) = L_{n-1} \cdot h_{n-1}^L - L_n h_n^L - V_n \cdot h_n^V + V_{n+1} \cdot h_{n+1}^V \quad (6.3)$$

Mole fraction sums

$$\sum_{i=1}^{nc} x_{i,n} = 1 \quad (6.4)$$

$$\sum_{i=1}^{nc} y_{i,n} = 1 \quad (6.5)$$

Vapour composition

$$y_{n,i} = y_{n+1,i} + E_M (K_{ni} x_{ni} - y_{n+1,i}) \quad (6.6)$$

Vapour flow

$$P_{n+1} = P_n + 3.226 \cdot 10^{-5} \frac{(V_n)^2 T_n}{A_p} + d_{Tn} \rho_L \quad (6.7)$$

Tray depth of liquid

$$d_{Tn} = 0.95 \left(d_w + 0.0392 \frac{L_n M_{WT}}{\rho_L l_w}^{0.667} \right) \quad (6.8)$$

Downcomer depth of liquid

$$P_n + d_{Tn} \rho_L = P_{n-1} + d_{Dn} \rho_L + 3.192 \cdot 10^{-5} \frac{1}{\rho_L} \left(\frac{L_n M_{WT}}{A_{DCC}} \right)^2 \quad (6.9)$$

Total tray holdup

$$H_n = \rho_L (d_{Dn} A_D + d_{Tn} A_T) / M_{WT} \quad (6.10)$$

The following symbols have been used:

H	Molar liquid holdup (Kmoles)
L	Molar liquid flow leaving tray (Kmoles/min)
V	Molar vapour flow leaving tray (Kmoles/min)
x	Mole fraction of component on tray
y	Mole fraction of component in vapour leaving tray
h	Molar enthalpy (GJ/kmol)
E _M	Murphree efficiency
K	Vapour liquid equilibrium K-value
P	Pressure of vapour on tray (Nm ⁻²)
T _n	Temperature on tray (Kelvin)
ρ _L	Liquid density (kg/m ³)
d	Liquid depth (m)

- A_p Perforated area of tray (m^2)
 A_D downcomer area (m^2)
 A_{DCC} Downcomer cutaway area (m^2)
 l_w Weir length (m)
 A_T Total area of tray (m^2)
 M_{WT} Molecular weight (kg/Kmole)
 n Subscript - tray number 1 to NP
 i Subscript - component number 1 to NC




Figure 6.1 Simple distillation tray

The hydraulic correlations for vapour and liquid flow (6.7 to 6.9) are based on formulae given in Van Winkle (1967). The expressions have been converted to compatible units for this work, and vapour density is computed for simplicity from the perfect gas assumption. Feed streams, product streams and external heat loads are omitted from the equations above. They may be easily included by augmenting the appropriate mass and energy balances (6.1 to 6.3). The resulting system has $2NC+7$ equations in $2NC+7$ variables for each equilibrium stage. Equation 6.1 only appears $NC-1$ times to avoid an inconsistency with equation 6.4. This model will be referred to as the CHE model, since it uses composition, holdup and enthalpy as the state variables.

6.2 EXTENSIONS OF THE FULL-ORDER CHE MODEL

The complexity of the CHE model has not deterred some workers from investigating the validity of some of the assumptions used in the formulation.

Assumptions 3 and 4 are challenged by Rademaker, Rjinsdorp and Maarleveld (1975) with a rather unusual model which includes a finite vapour holdup. Since the vapour density is closely coupled to both pressure and temperature, the extra term is a combination of vapour temperature and molar holdup. The model also uses tray pressure as the third state variable instead of enthalpy. This is perfectly acceptable because the pressure affects the bubble point of liquid which in turn determines temperature, and enthalpy. They show that the time constant of vapour holdup, τ_v , is due to thermal coupling (interphase thermal resistance) between the tray liquid and vapour. They also show that τ_v is simply related to the pressure time constant τ_p by the empirical formula:

$$\tau_v \approx 4 \tau_p (P_n - P_{n-1}) \quad (6.11)$$

Now, in normal column operation, the pressure difference between plates is considerably less than 0.01 bar, and τ_p itself is small (see later) so τ_v can normally be completely neglected. This is tantamount to assumption 4 above. Although Rademaker et al only consider linear models, it seems reasonable to extend this result to all systems.

Assumption 5 has been examined by Gautreaux and O'Connell (1955) for steady state column models and Baber and Gerster (1961) for simple dynamic models. Imperfect liquid mixing gives rise to a composition profile along a tray which will radically alter the apparent Murphree efficiency. Gautreaux and O'Connell show that the point efficiency E_p is related to the apparent plate efficiency E_m for a linear equilibrium system by

$$E_m = \frac{L}{mV} \left[1 + E_p \left(\frac{mV}{kL} \right)^k \right]^{-1} \quad (6.12)$$

where the liquid has been split into a cascade of k perfectly stirred pools and m is the slope of the vapour-liquid equilibrium line. This formula shows that the point efficiency will always be lower than the plate efficiency - indeed, plate efficiencies greater than unity are readily obtained.

For steady state simulation the plate efficiency may be modified to correct for the difference, or simply taken from experimental data. The effect of this on the dynamic response was reported by Baber and Gerster (1961) for a two-pool five tray CMO model (see later) to be insignificant. Although the result for such a simple example cannot be regarded as conclusive, it has been widely used to justify assumption 5 in later work.

6.3 SIMPLIFIED FULL ORDER MODELS

The comprehensive model given in section 6.2 is far too complex for the current work, but extensive research efforts over the last two decades have provided detailed evaluation of simpler alternatives. The three useful variants are obtained by neglecting accumulation of enthalpy, holdup, or both.

CE model; assume no mass accumulation. Replace 6.2 with

$$0 = L_{n-1} - L_n - V_n + V_{n+1} \quad (6.13)$$

CH model; assume no enthalpy accumulation. Replace 6.3 with

$$0 = L_{n-1} h_{n-1}^L - L_n h_n^L - V_n h_n^V + V_{n+1} h_{n+1}^V \quad (6.14)$$

C model; no enthalpy or mass accumulation. Replace 6.2 with 6.13 and 6.3 with 6.14.

It is important to realise that one is not equating the time derivative of the mass or enthalpy to zero. This would imply that the mass or enthalpy on each tray is constant, when it is actually determined algebraically.

The final important full-order model is the Constant (or Equal) molar overflow, abbreviated to CMO. This is a C model in which the hydraulic equations are replaced by constant flow assignments, and equimolar interphase component transfer occurs. The describing equations are 6.1, 6.4, 6.5 and 6.6. Although this differs considerably from the previous full order types it is always classified with them by other workers.

There is considerable variation in the literature over how models should be classified and what assumptions may be made. Popular assumptions are - neglecting downcomer effects, constant relative volatility, theoretical trays (efficiency = 1), and constant pressure.

This last one deserves special mention because it is only recently that a constant pressure CHE model has been reported to be directly equivalent to a CH model. Prokopakis and Seader (1980) derive a CH model from the CHE equation set by assuming that the vapour composition and enthalpy are not functions of pressure. It can be seen that if the pressure is constant, any pressure dependent terms become constants, so an algebraic expression for enthalpy derivatives can always be derived.

Peiser and Grover (1962) use a CH model for the numerical simulation of a distillation column processing a 4-component mixture. Their model has been reported as a CHE model in other literature surveys, although the authors themselves state that the enthalpy balance is algebraic, as explained above. Oscillatory hydraulic behaviour is reported which confirms later observations.

Levy, Foss and Grens (1969) compare the CHE, CH and CMO models of a binary 10 tray column using the eigenvalues and eigenvectors of the linearised system. The CMO model is reported as a C model, since it appears that the abbreviations were not defined then. The dominant modes of all examples are composition dependent, with small contributions from the other states, lending support to the CMO model. One anomalous slow enthalpy based mode is due to the thermal capacity of the condenser accumulator drum. The authors show that this mode will not be strongly excited by any normal operating disturbances. Complex eigenvalues only occurred in predominantly hydraulic modes and are regarded as purely mechanical phenomena.

Sourisseau and Doherty (1980) provide a direct comparison of variants of the five full-order models in numerical simulations. The C and CMO models assume constant pressure, while the CE, C and CMO models assume constant holdup. The time savings in neglecting enthalpy effects are very significant, and the C model performs adequately if interest does not extend beyond composition effects. For the test problems, the CMO model compared favourably with the C model; however, the small time

saving was not justified by the loss of accuracy.

Gani (1980) also used numerical simulation to compare the CH and CHE models, of a 6-tray column separating a 4 hydrocarbon mixture. The CH model is constructed with constant tray pressure, and two adverse effects of this assumption are illustrated. Firstly, the new steady states are not in agreement following a step increase in reboiler duty, because the enthalpy is computed at the previous constant pressure. Secondly, the hydraulic oscillatory modes are strongly excited in the CH model but not in the CHE model. This probably occurs because a vapour flow change is not subject to any damping through pressure interactions as it propagates up the column. At the top, it condenses to give a surge of liquid which sets the tray liquid oscillating as it returns.

Doherty and Perkins (1982) consider the qualitative properties of the full-order models to identify features of the dynamic equations which give rise to multiple steady states, whose importance is covered in chapter 2. They show that multiple steady states are a consequence of both multiple stages and multiple components. An azeotropic CMO problem is cited as an example. Other proofs are mentioned which support the theories of Levy et al on complex hydraulic modes.

6.4 REDUCED ORDER MODELS

The simplest full order model is still excessively complicated for the study of switchability through optimal control, so further means of simplification must be found. Reduced order models, including lumped section or collocation models, provide a suitable route. The term 'reduced order' covers any model where the number of equations is not directly related to the number of stages, and the techniques are applicable to any stagewise separation scheme.

A large number of reduced order models have been proposed for distillation columns and the following survey will only serve to illustrate the range of models available. The common element in all these is that the process is regarded as a group of 'black boxes' each possessing rudimentary dynamic characteristics and representing a significant part of the column. The overall column response is generated by suitably combining the output from each box.

The major contribution of reduced order model research has been towards synthesising transfer functions, since these are readily derived from either experimental measurements or linearised full order models. Extensive investigations have produced useful theoretical and empirical relationships between model parameters and operating conditions.

Moczek, Otto and Williams (1965) present an empirical model with instructions for estimating the two lag terms and single pure delay term from plant data. Wahl and Harriot (1970) and Waller (1979) use full-order CMO models to estimate time constants and steady state gains for a third order transfer function model. Tables and graphs are used to correlate key parameters with operating conditions.

Osborne and Maddox (1965) split the column at the feed tray and model the stripping and rectifying sections as packed columns. The transfer of a component between vapour and liquid streams is then determined by the driving force within the column, which is found from equilibrium relationships. This has more potential for modelling complex columns although still restricted to small deviations from a nominal steady state. Sullivan (1978) used a commercially developed model based loosely on this principle, but achieved a wider operating range by adopting a mass transfer formula suggested by Shinskey (1967). The empirical nature of Sullivan's model makes it difficult to interpret the state variables as physical quantities. Figure 6.2 shows how the column is analysed.




Figure 6.2 Lumped section binary column

Dahlquist (1980) uses a slightly different method of lumping the column sections together, see figure 6.3. Each node where an external stream meets the column is allocated a holdup of one tray, and the interconnecting column sections absorb the remaining holdup of the column. Each section is modelled as a CMO section with constant K values. The model is presented without supporting evidence for its accuracy, but in the associated application, namely for developing an adaptive control system, it appears satisfactory.




Figure 6.3 Dahlquist Column model

Benallou, Seaborg and Mellichamp (1982) propose a model which is similar to the Dahlquist model but uses more elaborate techniques to subdivide the column. In particular, the column section boundaries are no longer restricted to interconnecting feeds, and the section conditions are represented by a single tray within each section, called the 'sensitive stage'. The column dynamic characteristics are expressed by differential equations which only involve the sensitive stages of each section, and the conditions on any specific tray are found by interpolation. A considerable body of evidence is presented to show that steady states are correct and that overall accuracy is preserved.

6.5 NODE AND SECTION COLUMN MODEL

Although the reduced order models of Dahlquist and of Benallou et al discussed above offer considerable simplification over the full order models, the requirements for this work are not fully satisfied. The model must fulfil the same requirements as for heat exchangers, which are,

1. The steady states must be accurately represented.
2. The dynamic response must be qualitatively correct.

The sensitive stage of Benallou et al is used to represent the holdup of the entire column section, and the location of the sensitive stage is directly equivalent to the term e in the classical double convection heat exchanger model (section 5.2). As shown, this can introduce an inverse response characteristic, which violates requirement 2. One of Benallou's examples clearly illustrates this phenomenon. The inverse response is recognised by the authors, although the cause is not identified.

The results of Levy et al suggest that only the first two or three eigenvalues are sufficiently slow to affect the qualitative response, so in theory only two or three state variables would be needed to model the column (since each state gives rise to one eigenvalue). On this basis the Dahlquist model is unnecessarily complex.

For these reasons it was decided to develop a model specially for this work. The objectives were as follows:

1. Steady states must be accurately represented.
2. The dynamic response due to the first 3 eigenvalues (in a simple binary column) must be qualitatively correct.

3. The model must be able to represent complex columns with side strippers, multiple feeds etc.
4. The model must handle multicomponent mixtures.
5. The number of equations seen by the equation solves must be the minimum number consistent with the previous objectives.

The resulting model uses the essential features of both the Dahlquist and Benallou models. The points where feed or product streams join the column are designated 'nodes', as Dahlquist does, and the mixing and splitting of streams only occurs at these nodes. This will introduce the minimum number of equations consistent with mass, energy and component balances. The nodes are linked by 'column sections', which transfer material and energy between the nodes. The 'sensitive stage' of Benallou is split between the terminating nodes, avoiding any inverse responses. The conditions in each section are computed by a procedure from the composition of the adjoining nodes, thus introducing only one extra equation per column section per component to describe the separation properties of the column.

Other assumptions within the model are consistent with the CMO model; the accumulation of mass is zero (constant holdup) and the changes in enthalpy are due to composition alone (equal molar overflow). To reduce computation in the column section procedure, the model uses constant relative volatility and theoretical trays.

For simplicity the model will be presented for a simple binary column. Extensions to multicomponent complex columns are given in subsequent sections. The physical arrangement of the nodes and sections is illustrated in figure 6.4 overleaf.




Figure 6.4 Simple binary node and section column

Mass balances

$$V_R = L_R + D \quad (6.15)$$

$$F + V_S + L_R = L_S + V_R \quad (6.16)$$

$$L_S = V_S + B \quad (6.17)$$

Enthalpy balances

$$\lambda V_R + T_2 V_R C_p = T_1 D C_p + T_2 L_R C_p + Q_C \quad (6.18)$$

$$\lambda V_S + T_F F C_p = \lambda V_R + T_2 F C_p \quad (6.19)$$

$$L_S T_2 C_p + Q_R = V_S T_2 C_p + \lambda V_S + B T_3 C_p \quad (6.20)$$

Component balances

$$H_1 \dot{x}_1 = V_R y_R - L_R x_1 - D x_1 \quad (6.21)$$

$$H_2 \dot{x}_2 = V_S y_S + L_R x_1 + F x_F - V_R y_R - L_S x_2 \quad (6.22)$$

$$H_3 \dot{x}_3 = L_S x_2 - V_S x_3 - B x_3 \quad (6.23)$$

where

D is the distillate rate

B is the bottoms (residue) rate

F is the feed rate

λ is the latent heat of vapourisation

C_p is the specific heat of liquid and vapour

Q_R is the reboiler heat duty

Q_C is the condenser cooling duty

Subscripts S,R refer to stripping and rectifying sections respectively and all other symbols remain as defined in section 6.1. As stated in section 6.1, only $nc-1$ composition variables are necessary to define an nc component mixture.

The vapour compositions y_R , y_S are found as functions of the liquid and vapour flows in the respective sections, the liquid composition of the adjoining nodes, and the relative volatility of the mixture. The temperatures T_1 , T_2 and T_3 of each node are functions of the liquid composition, relative volatility and pressure, computed as the bubble points. Procedures COLUMN and BUBBLE have been produced to find these values as part of the SPEED-UP process simulation package which is described in section 4.6. The sections have no holdup themselves - all composition holdup in the sections is redistributed equally between adjacent nodes. There is no evidence for a better system, and it is compatible with the suggestions of Benallou et al, who place the sensitive stages in the centre of the main column sections.

6.6 COLUMN - COLUMN SECTION PROCEDURE

The vapour composition leaving the column section is determined by solving the steady state CMO equations for the section assuming constant relative volatility throughout. Consider the column section shown in figure 6.5.




Figure 6.5 Conditions in a column section

The mass balances around each tray give rise to N simultaneous algebraic equations for NC components

$$0 = x_{i,n-1} - (1 + A_i B_n) x_{i,n} + A_i B_{n+1} x_{i,n+1} \quad (6.24)$$

where A_i is the relative volatility of component i with respect to a reference component, normally the heaviest, and B_n is given by

$$B_n = \frac{V}{L} \cdot \frac{1}{\sum_i A_i x_{i,n}} \quad (6.25)$$

This set of $N \times NC$ equations is non linear, but if the dependency of

B_1 on x is ignored, they form NC sets of N linear equations with a tridiagonal jacobian. Such systems are efficiently solved by Gaussian elimination. Parameter B can then be found by successive substitution. The algorithm used in COLUMN is as follows;

1. Estimate $\underline{x}_n$ for all N by linear interpolation between $\underline{x}_0$ and $\underline{x}_{N+1}$

2. Compute B_n for all N using 6.25

3. Solve 6.24 for $\underline{x}_n$ with constant B_n

4. If

$$\sum_n \left| \sum_i x_{i,n}^{-1} \right| < \text{error tolerance}, \quad (6.26)$$

then go to step 7

5. Compute new B_n from

$$B_n = \frac{V}{L} \frac{\sum_i x_{i,n}}{\sum_i A_i \cdot x_{i,n}} \quad (6.27)$$

6. Return to step 3

7. Compute $\underline{y}_1$ from

$$y_{i,1} = A_i \cdot B_i \cdot x_{i,1} \quad (6.28)$$

8. Stop

The error tolerance (6.26) was $2 \cdot 10^{-6}$ in this work. The algorithm is similar to the one used by Eliceche and Sargent (1981) for steady-state column design. The parameter B above corresponds directly with temperature on each tray which in Eliceche and Sargent is the torn variable. This choice of torn variable precludes the use of a general equation solving method to speed up convergence in this procedure. The jacobian matrix for this system has NP squared elements, which for a section of 20 trays would require 21 steps to initialise by finite

difference. The algorithm above will usually converge within this time. An alternative tear set, the composition of the bottom tray, $\underline{x}_0$, will only need a jacobian of $NC-1$ squared elements. The following algorithm could then be used;

1. Set $\underline{x}_N = \underline{x}_{N+1}$
2. Set $i = N$
3. Solve (6.24, 6.25) for $\underline{x}_{i-1}$
4. Decrement i
5. If $i > 1$ go back to step 3
6. Solve (6.24, 6.25) for $\underline{x}'_0$
7. Compute residuals $\underline{f}(\underline{x}_N)$ from

$$\underline{f}(\underline{x}_N) = \underline{x}_0 - \underline{x}'_0 \quad (6.29)$$
8. Apply general purpose equation solver to find new $\underline{x}_N$ so that

$$\underline{f}(\underline{x}_N) = 0 \quad (6.30)$$
9. If equation solver has not converged, go back to step 2.

This algorithm was implemented using both the secant method, and a Quasi Newton method (see section 4.6) for step 8. The results for a series of column sections are shown in table 6.6 overleaf. The sections are taken from the networks in chapter 8.

Column Section	Successive substitution		Torn x_N , Secant		Torn x_N , QN	
	iterations	time	iterations	time	iterations	time
A1R	7	.082	31	.179	9	.102
A1S	12	.079	62	.148	10	.096
A2R	9	.119	62	.366	10	.111
A2S	12	.054	51	.099	10	.095
B1R	4	.047	11	.069	7	.087
B1S	56	.352	100*	.287	15	.141
B2R	4	.041	9	.047	9	.093
B2S	52	.386	100*	.354	15	.152
C1R	2	.016	100*	.286	20*	.189
C1M	11	.075	45	.150	10	.097
C1S	11	.055	42	.091	10	.093
C2R	5	.065	13	.085	8	.097
C2S	4	.027	5	.015	6	.056
TOTAL(CP seconds)		1.398		2.176		1.359

* failed to converge

Table 6.6 Comparison of column solution methods

It is apparent that the successive substitution method works surprisingly well, and converges in every case. The small improvement of the QN method is negligible, especially as it failed to converge on one of the examples.

The algorithm as given will only handle an integral number of trays. While satisfactory for simulation it is unsuitable for column design where the number of trays is initially unknown. This problem is avoided in the procedure by linear interpolation between the nearest integer tray counts. An alternative method would be to modify equation 6.24 to allow for non-ideal trays. Unfortunately this would destroy the symmetry of the jacobian, necessitating special solution methods.

6.7 BUBBLE - BUBBLE POINT PROCEDURE

The Antoine pressure-temperature correlation is used with the definition of constant relative volatility to compute the bubble point of the tray liquid. If we assume that tray liquid and vapour are in thermal and material equilibrium, the bubble point is equivalent to the tray temperature. At a given pressure and temperature, the K-values for a key component are related to the partial pressure by (6.31),

$$K_k = \frac{P_k}{P} \quad (6.31)$$

and the relative volatility A_i of any component i against the key component k is (6.32),

$$A_i = \frac{K_i}{K_k} \quad (6.32)$$

The gas phase mole fraction y_i is known from the liquid mole fraction by (6.33),

$$y_i = K_i x_i \quad (6.33)$$

For an ideal gas mixture, the total pressure is the sum of the partial pressures of each component, which is proportional to the fraction of that component in the mixture. From this, and equations 6.31 to 6.33, it can be shown that,

$$P_k = \frac{P}{\sum_i A_i x_i} \quad (6.34)$$

Rearrangement of the Antoine equation for the vapour pressure of the key component gives

$$T = 273.15 - c - b/(10^{P_k} - a) \quad (6.35)$$

where a, b, c are the Antoine coefficients of the key component.

6.8 EVALUATION OF THE NEW MODEL

The quality of any simplified column model can only be assessed by comparing it with a rigorous model over a wide range of design and operating conditions. The time available for this work does not permit such a thorough investigation so a single column has been used, albeit over a significant operating range.

The example chosen is a 5-tray column described by Gani (1980) for the separation of a 4 component hydrocarbon mixture. The physical arrangement of the column remains unchanged (figure 6.7) but the mixture to be separated now has only two components, benzene and cyclopentane. These were chosen mainly because the physical properties were available on the SPEED-UP database. The full model is a CHE model, named CHE25, using the equations given in section 6.1. The SPEEDUP implementation of this model (Appendix G) has a separate unit for each tray, so columns of any size may be modelled without difficulty. Vapour-liquid equilibrium data and physical properties of the fluids are computed by the Chao-Seader correlation.

It may seem inappropriate to include level and pressure control loops but they are necessary for a practical implementation of the model. Levy et al assume constant pressure and constant holdup, which may be regarded as 'perfect control' within the loops shown. It is considered more realistic to use finite control gains and accept that this may modify the system behaviour slightly. (As it happens, only the fastest modes are affected).

The first stage in the study was to compute the steady state for CHE25 at the nominal operating point. The eigenvalues of the full model were then found, following the technique of Levy et al (1969), to check the accuracy of the formulation. Full details of the column construction and steady state are given in appendix C.




Figure 6.7 Model CHE25 flow diagram

The eigenvalues spanned a range of -1.6×10^6 to -0.52 at the nominal steady operating state, and were in good agreement with Levy et al. Figure 6.8 shows how they related to the physical model. The classification is based on the contribution of the state variables to the eigenvectors.




Figure 6.8 Eigenvalues of CHE25

The hydraulic modes contain two complex pole pairs, which is considered normal. The anomalous slow thermal mode is due to a non-equilibrium condenser accumulator. This feature also appears in Levy, and is unavoidable if the accumulator is to be correctly modelled (n.b. the literature contains several examples of equilibrium accumulators, where the overhead system is simply lumped in with the top tray). All other thermal modes are developed from pressure difference terms and equilibrium vapour pressures, so will be very fast.

For the purposes of this work, only the slowest of the composition modes are likely to affect switchability, as the others will rapidly converge to the steady state values. The condenser accumulator mode, although significant, shows very little interaction with composition modes and can probably be lumped into the following unit if necessary.

The simple model, CLTEST, was constructed from three nodes and two column sections, in accordance with the equations in section 6.5. Vapour and liquid flowrates were taken from the steady state solutions for CHE25.

Before any comparisons can be drawn between the two models, the number of trays in the simple model has to be found to give compatible steady states. The full model has trays with a Murphree efficiency of 0.8 and a theoretical reboiler, with the added complication of the relative volatility changing within each section. The SPEED-UP package was used to solve the design problem of finding the number of theoretical trays given the arithmetic mean relative volatility and the node compositions. The result is shown in table 6.9, and compares well with the rule-of-thumb which multiplies the tray efficiency by the physical tray count. (The figures do not include the nodes, which are treated as theoretical trays).

Having established the physical characteristics of the simple model, it is possible to examine how the dynamic and steady state solutions compare with the full model for a range of feed compositions. The steady state compositions and eigenvalues of the composition modes are shown in table 6.10. The agreement between the two sets of figures is regarded as fortuitous, but it shows that the new model fulfils the design objectives. The fastest mode shows the greatest error, but even so this is below 20 per cent, and the inaccuracy will be short lived as the mode decays quickly to the steady state. The same applies to higher modes which the dimension of the simple model is too low to exhibit. The condenser accumulator thermal mode is not supported by the simple model at all, although Levy et al regard it as unimportant, see section 6.3.

The final comparison is based on efficiency of computation, measured as the number of equations and as solution speed on a CDC 6600 computer. Although this comparison is very dependent on the starting point and numerical properties of the formulation, it shows a clear advantage for the simple model, which easily outweighs the loss of accuracy

Section	Physical Trays	Theor. Trays	Design Soln.
Rectifying	2	1.4	1.33
Stripping	2	1.6	1.83

Table 6.9 Column section sizes for CLTEST

Parameter	Feed Composition					
	0.3		0.5		0.7	
MODEL Name	CHE25	CLTEST	CHE25	CLTEST	CHE25	CLTEST
x_T (Molfraction)	.525	.525	.793	.793	.927	.925
x_F (Molfraction)	.168	.171	.417	.417	.714	.717
x_B (Molfraction)	.0258	.0258	.101	.102	.368	.370
Separation factor	41.9	41.7	34.1	33.7	21.8	21.0
λ_1 (Eigenvalue)	-.73	-.72	-.52	-.56	-.54	-.54
λ_2 (Eigenvalue)	-2.62	-2.88	-2.07	-2.07	-2.00	-1.78
λ_3 (Eigenvalue)	-8.14	-6.52	-6.4	-5.03	-5.7	-4.42

Table 6.10 Steady state compositions and eigenvalues of major dynamic modes compared

MODEL	No. of equations		Steady State (QN)	Dynamic (Sparse BDF)
	TOTAL	DIFFERENTIAL		
CHE25	187	21	80	400
CLTEST	14	3	1.6	3

Table 6.11 Run times in CDC 6600 CPU-seconds

CHAPTER SEVEN

HEAT EXCHANGER NETWORK CASE STUDY

7.1 NETWORK DESCRIPTION

This chapter compares the switchability of two heat exchanger networks, which are alternative solutions to a standard steady state heat recovery problem. The aim is to assess the effect on switchability of the following features;

- (i) Heat recovery effectiveness.
- (ii) Bounded control action.
- (iii) Closed loop control properties.

The classical steady state heat exchanger network design problem (Flower and Linnhoff (1980), etc) defines a single set of stream properties and temperatures, and the resulting designs may only be valid at one operating point. In this work, we are interested in switching between two or more feasible steady operating points so we must ensure that the networks are flexible - in other words, they can tolerate changes in operating conditions. The networks in this case study are solutions to the four-stream flexible network design problem given in table D.1 and reported by Beautyman (1984). These networks achieve the specified outlet temperatures over the full disturbance range using the minimum number of manipulated variables and at minimum overall cost subject to any special constraints. Two solutions, illustrated in figures 7.1 and 7.2, arise from different network structural constraints. Full numerical details of these networks are presented in appendix D. See Beautyman (1984) for a detailed account of the flexibility problem and description of the design method and cost estimation.

The first network, identified as HEX7, offers maximum energy recovery at the nominal design point using the minimum total heat exchange area and a minimum approach temperature within any exchanger of 36°F. The temperature control on streams 1,2,3 and 4 is through bypasses/utility flows on U1, U2, H1 and U4 respectively. Note that the control on stream 3 takes effect indirectly via H3.




Figure 7.1 Network HEX7

The second network, HEX8, sacrifices some energy recovery to reduce the total area and the number of exchangers. The temperature control on streams 1, 2, 3 and 4 is through bypasses/utility flows on H1, U2, H3 and U4 respectively.




Figure 7.2 Network HEX8

For the switchability study, the heat exchangers and utility exchangers were modelled as described in chapter 5. Bypasses on the controlled stream were used as recommended by Shinskey (1979) p230, as being both a robust and efficient means of temperature control. The effective thermal holdup of each unit was estimated from the overall heat transfer coefficient - area product. A total of 68 equations for HEX7 and 52 equations for HEX8 were required to represent the entire network and ancillaries in SPEEDUP (Chapter 4).

The models satisfy the conditions of theorem F (chapter 2) so the networks are unconditionally switchable.

7.2 DETERMINING THE SWITCHES AND OBJECTIVE

The index of switchability defined in chapter 3 will be dependent on the initial and final states of the switch. Since we wish to compare the networks for switchability we must emphasise the difference in switch cost by choosing the most difficult switch. It is reasonable to assume that this will be the one involving the greatest transfer of heat within the network. The model which represents each heat exchanger has one accumulation term in temperature on each stream flowing through it. The change of heat content of such an accumulator will increase with the temperature change during the switch. The total amount of heat transferred ΔE can therefore be estimated by considering all accumulation temperature changes throughout the network, according to (7.1).

$$\Delta E = \| \underline{T}_1 - \underline{T}_2 \| \quad (7.1)$$

where $\underline{T}_1$ is the vector of accumulation temperatures at state 1, and $\| \underline{T} \|$ is the Euclidean norm of $\underline{T}$, defined in chapter 3. The most difficult switch will then be a switch between $\underline{T}_1$ and $\underline{T}_2$ in (7.2);

$$\Delta E_{\max} = \max_{\underline{T}_1 \in T} , \quad \max_{\underline{T}_2 \in T} , \quad \| \underline{T}_1 - \underline{T}_2 \| \quad (7.2)$$

where T is the set of all feasible steady state accumulation temperatures of the network over the full range of operating conditions (flowrates), while maintaining the specified outlet temperatures.

As there are four streams, each with infinitely variable flowrate, the search space T is quite large. To reduce this space and render the problem manageable, a further assumption is that the solutions of (7.2) lie on the boundary of T . This would certainly be true of a linear system, in which the extrema must lie at the vertices of the bounded solution space. In practice this means that the flowrates will all change between minimum and maximum values during the switch.

There are thus 16 possible switches of flowrates, arranged as 8 pairs since T_1 and T_2 are interchangeable in (7.1) and (7.2).

This entire set of switches was checked on network HEX7. The solution of (7.2) occurs when both hot streams change from maximum to minimum flowrate and both cold streams change from minimum to maximum flowrate. This switch causes an overall reduction of temperature throughout the network. Furthermore one of the states (state A in table D.2) defines the overall heat exchanger area in the steady state network design (Beautyman, 1984). Neither fact is surprising, since we are seeking the limits of the operating range of the network.

The regulation of temperatures under flow disturbances is only one aspect of switching; changes of feed and target temperatures must also be considered. The original problem definition (Table D.1) does not include temperature changes, so an arbitrary change of -10°F has been chosen for all hot end temperatures. Such a change could be induced by, for example, a pressure reduction in an associated distillation column. Three states (A,B,C) have thus been defined, and are detailed in table D.2. Switches between A and B are the regulatory switches, and between A and C are the servo switches, making a total of four for each network.

All utility costs are expressed as the equivalent cost of the hot utility, which will usually be the most expensive. The cold utility was charged at 0.2 of this. Following the procedure described in chapter 3, the relative weighting for off-specification product was found from the effect on target temperatures of a cutback in utility usage. The technique adopted was to shift the setpoints so that the utility consumption fell slightly and observe the resulting utility saving. Only HEX7 (A) was examined since it is the most energy-efficient. A deviation on each stream of 0.1°F was found to save the equivalent of $0.004 \text{ tons min}^{-1}$ of hot utility. This corresponds to a quadratic weighting factor of 0.1, and an Euclidean norm deviation of 0.2°F .

The final objective function for all networks is given by (7.3);

$$J = \int_{t_o}^{t_f} (0.1 \sum_i (T_{\text{error},i})^2 + 0.2(F_C - F_{C,ss}) + (F_H - F_{H,ss})) .dt \quad (7.3)$$

where T_{error} is the output stream temperature error, °F,
 F_C is the total cold utility flow, Tons min⁻¹,
 F_H is the total hot utility flow, Tons min⁻¹,
 ss is the target steady state value, and
 i refers to the stream number (1-4).

7.3 CLOSED LOOP CONTROL

A closed loop control policy was found in order to verify the network model and to establish an initial trajectory for the optimal control package. The nature of the manipulated variables permits simple proportional plus integral controllers to be implemented on each stream. Controller tuning will affect the resulting control policy, so a consistent method was required for both networks. The 'Reaction curve' method of Cohen and Coon (1953) was applied to estimate the controller gains. This uses the observed step response of a system to compute the equivalent delay time and dominant time constant of the transfer function. The simple heat exchanger models do not exhibit any pure delay term; for the purposes of control tuning, this was estimated from the examples of chapter 5 to be one fifth of the main time constant. The controller gains were then found for each control loop using the appropriate formula. Tables D.5 and D.8 in appendix D contain the resulting control tuning parameters.

Having established a control scheme, we must consider the effect of bounds on the manipulated variables. In classical control theory all control elements are linear, so the control action available is unbounded. This restriction is satisfactory when the disturbances are small, but in the current work we are considering exclusively disturbances which are large. Under these conditions the control action will not be linear because of the physical limitations of the equipment. A valve cannot close past 0% or open beyond 100% of maximum flow.

The direct solution is to constrain the achieved position of the control element; the output of such a constraint function appears below (fig. 7.3).




Figure 7.3 Inequality control constraint function

Unfortunately such a function is discontinuous in first and higher derivatives at the bounds. The numerical integration routine (See chapter 4) is unable to step over such discontinuities efficiently. The prediction function will not be valid across the discontinuity since it uses a Taylor series expansion, which assumes continuity of function and derivatives; the corrector may not converge if the jacobian estimate is wrong. In consequence, the time step is progressively reduced until the predictor error across the discontinuity is of similar magnitude to the specified tolerance. This reduction and subsequent restoration of time step length is wasteful of computing resources.

The above-mentioned problems may be alleviated if the first derivative is continuous throughout the range of the control. A suitable function (7.4), plotted in figure 7.4, has proved satisfactory for all control limits in this work.

$$\begin{aligned}
 y &= 0.1 \exp(10x - 1), & x < 0.1 \\
 y &= x & 0.1 \leq x \leq 0.9 \\
 y &= 1 - 0.1 \exp(10(1-x) - 1), & x > 0.9
 \end{aligned}
 \quad \left. \vphantom{\begin{aligned} y &= 0.1 \exp(10x - 1), \\ y &= x \\ y &= 1 - 0.1 \exp(10(1-x) - 1), \end{aligned}} \right\} (7.4)$$




Figure 7.4 Graph of function (7.4) near origin

The effect of applying closed loop control for each of the four switches on HEX7 and HEX8 is plotted in figures 7.5-7.8. These show the control action (percent of maximum) and setpoint error ($^{\circ}\text{F}$) against time from the start of the switch. The cost of operation for these switches is shown in table 7.13. It can be seen that the closed-loop control can achieve a smooth switch without difficulty in both networks. Extreme control action is only evident in the servo switches where an abrupt setpoint change is required (streams 3 and 4 of switches A-C and C-A). This is discussed in detail later.





Figure 7.5 HEX7 Closed loop control A/B





Figure 7.6 HEX7 Closed loop control A/C




Figure 7.7 HEX8 Closed loop control A/B





Figure 7.8 HEX8 Closed loop control A/C

7.4 OPTIMAL CONTROL

The SPEEVOP package (chapter 4) will attempt to find the optimal control policy using piecewise constant control action. Initially the heat exchanger network switching problems were presented directly to SPEEVOP to find such a control policy for the switches. It was soon apparent that although an optimal solution could be found, it would not represent the optimal control policy using continuous controls, due to the error caused by discretisation. Within each control interval, the outlet stream temperature would typically start 1° high of the target temperature, and fall almost linearly to 1° below it by the end. This pattern would be repeated in every interval except the last, giving a saw-tooth shape to the temperature error during the switch. The only way to smooth out this error is to increase the number of control intervals. This approach was considered impractical on numerical grounds.

A much smoother control action was obtained by the closed loop control system - it was decided to adapt this control system for optimal control. The SPEEVOP package was used to find the optimal piecewise constant setpoint control for the closed loop system. The control action is determined explicitly by the difference between the actual and target temperatures, so by suitably adjusting the latter, any desired control action may be achieved. It is easily seen that the resulting optimal control will never be inferior to the closed loop control, irrespective of the number of control intervals.

As pointed out in chapter 3, the presence of closed loop controllers must detract from the true optimal control policy, because the associated controls are not available for manipulation. However the dimension of the optimising control vector is unaffected in this case, because each manipulated variable has a corresponding setpoint. The optimal control policy may still be constructed by SPEEVOP.

The SPEEVOP package was thus used to solve all the switches previously

examined using closed loop control. The control policies and output temperature error for each switch are shown in figures 7.9 - 7.12, using the same notation as for the closed loop results. The cost of operation during each switch is shown in table 7.14

In all cases the optimal control is very similar to the closed loop control policy, and extreme control action only occurs under similar conditions. The optimal control costs are substantially less than the closed loop costs, showing that the objective function is relatively sensitive to the control variables.

The optimisation process was carried out with a higher specified accuracy than necessary, to avoid premature termination of the algorithm. The termination conditions are discussed in chapter 4 section 5. The close temperature regulation throughout the switches is an indication that these are globally optimal. The resulting accuracy of the optimal solutions was inferred by examining the progress of the objective value as better than 0.001 absolute in all cases.





Figure 7.9 HEX7 Optimal control A/B





Figure 7.10 HEX7 Optimal control A/C





Figure 7.11 HEX8 Optimal control A/B





Figure 7.12 HEX8 Optimal control A/C

7.5 PERFECT CONTROL

Perfect control of the product temperatures of the heat exchange networks is defined as a control policy which maintains these temperatures exactly at the setpoints for all time. Such a control policy is not generally realisable in practice because of instrumentation and control limitations. Perfect control and related topics are discussed in Wong (1984) and Bristol (1981). In some circumstances, perfect control provides a mathematical tool to measure the utility cost of a switch independently of any artificially imposed penalties on the product specification. For the heat exchange networks under consideration, only the regulatory switches (A-B and B-A) are suitable for perfect control, because the control action is always well within the bounds.

The perfect control problem can be solved using the same set of equations as used to model the network for closed loop control. The difference lies in the selection of known and unknown variables. In the closed loop case, the controls were defined by the controllers, and the outlet temperatures were allowed to vary. Perfect control defines the outlet temperatures, so the controls must be allowed to change. The control action necessary to maintain the outlet temperatures at the target values may be computed by a simulation of the switch with constant outlet temperatures. This problem is readily set up and solved by SPEEDUP.

The objective function for the perfect control switch cost is exactly the same as previously defined. The quadratic term in the off-specification product will always be zero, as the product has been constrained exactly to specification. The switching cost will then represent the extra utility consumption incurred because the initial state did not coincide with the target state. The switch costs are shown in table 7.15. The control trajectories have been omitted because they are not visibly different from the optimal control trajectories, given in the previous section.

7.6 RESULTS AND DISCUSSION

The costs (in tons equivalent hot utility) of the switching policies derived by the various methods given previously are tabulated below. These costs are not operating costs, which must of course be positive in this study. The actual operating cost for each switch can be found by adding the appropriate steady state operating cost for the same period of operation, which is given in table 7.16. The cost of switching is as defined in chapter 3, and is to be regarded strictly as an index of switchability. A low or negative index indicates a potentially efficient switch, and vice versa.

Switch states	HEX7 switch cost	HEX8 switch cost
A-B	1.624	0.239
B-A	0.208	0.635
A-B-A	1.832	0.874
A-C	6.120	1.822
C-A	4.592	3.012
A-C-A	10.712	4.834
A-B-A-C-A	12.544	5.708

Table 7.13 Closed loop switch costs

Switch states	HEX7 switch cost	HEX8 switch cost
A-B	0.403	-0.389
B-A	-0.632	0.175
A-B-A	-0.229	-0.214
A-C	3.253	0.117
C-A	3.065	1.066
A-C-A	6.318	1.183
A-B-A-C-A	6.088	0.969

Table 7.14 Optimal setpoint control switch costs

Switch states	HEX7 switch cost	HEX8 switch cost
A-B	0.4994	-0.2969
B-A	-0.5278	0.2715
A-B-A	-0.0284	-0.0254

Table 7.15 Perfect control switch costs

Operating state	Cost HEX7	Cost HEX8
A	47.3098	48.2666
B	41.0906	45.7830
C	39.3227	41.2115

Table 7.16 Steady state utility costs

The first point to note from these tables is that the switch cost is relatively small compared to the normal steady-state operating cost - in the optimal regulatory switches (A-B and B-A) it is less than two percent. This suggests that the cost of switching is not an important consideration in heat exchanger network synthesis. A selection on the basis of minimum overall operating cost of the complete set of either closed loop or optimal switches will favour the design with the lowest steady state operating cost. The switching cost index derived in chapter three has been constructed to cancel out the steady state operating cost so that more detailed comparisons can be made.

The optimal control results will be considered first as they are the principal results of the study. There are considerable variations of the individual switch costs in table 7.13 which hinder comparison. In particular, the regulatory switch costs for each network have complementary signs. The switching cost expression makes no allowance for the change of inventory between two states. Since the initial and target states are predefined, this merely introduces a constant term into the switching cost. This term will be cancelled out if the overall change of inventory (which may be energy or material) is zero. Such cancellation can be arranged by considering a cyclic switch -

from A to B, then from B back to A. The cyclic switch costs are tabulated for all pairs of switches as A-B-A, etc.

The cyclic regulatory switch costs are slightly negative for both the optimal switches and the perfect control switches. In both cases the quadratic term is greater than or equal to zero, so the negative cost must arise because the utility consumption during the switch is less than that required for steady state operation. It is easy to see how a change of inventory can supplement the overall energy balance, but the cyclic cost cancels this effect out. There is no reason a priori why the operating costs should not be less during a switch than at either steady state. However the results should be carefully checked to eliminate other causes before reaching this conclusion. A number of possible explanations are considered.

(i) Numerical error

The solution of the model equations will only be accurate to within the tolerance specified on the residual error. Further errors will be introduced by the finite word length of the computer representation of the numbers. This cause was eliminated by repeating the calculation at higher accuracy, until no further change in the result was apparent. This occurred at a point well within the machine accuracy, with negligible difference from the earlier results.

(ii) Terminal state error

The final state may not coincide with the target state, introducing a small error in the cyclic inventory change. This was eliminated in the perfect control case by extending the time horizon until no further change was observed in switch cost. Although a small increase was obtained it did not account for all the observed cost saving.

(iii) Product specification offset

The marginal value objective function described in section 3.2 has to satisfy equation 3.12 to ensure that control action which perturbs the final state away from the target state does not improve the operating cost. The compound objective function of section 3.6, as defined by equation 7.3 relies instead on the quadratic terms to penalise any deviation from the target state. However the remaining linear terms are superimposed on this penalty function and the minimum is not coincident with the target state. The difference between the desired and actual minimum is a function of the relative weighting between quadratic and linear terms.

Consider the case of an objective function which should constrain x to 0, given by $ax^2 + x$. The minimum of this function occurs when $x = -1/2a$. As the weighting a is increased, the offset will reduce but it can never be completely eliminated for finite a . The choice of quadratic weighting given in section 3.6 was intended to limit this offset to an acceptable value.

The offset is indeed present on stream 4 in the optimal control runs, graphs 7.9 to 7.12. The effect is to reduce the utility consumption when possible by deliberately producing off specification product. The error on stream 4 is around 0.25°F , similar to the normalised tolerance of 0.2°F used to compute the weighting factor. The error on the remaining streams is insignificant because the relative costs of utilities are much less. The saving in utility cost caused by a temperature error of this magnitude is around 0.30 tons of hot utility in all cases. This is comparable to the observed reduction in the optimal control cost. The perfect control switch costs have no set-point deviation so the above argument is not applicable to them.

(iv) Intrinsically negative switch cost

The three explanations given above do not account for all the negative costs observed in all cases. This suggests that the system can

operate more efficiently at some intermediate set of states which are encountered during the switch. The following example shows how such a situation can arise.

The use of a cooler and a heater on a single process stream is an energy efficient way of smoothing out bidirectional fluctuations in stream temperature. The principle is employed for example by Tyreus and Luyben (1975) for a reboiler/condenser system. Consider the simple network below (figure 7.17) modelled as described in chapter 5.




Figure 7.17 Heater/cooler combination

Perfect control may readily be achieved for a restricted range of inlet temperatures above and below the target product temperature. The hot utility flow maintains the feed to the cooling exchanger at or above the target temperature. The cold utility flow can then always be set to achieve the target temperature. Let us suppose that the system operates with a high feed temperature for one hour then at a low feed temperature for a further hour. The total utility consumption will be 1 hour each of hot and cold utility at unit flowrate. Now suppose that the feed temperature changes very rapidly between hot and cold, say ten seconds of each, repeated 180 times. The thermal inertia of the first element in the network will limit the variation of feed temperature to a small fraction of the original so the amount of hot and cold utility required will be very much less. However the total amount of product is the same as before. The saving must therefore be due entirely to the large number of switches.

The occurrence of negative switch costs in the tables is therefore satisfactorily explained. The last three causes are probably all contributing a part to these in table 7.14, but only the final one applies to table 7.15. The setpoint offset is acceptable in terms of the offset used to define the weighting factor, and the final state error as shown by the final controls and stream temperatures is also small. The results may therefore be used confidently for the comparison of switchability of the two networks.

The regulatory (A-B-A) switch costs are very similar - in view of the previous discussion, it is impossible to make any distinction between the networks. The temperature is maintained close to the setpoint without incurring any penalty in utility consumption. This is apparent from the perfect control cases. Conversely, there is no significant saving to be made by sacrificing product quality. In this example regulatory switchability is not a significant consideration in the choice of networks.

The servo switches show much greater switching cost, due directly to the unavoidable temperature deviation at the start of the switch. The disparity between networks is also much greater, because stream 3 is not directly controllable in network HEX7. Limiting control action is evident at the start of all servo switches.

Obtaining the optimal control costs was a computationally expensive business, and if equivalent results could be obtained by a simple simulation, it would be more suited to screening of preliminary designs. The closed loop control costs were computed for this purpose and are shown in table 7.13. All switching costs are substantially higher, as expected, but the ranking of the regulatory switch has been reversed. This alone is sufficient evidence to discard simple closed loop behaviour as an indicator of switchability in general.

The product stream temperature responses show that in the absence of limiting control action, both closed loop and optimal control policies were able to maintain the target temperatures without difficulty.

This is a consequence of the bypass control system, which ensures that the dynamic control range of the heat exchanger always equals or exceeds the steady state control range. To illustrate this, consider the simple network shown below (figure 7.18).




Figure 7.18 Bypass control of a single exchanger

The temperature of the controlled stream T_{H3} depend on the fraction u of the total controlled stream flow which passes through the exchanger, and the instantaneous hot stream temperature leaving the exchanger T_{H2} .

$$T_{H3} = u \cdot T_{H2}(T_{H1}, T_{C1}, u) + (1 - u) \cdot T_{H1} \quad (7.5)$$

Let us define the control range as the set of all hot stream temperatures T_{H3} which may be achieved by manipulating control u , subject to constant inlet conditions. The dynamic control range is the instantaneous range of T_{H3} which will clearly comprise all temperatures between T_{H1} and T_{H2} . Hence,

$$T_{H1} \geq T_{H3} \text{ (dynamic)} \geq T_{H2} \quad (7.6)$$

The steady state control range will also have an upper bound of T_{H1} , but the lower bound will be dictated by the maximum heat transfer duty, which occurs when $u = 1$. This also defines the highest temperature of T_{H2} , because if the flow through the hot side of the exchanger is reduced, T_{H2} must fall. Hence,

$$T_{H2} \text{ (max)} \geq T_{H2} \geq T_{c1} \quad (7.7)$$

$$T_{H1} \geq T_{H3} \text{ (steady state)} \geq T_{H2} \text{ (max)} \quad (7.8)$$

From (7.6-8), it can be seen that the dynamic control range will never be less than the steady state control range. This conclusion will apply to any of the heat exchange models discussed in chapter 5 because it is completely independent of the heat exchanger dynamics.

The control of more complex networks will not be so direct because of interactions, but will benefit from this property if the bypass control system is employed. In the case study, HEX8 can control two streams directly by this means, whereas HEX7 has indirect control of all streams. The temperatures are thus subject to a lag (first order in the current models) which will smooth out the rapid change in outlet temperature from the controlled stages. This is identified as the main reason why in general HEX8 maintains closer control of the product streams than HEX7.

The choice of controlled stages is often fixed at the network design stages by the flexibility requirements (Beautyman, 1984). Although HEX7 could be controlled by a bypass on exchanger 3, the area of this exchanger would need to increase to 2.32 units to provide the flexibility required. It would be difficult to justify such an increase merely for switchability.




Figure 7.19 Steady state temperatures vs bypass control




Figure 7.20 Effect of a step change in u from 0.2 to 0.8

The effect of control u on T_{H2} and T_{H3} for the network in figure 7.18 is shown in figure 7.20. The flows and temperatures have been normalised so they lie between 0 and 1, with a heat transfer coefficient-area product of 2. The effect of a step change in u on T_{H3} has been plotted in figure 7.20. These illustrate the phenomena discussed above.

In conclusion this case study has highlighted several facts which are probably applicable to all heat exchanger networks where temperature and flow switchability is desirable.

1. If optimal switching policies are used the operating cost is unlikely to be either increased or decreased significantly by the presence of switches. However, the cost of switching is very sensitive to the exact control policy employed, and in practice the performance of the control system must be accounted for.
2. Current recommendations for temperature control allow target temperatures to be achieved despite considerable flow and temperature disturbances within the network.
3. Designs should be favoured when the bypass controlled stages are on the outputs of the network, providing that the temperature control range is sufficient to cope with the anticipated disturbances.
4. Closed loop control behaviour is not a reliable index of switchability, but optimal setpoint control of a closed loop system is recommended if a detailed switchability study is undertaken.

CHAPTER EIGHT

HEAT-INTEGRATED DISTILLATION SYSTEM CASE STUDY

8.1 SYSTEM DESCRIPTION

The previous case study (chapter 7) concerned systems in which thermal effects were dominant, and the system behaviour was approximately linear. This chapter looks at systems with highly nonlinear responses where temperature changes are only of secondary importance. The system chosen is the high purity separation of methanol and water by distillation in two heat integrated columns. The columns are operated at different pressures so that the reboiler of one is heated by the condensate of the other. This arrangement and associated problem was described initially by Tyreus and Luyben (1975) for comparing closed loop control strategies. It was further developed by Lenhoff and Morari (1980) to study general process network design to optimise dynamic and economic performance simultaneously. Three alternative configurations were analysed using a wide range of simple control schemes which were optimised as briefly described in section 3.3. These same configurations were used for this work, although the operating points have been altered to conform with the column model given in chapter 6. The networks are illustrated in fig 8.1 and further details may be found in appendix E.




Figure 8.1 Distillation networks

Key: _____ mass flow
 ----- energy flow
 * manipulated variable (control)

8.1.1 LENMOA description

The feed stream is split into two and each part is processed independently. The two distillate and two residue streams are then mixed to give the products. The only coupling of this system is through the reboiler-condenser. Four manipulated variables are available; feed split, reflux ratio of column 1, reflux ratio of column 2 and reboiler duty of column 2. The reboiler of column 1 is not independently controllable since it is part of the condenser of column 2. The SPEEDUP model contained 57 equations.

8.1.2 LENMOB description

The feed passes directly to column 2, where the bulk of the water is removed. The enriched distillate is then dehydrated in column 1. The two residue streams are combined. The system exhibits some interaction through the top composition of column 2, as well as the thermal coupling. Only three controls are present since the throughput of column 2 is fixed. They are; reflux of column 1, reflux of column 2 and reboiler duty of column 2. The SPEEDUP model contained 51 equations.

8.1.3 LENMOC description

The feed is partially separated in column 2 and the separation is completed in column 1. This arrangement is attributed to Petlyuk et al (1965) and is potentially the most effective configuration. It is unusual in that the second column has two feed points. Interaction occurs through top and bottom compositions of column 2 and in the reboiler-condenser. The manipulated variables are as for LENMOB. The SPEEDUP model contained 49 equations.

8.2 DETERMINING THE SWITCHES AND OBJECTIVE

The system studied by Lenhoff and Morari operated with a constant feed composition of 0.7 mole fraction methanol, 0.3 mole fraction water. The controller performance was measured correcting an arbitrary small composition disturbance applied simultaneously to all composition holdups in a linearised model of the columns. The integral of the squared product composition deviations and control action (as described in section 3.3 of this work) was used as the dynamic performance index. The controller gains were chosen to minimise this index for each feasible single-input single-output control scheme of each of the three column configurations.

In the present study, the disturbances caused by the switch must be relatively large, to emphasise the nonlinear features of the system. A switch in the final product compositions was discounted because of their comparatively high purity. The small absolute composition change would not involve a major transfer of species to the ends of the column. A change of throughput would not alter the compositions in the system either, because the simple models used neglect the dynamics of flow changes. A feed composition switch was therefore chosen, of 0.1 mole fraction either side of the nominal value of 0.7. The steady state operating conditions were computed for each feed composition to meet target purities of 0.99 mole fraction methanol in the distillate and 0.001 mole fraction methanol in the residue.

The networks each have only two target specifications, namely the product compositions, but three or four possible manipulated variables. This means that the manipulated variables are not completely defined by the steady state specifications, and possess at least one degree of freedom. This may be used to optimise the steady state. The steady state operating cost was the vapour flow leaving the reboiler of the high pressure column, as used by Lenhoff and Morari. For each feed composition in each network the reboiler duty was computed at a number of possible operating conditions. The minimum operating cost was found by inspection of these results.

The operating states used by Lenhoff and Morari differ significantly from the minimum cost operating states derived as above. Although there is no indication that the former were optimised, this deviation is greater than one would expect from a standard column design procedure. The difference is probably due to the models employed in this work. Lenhoff and Morari use a more sophisticated vapour-liquid equilibrium calculation than the constant relative volatility model employed here. A relatively small change in relative volatility will have a large effect in a high-purity column because the change is compounded at each tray. The relative volatility is known to change considerably throughout the columns, as the entries in tables E2, E5 and E8 show. The consequent difference in separation power of each section has shifted the optimal steady state to the new point. The difference is not important to this work because the same models and equilibrium calculations have been used throughout.

The linear term of the objective function was taken directly from the reboiler vapour load. To find the quadratic weighting factor on the off-specification products, the procedure in chapter 3 was applied at the nominal feed composition steady state of all 3 configurations. The distillate, 99 Mole per cent methanol, can be regarded as the marketable product, while the residue, 99.9 mole per cent water, is waste. The distillate impurity must be closely controlled, so a 5 per cent change in the fraction of water was permitted. The waste product is not so critical but the methanol in the waste should not rise more than 50 per cent over the nominal value. This sets an absolute deviation of 0.05 per cent on both streams. An Euclidean norm of product composition error of 0.0005 was found to correspond to the changes of operating cost given in table 8.2.

Model	Decrease in vapour flow (GJ/min)
LENMO A	66
LENMO B	17
LENMO C	21
Average	34.7

Table 8.2 Effect of product purity on operating cost

From equation 3.21, the required weighting factor is $1.38 * 10^8$. To avoid the numerical difficulties of manipulating very large and very small numbers, the objective function was scaled by 10^{-4} . The same objective function was used on all three configurations (8.1).

$$J = \int_{t_o}^{t_f} (1.38 * 10^4 (x_{error,D}^2 + x_{error,B}^2) + 10^{-4} (Q_{R2} - Q_{R2,ss})) \cdot dt \quad (8.1)$$

where x_{error} is the absolute composition deviation,
 Q_{R2} is the reboiler duty of column 2,
 ss signifies the value at the target steady state, and
 D,B indicates distillate or residue stream.

An alternative objective function with only quadratic terms was used to examine the importance of the weighting factor, equation 8.2. This is very similar to the one used by Lenhoff and Morari, but the penalty on control action is omitted.

$$J_{QUADRATIC} = \int_{t_o}^{t_f} 10^6 (x_{error,D}^2 + x_{error,B}^2) \cdot dt \quad (8.2)$$

8.3 OPTIMAL CONTROL

Due to the nonlinear sensitivity of product compositions to the control variables, closed loop control with linear two-term controllers proved most unsatisfactory. Although the controller could be tuned to give effective control of a switch in one direction, it would need retuning to cope with the reverse switch. It was clear that a practical control scheme needs some compensation for this nonlinearity. A typical solution would be to control the temperature three or four trays above the reboiler instead of controlling the residue temperature directly, with a similar arrangement for the distillate. This technique is demonstrated in Tyreus and Luyben (1975). The development of such schemes was considered inappropriate for this work, especially as closed loop control has no proven value in switchability (see chapter 7).

The optimal control problem was posed with piecewise constant manipulation of the control variables. This was successful because the product purities are not unduly sensitive to the control action. A switch duration of 16 minutes with a setting time of a further 16 minutes was found to be too short, giving a pronounced discontinuity as the settling period was entered. The switch duration was extended to 24 minutes and the setting period to 40 minutes in the final runs.

The size of the problem and relatively heavy computing load in the column section procedures gave long computation times, so the optimisation was carried out in blocks of around 50 iterations. The initial runs were carried out with nonlinear control constraints set up on all product flowrates to prevent either the reboilers or condensers 'running dry'. It soon became clear that the optimal state trajectory was unlikely to violate any of these constraints, so they were removed to improve the performance of the optimal control algorithm (see section 4.4.1). Inspection of the final state trajectories confirms that this simplification was correct in all cases.

The optimal controls (percent) and deviation of the product purities from the targets are shown in figure 8.3 - 8.8. As in the previous case study (chapter 7) the accuracy obtained for the optimal control runs was higher than necessary to confirm the ranking of the three configurations. Examination of the progress of the algorithm indicates that all results are within 50 GJ of the optimum cost.

Some further optimal control runs were performed using the purely quadratic objective function, equation 8.2. The previously obtained control policy was used as a starting point for the optimisation. The resulting control policies are not significantly different from the original policies obtained with the full objective function. However the product specifications show much smaller deviations from target. Again the obtained accuracy was substantially better than needed for separation of the cyclic switching costs.




Figure 8.3 LENMO A Optimal control A to B




Figure 8.4 LENMO A Optimal control B to A




Figure 8.5 LENMO B Optimal control A to B




Figure 8.6 LENMO B Optimal control B to A




Figure 8.7 LENMO C Optimal control A to B




Figure 8.8 LENMO C Optimal control B to A

8.5 RESULTS AND DISCUSSION

The results of the preceding computational runs are shown below, together with the results of Lenhoff and Morari, for comparison.

Configuration	LENMOA	LENMOB	LENMOC
Optimal nominal steady state vapour flow (GJ/min)	1300	1507	1040
Full switch cost (GJ)			
A-B	-928	289	-366
B-A	-637	2766	486
A-B-A	-1565	3055	120
Quadratic switch cost $\times 10^6$			
A-B	0.022	0.230	0.487
B-A	0.234	1.384	0.238
A-B-A	0.256	1.614	0.725
Lenhoff and Morari economic performance index (GJ/min)	1281	1686	1569
Lenhoff and Morari dynamic performance index minimum $\times 10^6$	39.9	13.2	5.4

Table 8.9 Steady state and dynamic costs

The steady state operating cost for each system is comparable to the economic performance index employed by Lenhoff and Morari for configurations A and B, but not C. As stated in section 8.2, the operating conditions given by Lenhoff and Morari were not optimal for the simple models used in this work, and configuration C was capable of the greatest improvement. This difference will not affect the

optimal control results with either full or quadratic objective because the steady state operating cost is either cancelled out or absent from the objective.

The problem of weighting the quadratic terms in the objective function can be resolved in this case study by comparing the results for the two types of objective function. The ranking of the three configurations is the same for both the normal objective (8.1) and the cost with effectively infinite weighting of the quadratic terms (8.2). This suggests that the utility costs are not critical in determining a switching policy to maintain product specifications throughout the switch. It confirms the conclusion of the 'perfect control' runs of chapter 7, which show that there is no operating cost penalty for close control of the product specifications. The method used to determine the weighting factor appears to give a lower bound on range of suitable values.

The cost of switching can be directly related to the normal energy usage of the process. It can be seen from table 8.9 that the switch cost corresponds to less than 40 seconds of normal operation in all but one case, and always less than 120 seconds. This is a very small fraction of the cost of operating the plant at steady state for say 6 hours or more between switches. The cost of switching would thus be irrelevant to the selection of a system from the three presented. The negative switching cost which occurs in some instances is subject to the same considerations as in chapter 7. The product purity deviation is in all cases compatible with the specified nominal deviation of 0.0005 mole fraction.

Although the same ranking is obtained for both types of objective function, this ranking is strikingly different from the dynamic performance ranking of Lenhoff and Morari. The reasons for this will now be considered.

The dynamic performance index is concerned with the performance of each system under simple feedback control, so it is measuring some

property of the transfer function of the linearised system. The piecewise constant optimal control cost measures the ability of the manipulated variables to effect rapid changes to the operating state under deterministic open-loop control. These measures are fundamentally different in two ways.

Firstly the dynamic performance index is based on a 'square system', i.e. the number of regulated variables is equal to the number of manipulated variables. For systems with more manipulated variables, Lenhoff and Morari constrain the surplus variables to the steady state values. The choice of manipulated variables has a marked influence on the dynamic performance index - for each system, the published indices cover a range up to five times the minimum quoted in table 8.2. The optimal control problem is posed with all manipulated variables as controls, so there are fewer constraints on the resulting control policy. This can only decrease the resulting optimal control cost. A suitable multivariable controller may have a similar effect. Lenhoff and Morari report that independent control of the compositions of the four streams leaving the column (presumably with proportional controllers) does improve the dynamic index slightly, but not enough to alter the ranking of the three configurations. They do not consider the use of more sophisticated controllers, or of state feedback control, which would offer a greater improvement of the index, and may be sufficient to change the ranking.

Secondly there are properties of the system which are known to cause difficulties in closed loop control but which would have little effect of switching. In this instance attention is focussed on the presence of zeroes in the right-half-plane of the transfer function of the linearised system. These affect the implementation of a control scheme because the inverse of the plant transfer function will have a right-half-plane pole for every such zero. A controller which maintains very tight control would ideally have the same transfer function as this inverse (Bristol, 1981), but in practice the right-half-plane poles would make such a controller unstable. This limits the effectiveness of the control scheme.

Switching, on the other hand, is concerned with transferring the system to a new operating state as efficiently as possible. Short-term deviations in the controlled product specifications may be inevitable for the reasons given above, but it is not mandatory to maintain specifications throughout the switch. The switching policy for a deterministic system can minimise the cost of any deviations by a feed-forward strategy determined by the optimal control package.

The ranking of dynamic performance indices is allied to the existence of right-half-plane zeroes in each system (Wong, 1984). All of the control schemes examined by Wong for configuration A contain one or more right half plane zeroes, whereas only one scheme of configuration C, and none of B, exhibit such features. (The zero of C was very close to the imaginary axis, and would have a negligible effect in practice).

The origin of right half plane zeroes was mentioned in chapter 5, but will be illustrated here in context for convenience. Consider the system below (Figure 8.10)




Figure 8.10 System with a right half plane zero

A step increase of the input will cause a small decrease in the output almost immediately through the top signal path. The bottom signal path will take longer to respond but eventually it will overcome the decrease caused by the top path and the output will increase. This 'wrong-way response' is a feature of all systems with a right-half plane zero in the transfer function.

The distillation column models used in this study neglect the dynamics of flow changes, so the effects of the manipulated variables on the column product flows are instantaneous. The final product composition in configuration A (and partly in B) is the result of blending two streams of unequal composition, so it will be a function of the relative flowrates. The fast signal path of figure 8.10 is thus the blend ratio of the product streams, and it is instantaneous with respect to the controls. The change of vapour and liquid flows in the columns will also have an effect on the separation efficiency of each column section. The composition of each stream will therefore alter but will be subject to a relatively long exponential delay. This is the slow path in figure 8.10. The presence of one or more right half plane zeroes is readily explained in configuration A. Why are there none in configuration B ? The gain of the slow path is evidently less than the gain of the fast path in this instance, so the step response will have a pronounced 'kink' but will not change direction.

The excellent switching properties of configuration A may be explained in similar terms. The instantaneous effect of the controls on the blending ratio of each product stream enables the compositions to be maintained very close to the target specifications throughout the switch. There are sufficient (4) degrees of freedom to allow this without interfering with the slow shift of the compositions within each column to the target states. Note the similarity of this type of control to the bypass control of the heat exchanger networks in chapter 7, which are also capable of efficient switches.

The sign of the deviation of the product from the target specification will not affect the cost of a switch if a purely quadratic objective function similar to (8.2) is used. This means that for a linear system, reversing the direction of a switch will not affect the switch cost, so a switch from A to B should cost the same as a switch from B to A. Inspection of the results of table 8.9 reveals that the switch direction has a very strong influence on the optimal switch cost, so the nonlinear features of the systems must be influencing the switch

trajectory. The association of infeasible switches with nonlinear behaviour was mentioned in chapter 2. This example illustrates how such behaviour can also affect switchability in systems even where the switch is known to be feasible.

The preceding discussion of switchability and control properties was based on the properties of a linearised system. This suggests that the switchability of configuration B should be better than configuration C, because it has an instantaneous control of the blend ratio of the residue streams. The results show clearly that the asymmetry of the switches in configuration B outweighs the advantages conferred by the control scheme.

The case study has shown that all of the conclusions for the heat exchange networks of chapter 7 are also applicable to the heat integrated distillation schemes of this chapter.

1. Switches implemented using the optimal policy will have very little effect on the overall operating cost of the process. The cost will typically be equivalent to less than one minute of normal operation per switch.
2. The product specifications are constrained within acceptable limits without difficulty throughout the switch for all configurations by the optimal switching policy.
3. Indices of controllability based on the transfer function of the linearised model are unsuitable for estimating the switchability of a process. This implies that closed-loop control will also be unsuitable, since the published controllability measures are intended for such control schemes.
4. Non-linear behaviour of the system can have a major effect on switchability even if all switches are feasible.
5. Control schemes which allow rapid control of product specification will promote switchability, providing the control range is sufficient.

CHAPTER NINE

CRUDE OIL DISTILLATION SYSTEM CASE STUDY

9.1 SYSTEM DESCRIPTION

The studies presented in the previous two chapters concerned systems where the dominant interactions were temperature and composition induced respectively. It was found that the differences in switchability were due in part to the manner in which the manipulated variables could affect the key properties of the products. The choice of controls would therefore influence the switchability to a similar degree as the actual heat integration scheme. This final case study aims to measure the effect of energy integration independently of the control scheme by using a consistent set of manipulated variables throughout. Both thermal and composition effects are represented in the system chosen for investigation.

The system consists of a multicomponent complex distillation column and the associated heat recovery network for the separation of crude oil. It is based on actual plant data kindly provided by BP International PLC. The principal features have been retained but the numerical data has been altered as necessary to preserve confidentiality. The column structure is shown in figure 9.1. The feedstock is split into four fractions, the exact ratio being determined by the feed composition. Liquid is drawn off at intermediate stages and cooled, then returned to the column to provide inter-section cooling. These 'pump-around' streams, as they are called, regulate the reflux ratio in each section. The main vapour flow comes from partial vapourisation of the feed, but direct steam injection is used to strip valuable light products out of the residue and side streams. The modelling of the main column is discussed in the next section.

Crude oil separation consumes a substantial amount of energy, and the

recovery of waste energy is a very important aspect of the operation. The feed is heated to a temperature equivalent to around 700K in the unvapourised liquid before injection into the main column. Most of this heat is in practice recovered from the product streams, and the oil fired heater only accounts for about one third of the final heat content of the feed. There are many different heat exchanger networks which will accomplish this task. The purpose of this case study is to investigate how the choice of heat exchange network affects the switchability of the combined column and heat exchange system.

Three networks have been selected for study (figure 9.2). The first is a simplified version of the original network currently in use at the refinery. The temperatures and flowrates within this network were used to define the heat recovery problem. The second network is a solution to this problem reported by Hanson (1984), subject to a minimum approach temperature of 50K. The third network is a 'minimum equipment' solution developed to represent a scheme with virtually no interaction between column and network. Even so, this network recovers over two thirds of the heat which either of the complex networks recover. The modelling of the heat exchanger networks is discussed in section 9.3.




Figure 9.1 Main column schematic

Key: DIST - Distillate
 LGO - Light gas oil
 HGO - Heavy gas oil
 RES - Residue

TPA - Top pump around
 MPA - Mid pump around
 BPA - Bottom pump around






Figure 9.2 Heat recovery networks

(These are described in detail in appendix F)

9.2 THE MAIN COLUMN MODEL

The basis for the main column model is the node and section model presented in chapter 6. There are several aspects of the system which require attention before the main column is implemented in this form.

Crude oil is a typical example of a complex mixture. The number and identity of the individual components is unknown and therefore individual pure component properties cannot be used to find the physical properties of the mixture. The principal solution to this problem is to consider it as a mixture of a relatively small number of cuts; each cut is characterised by a pseudo-component with average properties of the cut. A variant on this is to select a mixture of pure compounds whose physical properties approximate those of the crude oil. The characterisation problem is covered in detail by Gani (1980).

It is essential to reduce the number of pure components in our representative mixture to a minimum, since each component introduces another 13 equations. Against this, the number of components must be at least the number of product streams, because a perfect separation will result in pure products. Four components were used in the model, and in the absence of further information they were selected from the straight-chain homologous paraffin series. Figure 9.3 shows the boiling point relationship for C_8 to C_{20} of this series at 2 bars pressure.

It can be seen that the four product streams would, if pure, correspond to C_{10} , C_{13} , C_{17} and C_{18} respectively. To complete the characterisation, the relative quantities of each component need to be determined. Unfortunately the TBP (true boiling point) data for the two feed types was not available, so an alternative method was to base the fractions on a perfect separation. This means that the fraction of each component in the feed was the same as the fraction of the feed leaving from the associated product stream. Steady state simulations were used to check the characterisation, and these showed that the




Figure 9.3 Boiling points of n-paraffins at 2 bars

mixture C_9 , C_{14} , C_{17} , C_{19} was a closer fit to the observed operating conditions.

The main column differs in several respects from the simple columns of chapter 8. It features multiple products, partial condensation of distillate, side columns and interstage cooling. These are easily modelled by extending the node and section concept of chapter 6. The side columns and pump-around streams are connected by defining a node at every point where a stream enters or leaves the main column. However, to produce an exact model would require an excessive number of nodes, since the entry and exit tray positions of these streams are rarely coincident, and each node introduces another seven equations in the model. The column was simplified slightly so that each node connects a side column and pump-around stream simultaneously as required.

The physical property calculations in the main column needed some adjustment to give temperatures sufficiently close to the observed values for correct operation of the BP heat exchange network. The major change from the model of chapter 6 was to use mass flows and

holdups for all calculations instead of molar flows. The main reason for this change was the wide range of molecular weights (128-268) in the column. There is no reason why molar units must be used; the component and enthalpy balances are still valid, and the equilibrium data are readily converted into mass fraction units.

For hydrocarbon species, mass units will often be more appropriate. Many physical properties (e.g. specific liquid volumes, vapour enthalpy) show little variation over a considerable molecular weight range when expressed in mass units (API, 1981). If molar units were used, these values would need to be divided by the average molecular weight. Additionally, the heat recovery network requires heat capacity flowrate information which would also need conversion. To ensure that the overall mass balance in the column was preserved, the molecular weight would need to be computed from the compositions of each stream. It is felt that such complications are unwarranted for this work.

The equilibrium data were initially estimated from the vapour pressure ratios of the representative mixture components. Again, it proved impossible to meet the target temperatures without adjusting these by trial and error. Two relative volatility vectors were prepared. The first defines the separation efficiency of the column sections, taking into account the activity of the various species. The second determines the bubble point of the mixture at the nodes, which defines the output temperatures.

Numerical details of the column model may be found in appendix F.

9.3 THE HEAT EXCHANGER NETWORKS

The original system as presented by BP contained the main column as described together with a group of smaller columns to separate the light products in the distillate. These columns were integrated with the heat exchange network through their reboilers. Initially it was hoped to include these columns in the system for study, but the computational demands were found to be excessive. Instead, the reboiler effects have been represented by constant heat loads on the network. This simplification was also used by Sullivan (1978) to reduce the complexity of a similar system. The reboilers have been included as cold streams in the heat exchanger network synthesis problem solved by Hanson (network 2).

Although the original system featured utility exchangers at the end of some product streams to meet certain product target temperatures, these have been omitted from the model used here. This is acceptable because the cost of operating these utility exchangers is negligible (air cooled) and the product temperatures from the network do not appear in the objective function. This simplification has been extended so that any exchanger which is completely downstream of the column, so that it will never affect the column operation, has been omitted. In networks 2 and 3, the constant heating loads are satisfied by streams which have no influence on the column, so they do not all appear in the model.

Modelling of the first heat exchange network was straight-forward because much of the network data was provided by BP. This enabled heat transfer areas and holdups to be calculated. The only significant changes required were to the cold utilities. The models used in chapter 7 are appropriate when cooling water is used, but the thermal holdup in an air cooler is very much less, since it has liquid on only one side. The model therefore expresses the heat removal as a fraction of the maximum duty of the cooler, and all holdup is included in the preceding or following exchanger. A similar line of reasoning

was applied to the hot utility, which is a furnace for partial vapourisation of the feed; the heat is transferred by gas convection and radiation to the feed tubes so the overall holdup is small.

The thermal capacity of most of the exchangers in the first network was evaluated from the physical construction data provided. These heat exchanger designs took account of the effect of the physical properties of the fluids on the overall heat transfer coefficient. The second and third networks were designed with simplified countercurrent exchangers similar to the model proposed in chapter 5. The thermal capacity of these exchangers could have been found by performing a detailed design on each, but this was unnecessary because the thermal capacity of exchangers operating under similar conditions was known from the first network. The thermal capacity of all exchangers with similar process streams was taken to be proportional to the product of heat transfer coefficient and area, which is readily computed from the operating conditions.

The latent heat of vaporisation is important because the temperature change in the pump-around streams determines the rate of condensation of vapour in the top three nodes of the column, and consequently the distillate (vapour) flow. The latent heat of vaporisation of the column fluids is not constant in practice but the model as described in chapter 6 requires an average value for all stages. The column operating conditions are known from the data supplied for network 1. the value of 200 GJ/ton for the the latent heat was found to give the correct column internal flows under these conditions.

Numerical details of the heat exchange networks are given in appendix F.

9.4 SWITCHES AND OBJECTIVE FUNCTION

The crude column can handle a wide range of feed compositions. The operating conditions of the column and associated network will be adjusted to suit each feed type to give the most profitable separation. As stated in chapter 7, a widely separated pair of steady states are desirable for switchability studies. A suitable switch is the changeover from a heavy crude to a light crude, or vice versa. Heavy and light refer to the average molecular weight of the crude oil. This type of transition was used by Sullivan (1978) for his changeover studies for similar reasons.

The operating data for the system was provided by BP for suitable extreme conditions. This is assumed to correspond to the most profitable separation for each feed, using the representative product values which were also given, see table F.4, appendix F. These values only apply to product of a certain purity, so the quadratic off-specification penalty must be used to constrain the compositions to the nominal values. This raises the question of purity measurement of a multicomponent mixture.

One may measure the purity of a mixture with a small number of components by the fraction of unwanted components in each stream. However the crude feed is a mixture of a very large number of individual species, so it will generally be impractical to define wanted and unwanted components of a particular stream. The practical alternative is to measure some characteristic in the column which is representative of the desired composition. The bubble point is frequently used for this purpose (Sullivan (1978), Rademaker et al (1975), Tyreus and Luyben (1975) etc), and it is justified because the separation is to split the feed according to molecular weight of the component species. In a homologous series such as of hydrocarbons, the molecular weight is closely associated with the bubble point at a given pressure. It is convenient for this work because the bubble point is already computed in the model.

The objective function for the assessment of switchability as presented in section 3.9 consists of a linear and a quadratic part. The linear part, which represents the on-specification instantaneous operating cost, has until now been equated to the utility consumption. This essentially neglects the increase in value of the products due to processing. In the previous examples this has been of little consequence, because the product flows have been constrained within relatively close limits, so the value can be adequately represented through the quadratic terms. In this example, the value of the product streams is proportional to their flowrate, which in turn is directly controllable by the switching policy. The distribution of the feed to the product streams is thus equally important as the stream composition characteristic.

The steady state operating cost for each of the feed types is shown in tables F.9, F.12 and F.15 of appendix F. As expected it is dependent on the amount of energy recovered by the network. The negative sign indicates that according to the method of calculation employed, the refining operation is profitable.

The weighting factor for the quadratic term was calculated as described in section 3.9. The presence of feed and product value terms in the linear part of the objective function does not affect the validity of equation 3.20. The sensitivity of the heavy feed steady state operating cost and the four product bubble points to each control was calculated by a finite difference analysis. The change in operating cost corresponding to an Euclidean norm shift in product bubble point of $0.5K$ was computed for each control. The mean value of this over all controls was substituted into equation 3.20 to give a weighting factor of 0.0174.

The resulting objective function used for this study was;

$$J = \int_{t_0}^{t_F} (0.0174 \sum_i (T_i - T_{i,ss})^2 + \sum_j c_j (u_j - u_{j,ss}) - \sum_i c_i (F_i - F_{i,ss})) .dt \quad (9.1)$$

where c is the stream cost given in table F.4,

T is the stream temperature

F is the stream flowrate

i is the product stream identification,

j is the utility stream identification,

ss indicates the final steady state value.

(The feed cost does not appear in 9.1 because the feed flow is always equal to the final steady state value, and thus cancels out)

9.5 OPTIMAL CONTROL

Before the optimal control policy can be found the set of manipulated variables must be chosen. The column model has eight suitable variables, the heat exchange networks each have up to four. Since there are only four product specifications to be met, there are up to eight degrees of freedom in the system at steady state. The choice is between a small number of controls over many time intervals or a large number over few intervals. It is inadvisable to have too many decision variables in the optimisation problem because at least the same number of gradient evaluations are generally required to define the Hessian matrix through the updating method employed by VF02AD.

The available controls for the main column are;

Flowrates: *Top pump around
 *Middle pump around
 Bottom pump around
 *LGO product
 *HGO product
Stream: *LGO stripping
 *HGO stripping
 Residue stripping

The heat exchanger controls are,

Network 1 : utility cooler on TPA (no. 2 fig. F6)
 utility cooler on MPA (no. 8 fig. F6)
 stream split on feed
 *feed preheater furnace duty (no.16 fig. F6)
Network 2 : *bypass on TPA (no. 3 fig. F9)
 bypass on Distillate (no. 5 fig. F9)
 *feed preheater furnace duty (no.13 fig. F9)
 stream split on feed
Network 3 : *feed preheater furnace duty (no. 5 fig. F12)

The controls marked with an asterisk are used in the optimal control problem. All members of this set were observed to have a strong influence on the compositions within the column or the product flowrates. The set is also consistent for the three heat exchange networks, except that TPA flow is replaced by TPA bypass on network 2.

A simulation of each switch was performed using constant controls to check the modelling and final control settings. If the system is globally asymptotically stable it will converge to the desired steady state under such a control policy. We can be reasonably confident that the systems are stable by applying theorem F of chapter 2. The flowrates from the column are defined either explicitly or indirectly through the vaporisation fraction of the feed, and the composition changes do not have a dominating effect on the temperature of the two main product streams. Convergence to the target state was achieved in all cases, with a negligible error (measured as temperature offset) after 150 minutes operation. The switch costs using this simple policy are shown in the next section.

The control bounds were determined by inspection of the flowsheet and applying engineering judgement. Clearly an air cooler can be shut down completely by closing the draught flaps, but a preheat furnace would be limited by the risk of damage from temperature changes if turned down too much. Column conditions (flooding, weeping and entrainment, see e.g. Rademaker et al, 1970) would constrain the vapour and liquid flows to and from the column. Consideration of these and other factors led to the control limits given in appendix F. The exact limits are unimportant providing they are consistent for all cases.

The optimal control policy was sought for all switching problems using four control intervals between 0 and 32 minutes, with a settling time of 118 minutes. These times were chosen by inspection of the previously obtained simulation results. This number of control intervals

presents a total of 31 decision variables to the minimisation package. Previous experience shows that this is close to the practical limit for problems of this size. Fifty iterations of the SPEEVOP algorithm took around 20000 equivalent CDC 6600 seconds of computer time. This large computing load prevented further runs being made, even though they were required in most instances to place tighter bounds on the cost of the optimal switching policy. The accuracy obtained was sufficient to place the three networks into a well-defined ranking of cyclic switching cost (see section 7.6).

The control policies and temperature deviations for each switch are shown in figures 9.4-9.7. It will be noted that the switching policies apply limiting control action in many cases to force a rapid change of states. Even so, there is an unavoidable temperature deviation on certain streams which exceeds the permissible deviation of 0.5K. The heat exchange network has a significant effect on this deviation, but it appears to be essentially a column-induced phenomenon, since it is present even in the minimal network 3. The validity of the optimal control policy for the current work is thus unaffected.




Figure 9.4 Optimal control policies for switch A-B




Figure 9.5 Optimal target deviations for switch A-B




Figure 9.6 Optimal control policies for switch B-A




Figure 9.7 Optimal target deviations for switch B-A

9.6 RESULTS AND DISCUSSION

The switch costs and steady state operating costs are given in table 9.10 below.

Network	1	2	3
Steady state switch cost			
A - B tons	10.692	9.162	16.191
B - A "	21.845	47.001	18.332
A - B - A "	32.537	56.631	35.293
Optimal switch cost			
A - B tons	-7.405	-7.912	-6.953
B - A "	-1.687	-0.694	-1.089
A - B - A "	-9.092	-8.606	-8.042
VF02AD accuracy			
A - B	-0.052	-0.006	-0.007
B - A	-0.018	-0.021	-0.062
Steady state cost			
A tons/min	-0.168	-0.170	-0.137
B "	-0.089	-0.087	-0.055

Table 9.10 Relative costs of systems 1-3

key:- state A - processing heavy crude
state B - processing light crude

The steady state operating costs are a useful guide to the level of energy integration. The product values and flows are independent of the heat exchange network, so the difference must be due to the network energy consumption. The costs are negative because the plant is operating at a profit - the increase in value of the feed exceeds

the energy consumed. Networks 1 and 2 are equally good in this respect, and the minimum network is a poor third, as anticipated.

The cost of the switch using a very rudimentary switching policy is given to illustrate the effect of the heat exchange network on the system dynamics. The 'steady state switch' cost is the cost of simply moving all controls to their target steady state values and allowing the system to converge to the target operating state. The time constants within the system, related to the poles of the open-loop transfer function, determine how rapidly this occurs. The effect on the objective function is more complex because it involves quadratic terms but the difference in steady state switch costs indicates the magnitude of the effect.

The optimal cyclic switching costs in table 9.10 are all negative. This means that on the basis of the objective function (9.1) a period of operation containing a switch will be considerably more profitable than steady state operation. The difference in costs is not negligible, unlike the previous case studies. The quadratic terms in (9.1) cannot be negative, so the increased profitability must be due to the linear terms, either the energy costs or the product values. The maximum possible saving through energy costs is around a third of the observed increase in profit, so the main contribution must arise from the redistribution of material in the product streams. The level of redistribution is relatively small because of the large differences in stream values. An average increase of only 0.4 tons/minute of the more valuable distillate at the expense of the residue during the switch period would account for all the observed profit.

The first explanation is that there is some feature of transient operation of the column which permits greater efficiency. The existence of periodic control policies which are superior to steady state operation of distillation columns has been reported, see for example Pollard (1967). However, Sullivan (1978) demonstrated that the steady state of a multicomponent crude column was superior to periodic operation, and no evidence of oscillations (other than induced by

discretization of the controls) was found in the optimal control runs. Transient increases in efficiency are therefore unlikely to be the cause.

An alternative explanation is that the negative switch costs are an artefact of the objective function employed in this study. The instantaneous operating cost has been shown to be very sensitive to the product distribution, which in turn affects the purity of the product streams. The weighting on the quadratic terms of the objective function has been chosen to constrain the temperatures in the column, which characterise the stream compositions, within acceptable limits. The resulting objective function is thus a relatively crude approximation to the full economic objective function which would have been used to define the target steady state. The simplifications give rise to an objective function whose minimum is not coincident with the target state.

The application of final steady state controls following the switching period ensures that the target state is reached, so the switching policy satisfies the original requirements. However, the optimal control package will minimise the cost of the switch by adopting certain features of the optimal steady state during the switching period. This behaviour is noted and discussed in section 7.5. The cost of the same switch as measured by the full economic objective will be substantially higher.

It appears from the temperature responses of the residue and distillate streams that there is indeed a tendency to establish an intermediate steady state as described above. This being so, the absolute cost of the switch cannot be reliably related to the normal operating cost. Nevertheless, the differences in switch cost are still meaningful, since the redistribution effect is the same in every case.

The main purpose of this case study is to assess the effect of energy integration on switchability. The influences of the control structure

and the differences in the major nonlinearities of the systems, which dominated the earlier studies in chapters 7 and 8, have been eliminated from this case. The differences in optimal switch cost are only attributable to the energy integration scheme itself.

Before the optimal switch costs can be compared the accuracy of each result must be considered. The projected descent step at the end of the optimal control run is included for this purpose, as described in section 4.6. The accuracy is considerably worse than obtained for previous studies but is sufficient to rank the three networks with confidence. Because the energy stored in the networks is directly related to the network construction, the cyclic switching cost must be used. As explained in section 7.6, the cyclic cost cancels out the effect of the difference in potential energy of the initial and target states.

Network 1, the original BP network, is the most switchable, with the Hanson network (2) second and the minimal network (3) trailing in last place. This clearly shows that heat integration is beneficial to switchability. Network 1 was designed to offer good performance over a wide range of operating conditions. Although network 2 gives better energy recovery at the design point (state A, heavy crude) the performance deteriorates rapidly as the operating conditions move away from it. Network 3 gives consistently poor energy recovery.

It is not obvious why network 2 should be so much worse than network 1 in view of the relatively small difference in steady state operating costs. One possibility is that the flexibility, as defined in chapter 1, is inadequate for the switching process although it will allow operation at the two steady states. This is indicated by the wide variation in the control values of major items in the network (units 3 and 5, table F.12) between the two states.

One must not attach too much importance to the effect of energy integration on the switchability of this process. The difference in cost as a result of choosing a better heat integration scheme will be in the

order of eight minutes equivalent steady state operation. This is too small to be considered during the design of a process which may only be switched at intervals of several days.

The results of this study may be summarised as follows.

1. The choice of energy integration scheme can have a profound effect on the system dynamic behaviour, but the switchability of the process considered was virtually unaffected. Switchability is not an issue in the design of an energy integration scheme in this instance.
2. The switchability improves with energy integration. This contradicts the theoretical results of chapter 2, which show that switches may become infeasible through energy integration. It is clear that there are no general rules for this matter.
3. Network flexibility, over and above the steady state requirements, may be advantageous for switchability of complex processes.

CHAPTER TEN

CONCLUSIONS

10.1 CONCLUSIONS

The feasibility of switches in chemical processes is something that the chemical engineer cannot take for granted. The existence of unstable steady states and limit cycles are potential sources of infeasible switches and can arise in many simple process networks through heat integration. Fortunately there are certain classes of system which are always switchable, so feasibility is not an issue. Theorems to establish feasibility have been presented for globally asymptotically stable systems and for certain types of heat integration networks.

Given that a group of alternative designs are switchable, a method of measuring the switchability of these designs has been proposed, to assist in selecting a suitable design. This consists of finding the optimal control policy to achieve an extreme switch and measuring the associated cost in energy and materials. An optimal control package has been interfaced to an equation oriented flowsheeting package to investigate a number of energy integrated networks using this method.

The modelling of the networks has been considered in detail, and a review of existing models was conducted. This revealed undesirable characteristics in all previously published models which rendered them unsuitable for this work. New models were developed and tested for heat exchangers and distillation columns. The heat exchanger model is unconditionally robust and suitable for a variety of purposes. The column model can represent a wide range of column structures and operating conditions and has also proved sufficiently robust for general applications.

Three case studies were conducted on networks which incorporate these units. The first case study, on a simple heat exchange network, showed that closed loop control costs were not a reliable index of switchability. The presence of control elements which can directly affect the controlled product specifications was a definite asset for switchability. The second case study investigated a pair of heat integrated distillation columns which had been previously analysed for closed loop control. The closed loop results were significantly different to the optimal control results, reinforcing the conclusion from the first study. More evidence was obtained for the use of direct-acting manipulated variables.

The final case study considered a range of heat integration schemes for a crude oil refinery main column. The system avoided the effects of the choice of control variables previously mentioned so the effect of the energy integration was measured in isolation. The results suggest that energy integration will improve the switchability of this class of system, although the improvement obtainable was insignificant in the system examined. Unfortunately the computational demands of this type of system prevented further work on it.

The results from the case studies showed that for the systems considered, the cost of switching using a piecewise constant optimal control policy was small compared to the normal operating cost of the process. Furthermore, the switches examined could easily be implemented without interrupting production irrespective of the integration scheme employed. It would therefore appear that switchability is not an important factor in the design of such energy integration schemes, provided the switching policy can be carefully chosen.

10.2 RECOMMENDATION FOR FUTURE WORK

The optimal control package has been designed to optimise nonlinear systems with arbitrary objective functions. It can do this for relatively small systems but has proved slow and expensive for systems of more than around fifty equations. If the problem definition was simplified a more efficient solution method may be applicable. The switchability of linear systems with very simple objective functions like minimum time has not been considered in this work. Relatively large problems of this type could be solved and systems which possess interesting properties may be identified for detailed study with a more comprehensive method.

The iterative tray-by-tray calculation in the column model is a major consumer of computer time. A more efficient method based on the short-cut techniques in steady state column design may exist. The distillation columns are the central feature of the non-trivial switching problems identified by this work. A faster model will allow a wider range of such systems to be examined.

Finally, as mentioned in chapter 2, a switch can become infeasible as a result of heat integration. It may not be easy to determine the best way of restoring switching feasibility in such cases. The switchability of systems which do not satisfy the criteria of the theorems on feasibility should prove an interesting field of research.

REFERENCES

API (1981)

Technical data book - Petroleum refining, 3rd edition

Arkun, Y. and Stephanopoulos, G. (1980)

'Studies in the synthesis of control structures for chemical processes part IV - Design of steady state optimising control structures for chemical process units', A.I.Ch.E.J. 26 p975

Athans, M. and Falb, P.L. (1966)

'Optimal Control'

McGraw-Hill, New York.

Baber, M.F., Edwards, L.L.Jnr, Harper, W.T.Jnr, Witte, M.D. and Gerster, J.A. (1961)

'Experimental transient response of a pilot plant distillation column' Chem. Eng. Progress Symposium Series no.36 57 p148

Beautyman, A.C. (1984)

'The design of flexible heat exchanger networks'

Ph.D. Thesis, Imperial college, London.

Benallou, A., Seaborg, D.E. and Mellichamp, D.A. (1982)

'Low order physically lumped dynamic models for distillation column control'. Paper presented at the 95th Annual A.I.Ch.E.J. meeting, Los Angeles, California, Nov 14-18th 1982

Bilous, O. and Amundsen, N.R. (1955)

'Chemical reactor stability and sensitivity'

A.I.Ch.E.J. 1 p513

Bogle, I.D.L. (1983)

'The numerical solution of flowsheeting problems characterised by large sparse jacobian matrices'

Ph.D. Thesis, Imperial college, London.

Bristol (1981)

'The right half plane 'll get you if you don't watch out'

Paper presented at J.A.C.C.

Bryson, A.E. and Ho, Y-C. (1975)

'Applied optimal control'

Wiley, New York.

Cameron, I.T. (1981)

'Numerical solution of differential-algebraic systems'

Ph.D. Thesis, Imperial college, London.

Campbell, D.P. (1958)

'Process Dynamics'

Wiley, New York.

Cohen, G.H. and Coon, G.A. (1953)

Trans. A.S.M.E. 75 p827

Curtis, A.R., Powell, M.J.D. and Reid, J.K. (1974)

'On the estimation of sparse jacobian matrices'

J.Inst.Math.Appl. 13 p117

Dahlquist, S.A. (1980)

'Control of the top and bottom compositions in a pilot distillation column'

I.Chem.E. Symposium Series no.56, Ch. 2.6, p25

Doherty, M. and Perkins, J.D. (1982)

'Distillation dynamics IV - Uniqueness and stability of the steady state'

Chem.Eng.Sci. 37 p381

Duff, I.S. (1979)

'MA28 - A set of Fortran subroutines for sparse unsymmetric linear equations'

Report AERE-R8730, A.E.R.E., Harwell.

Eliceche, A.M. and Sargent, R.W.H. (1981)

'Synthesis and design of distillation systems'

I.Chem.E. Symposium series no.61 p1

Farber, E.A. and Scoria, R.L. (1948)

'Heat transfer to boiling water under pressure'

Trans.Am.Soc.Mech.Eng. 70 p369

Finlay, I.C. and Dalgleish, N. (1965)

'Response of single pass heat exchangers to disturbances in flow rates'

J.Mech.Eng.Sci. 7 p318

Flower, J.R. and Linnhoff, B. (1980)

'A thermodynamic-combinatorial approach to the design of optimum heat exchanger networks'

A.I.Ch.E.J. 26 p1-9

Gani, R. (1980)

'Crude oil characterisation and distillation dynamics'

Ph.D. Thesis, Imperial college, London.

Gautreaux, M.F. and O'Connell, H.E. (1955)

'Effect of length of liquid path on plate efficiency'

Chem. Eng. Progress, 51 p232

Gear, C.W. (1971)

'Simultaneous numerical solution of differential-algebraic equations'

I.E.E.E. Trans.Circ.Theory CT18(1) p89-95

Hanson, K.W. (1984)

'The synthesis of heat exchanger networks subject to variable sets of constraints'

Ph.D. Thesis, Imperial college, London.

Koppel, L.B. (1962)

'Dynamics of a forced flow heat exchanger'

IEC Fund., 1 p131

Lefschetz, S. (1957)

'Differential Equations - Geometric Theory'

Interscience Inc., New York.

Lenhoff, A.B. and Morari, M. (1982)

'Design of resilient processing plants I - Process design under consideration of dynamic aspects'

Chem. Eng. Sci. 37 p245

Levy, R.E., Foss, A.S. and Grens II, E.A. (1969)

'Response modes of a binary distillation column'

I.E.C. Fund., 8 p765

Maarleveld, A. and Rjinsdorp, J.E. (1970)

'Constraint control of distillation columns'

Automatica, 6 p51

Matsubishi (1960)

'Dynamic response of multipass heat exchangers'

Trans ASME, J. Basic Engg. 82, 1 p51

Moczek, Otto and Williams (1965)

'Approximate models for large distillation columns'

Chem. Eng. Progress Symposium series 61 no.55 p136

Morari,M., Marselle,D.F. and Rudd,D.F. (1982)

'Design of resilient processing plants II'

Chem.Eng.Sci. 37 (2) p259

Morari,M. (1982)

'Operability measures for process design'

I.Chem.E. Symposium series 74 pl31

Morison,K.R. (1984)

'Optimal control of processes described by systems of differential-algebraic equations'

Ph.D. Thesis, Imperial college, London.

Mozely,J.M. (1956)

'Predicting dynamics of concentric pipe heat exchangers'

IEC pl035 June 1956

Murray,W. (1978)

'Aspects of modelling which affect algorithmic choice and performance'

Paper presented at conference for numerical optimisation, Russell Hotel, London, 20-22nd March 1978 (NPL).

Nishida,N. and Ichikawa,A. (1975)

'Synthesis of optimal dynamic process systems by a gradient method',

Ind. Eng. Chem. PD&D 14 p237

Nishitani,H. and Westerberg,A.W. (1982)

'Progress report on process start-up'

Paper presented at Process Systems Engineering 82, Kyoto,Japan.

Osborne,W.G. and Maddox,R.N. (1965)

Control of distillation columns using a simplified dynamic model'

Chem. Eng. Progress Technical manual, Systems and Process Control, A.I.Ch.E.J., New York, pl19

Paloschi, J.R., Perkins, J.D. and Sargent, R.W.H. (1983)

'Steady state simulation with SPEEDUP'

Presented at A.I.Ch.E.J. Spring National Meeting, Houston.

Paloschi, J.R. (1982)

'The numerical solution of nonlinear equations representing chemical processes'

Ph.D. Thesis, Imperial college, London.

Perkins, J.D. and Sargent, R.W.H. (1982)

'SPEEDUP - A computer program for steady state and dynamic simulation and design of chemical processes'

A.I.Ch.E.J. Symposium series no. 214 78 p1-11

Petlyuk, F.B., Platonov, V.M. and Slavinski, D.M. (1965)

'Thermodynamically optimal method for separating multicomponent mixtures'

Int.Chem.Eng. 5 p555

Pieser, A.M. and Grover, S.S. (1962)

'Dynamic simulation of a distillation tower'

Chem. Eng. Progress 58 p65

Pollard, G.P. (1967)

'Prediction and optimisation of dynamic performance of a distillation column'

Ph.D. Thesis, Imperial College, London

Pollard, G.P. and Sargent, R.W.H. (1970)

'Offline computation of optimal controls for a plate distillation column', Automatica 6 p59

Powell, M.J.D. (1978)

'A fast algorithm for nonlinearly constrained optimisation calculations'

Lecture notes in mathematics 630 - Numerical analysis p144

Springer-Verlag, New York.

Prokopakis,G.J. and Seader,W.D. (1980)

'Dynamic simulation of a distillation tower'

A.I.Ch.E.J. Annual meeting, Chicago, Illinois (November 1980)

Rademaker,O., Rjinsdorp,J.E. and Maarleveld,A. (1975)

'Dynamics and control of continuous distillation units'

Elsevier Scientific, Amsterdam.

Rosenbrock,H.H. (1963)

'A Lyapunov function with applications to some nonlinear physical systems'

Automatica 1 p 131

Shinskey,F.G. (1967)

'Process control systems' 1st edition

McGraw-Hill, New York.

Shinskey,F.G. (1979)

'Process control systems' 2nd edition

McGraw-Hill, New York.

Sinovec,R.F., Erisman,A.M., Yip,E.L. and Epton,M.A. (1981)

'Analysis of descriptor systems using numerical algorithms'

I.E.E.E. Trans. on Automatic Control AC-26 1 p131 (Feb. 1981)

Sourisseau,K.D. and Doherty,M.F. (1980)

'Dynamic modelling of stiff distillation systems'

Presented at J.A.C.C. San Francisco 1980.

Stermole,F.J. and Larsen,M.A. (1963)

'Dynamic response of heat exchangers to flowrate changes'

I.E.C. Fund. 2 p62

Sullivan,G.R. (1978)

'Development of feed changeover policies for refinery distillation units'

Ph.D. Thesis, Imperial College, London.

TEMA (1978)

'Standards of the tubular exchangers manufacturers association'

6th edition, New York.

Tyreus,B.D. and Luyben,W.L. (1975)

'Controlling heat integrated distillation columns'

Chem. Eng. Progress, September 1975 p59

Van Heerden,C. (1953)

'Autothermic processes'

I.E.C. 45 p1242

Van Winkle (1967)

'Distillation'

McGraw-Hill, New York.

Wahl,E.F. and Harriot,P. (1970)

'Understanding and prediction of the dynamic behaviour of distillation columns'

I.E.C. Proc. Des. Dev. 9 p396

Waller,K.V. (1979)

'Simple models for distillation dynamics'

Presented at A.I.Ch.E.J. 86th National meeting, Houston.

Wong,M.P.F. (1984)

'Assessment of controllability in chemical processes'

Ph.D. Thesis, Imperial college, London.

TABLE OF COMMONLY USED SYMBOLS

(Numbers in brackets refer to chapter numbers, if required)

A	System matrix of state space model (3)
A	Derivative premultiplier for mixed systems (4)
A_i	Relative volatility (6)
A_D, A_T	Areas on distillation tray (6)
B	Control matrix of state space model (3)
B_j	Component balance coefficient (6)
$\underline{c}$	Vector of constraints for minimisation
C_p	Specific heat
$\underline{d}$	Step size of minimisation algorithm (4)
d_T, d_D	Liquid depth on distillation tray (6)
E	Thermal efficiency (5)
E_M, E_p	Tray efficiency (6)
e	Weighting parameter for computing mean value
F	Flow (mass, molar or specific heat)
f	Vector function
G	Controllability matrix (2)
$G(s)$	Transfer function (5)
H	Holdup (2,5,6)
H	Hessian matrix; second derivative of J w.r.t $\underline{y}$ (4)
h	Step length of integrator (4)
I	Unit matrix
J	Objective function
K_i	K-value of pure component
L	Instantaneous cost function (3-4, 7-9)
L	Liquid molar flow (6)
$\underline{m}$	Vector of marginal values (3)
m	Dimension of decision variable vector $\underline{y}$
n	Dimension of X (2)
P	Pressure
$\underline{p}$	Vector of adjoint variables
$\dot{p}$	First time derivative of p

Q	Heat flow (2,5)
Q	Quadratic weight on state or output deviation (3,7-9)
q	Elements of Q (3)
q	Dimension of constraint function vector <u>g</u>
R	Real vector space
R	Quadratic weight on control deviation
r	Dimension of U
S	Solution of matrix Riccati equation
s	Dimension of <u>v</u> (4)
s	Laplace transformation independent variable (5)
T	Temperature
T	Space of feasible temperatures.
<u>T</u>	vector of temperatures
t	Time
U	Feasible space of <u>u</u>
UA	Product of heat transfer coefficient and area
<u>u</u>	Vector of control variables
<u>w</u>	Vector of gradients of J. w.r.t <u>y</u>
X	Feasible space of <u>x</u>
<u>x</u>	Vector of states (2-4)
<u><u>x</u></u>	First time derivative of <u>x</u>
<u>x</u>	Vector of liquid composition (6)
<u>y</u>	Vector of decision variables for minimisation (4)
<u>y</u>	vector of vapour composition (6)
α	Step length reduction factor
ρ	Density of liquid
τ	Time constant
λ	Eigenvalue (4-5)
λ	Latent heat of vapourisation (6)

Subscripts

F	Final state
ss	Steady state
E	Escape space
C	Capture space

APPENDIX A

ATMOSPHERIC FLASH PROBLEM - (Chapter 2)




Figure A.1 Process flow diagram

$$\frac{d}{dt} T3 = (T2(T1, T5, 1, F4) - T3)$$

$$F4 = (1 - \phi(T4))$$

$$\phi(T4) = (T3 - T4) / 50$$

$$T5 = T5(T4, T6, F4, F6)$$

Where $T1$ is temperature of stream 1, etc.

$F1$ is heat capacity flowrate of stream 1, etc.

ϕ is the fraction of the feed vapourised.

Heat exchangers are modelled as described in chapter 5.

T1 (°C)	_____	160.000	_____
T2 (°C)	99.729	128.584	159.999
T4 (°C)	98.707	101.689	109.999
T5 (°C)	93.051	92.970	92.000
T6 (°C)	_____	92.000	_____
F4	0.9796	0.5397	2.26×10^{-7}
F6	_____	1.0	_____
heat transfer coeff 1.	_____	10	_____
heat transfer coeff 2.	_____	5	_____
Eigenvalue	-0.5866	+0.2332	-0.9999

Table A.2 Steady state solutions




Figure A.3 Vapour fraction $\phi(T_4)$

The vapour fraction rises linearly from 0.1 to 0.9 over the temperature range 99.5 - 104.5 °C and approaches the limits (0 and 1) asymptotically below and above this range respectively (see Fig. A3). This is accomplished using the 'continuous constraint' function 7.4.

APPENDIX B

HEAT EXCHANGER MODEL EVALUTATION




Figure B.1 Process Flow diagram

	Hot stream	Cold stream
Initial steady state		
Inlet °C	300	200
Outlet °C	274	296
Heat capacity flowrate*	6.0	1.0
Final steady state		
Inlet °C	300	400
Outlet °C	326	304
Heat capacity flowrate*	6.0	1.0
Effective heat capacity*	1.0	1.0
Duty (percent)	60	100

* Units compatible with Tons, seconds, °C

Table B.2 Steady state operating conditions




Figure B.3 Network of units for model I

Network of units for models II and III is as figure B1. Heat transfer coefficient - area product $UA = 4.0$ nominal. In model I, each unit has $UA = 2.474$ to compensate for reduced efficiency of perfectly mixed exchanger.

eigen -value	eigen vector components (normalised)							
	T11	T21	T12	T22	T13	T23	T14	T24
-2.69	0.115	0.248	0.265	0.407	0.398	0.476	0.425	0.352
-7.3	-.244	-.417	-.256	-.088	0.064	0.460	0.400	0.563
-11.6	-.372	-.472	0.083	0.631	0.213	0.150	-.096	-.395
-13.2	-.325	-.362	0.232	0.716	-.089	-.416	0.010	0.127
-25.0	0.006	0	0.070	-.014	0.406	-.125	0.693	-.578
-26.6	-.017	0.004	-.021	0.029	0.304	-.038	0.767	-.562
-30.9	0.031	-.021	-.083	0.011	-.054	0.150	0.863	-.471
-35.4	-.030	0.034	0.134	-.107	-.388	0.241	0.801	-.344

Table B.4 Eigenvectors of model I

APPENDIX C

DISTILLATION COLUMN MODEL EVALUATION




Figure C.1 CHE25 Unit interconnections

Average molecular weight M_w (/Kmol)	80
Downcomer area A_D (m^2)	0.0455
Downcomer cutaway area A_{DCC} (m^2)	0.0211
Weir height d_w (m)	0.0381
Weir length l_w (m)	0.532
Tray area A_T (m^2)	0.366
perforated area A_p (m^2)	0.0366
Density ρ (Kg/m^3)	800
Murphree efficiency	0.8

Table C.2 Tray parameters




Figure C.3 Condenser unit




Figure C.4 Reboiler unit

Unit	condenser	reboiler
Holdup setpoint (Kmol)	1.0	1.0
Pressure setpoint (Bars)	1.0	-
MGAIN (min^{-1})	100	100
PGAIN (GJ/Bar)	100	-

Table C.5 Control loop parameters

Tray no.	Temperature K	mol fraction	Vapour Kmol/min	Liquid Kmol/min	Holdup Kmol
Condenser	326.8	0.793	-	2.580	1.037
1	329.7	0.658	3.732	2.528	0.244
2	333.0	0.522	3.681	2.485	0.276
3	335.9	0.417	3.637	4.317	0.317
4	339.3	0.308	3.469	4.271	0.324
5	543.2	0.202	3.423	4.226	0.323
Reboiler	347.7	0.101	3.378	0.848	1.008

Table C.6 Column conditions for $X_F = 0.5$

Feed Composition (cyclopentane)	0.3	0.5	0.7
Reboiler duty (MJ/min)	96.83	96.82	96.83
Condenser duty (MJ/min)	98.7	99.92	102.0
Distillate composition	0.525	0.793	0.927
Distillate rate (Kmol/min)	1.098	1.152	1.189
Distillate temperature (K)	332.9	326.8	324.2
Residue composition	0.0258	0.101	0.368
Residue rate (Kmol/min)	0.902	0.848	0.811
Residue temperature (K)	351.5	347.7	337.4
Stripping section vapour (Kmol/min)	3.36	3.38	3.50
Stripping section liquid (Kmol/min)	4.26	4.22	4.32
Average relative volatility	2.574	2.450	2.288
Rectifying relative vapour (Kmol/min)	3.56	3.73	3.85
Rectifying relative liquid (Kmol/min)	2.46	2.58	2.66

Table C.7 Column conditions for all feeds

Misc. operating conditions: Feed rate - 2.0 Km³/min

Feed temp.- 350 K

	$x_F = 0.3$		$x_F = 0.5$		$x_F = 0.7$	
	CHE25	CLTEST	CHE25	CLTEST	CHE25	CLTEST
λ_1	-.73	-.72	-.52	-.56	-.54	-.54
z_{11}	0.644	0.561	0.192	0.212	0.062	0.072
z_{12}	0.413	0.454	0.385	0.391	0.188	0.198
z_{13}	0.644	0.692	0.903	0.896	0.980	0.978
λ_2	-2.62	-2.88	-2.07	-2.07	-2.00	-1.78
z_{21}	0.152	0.119	0.233	0.223	0.258	0.351
z_{22}	-.177	-.197	-.078	-.096	0.128	0.152
z_{23}	-.972	-.973	-.970	-.970	-.960	-.924
λ_3	-8.14	-6.52	-6.40	-5.03	-5.70	-4.42
z_{31}	-.122	-.084	-.114	-.099	-.093	-.151
z_{32}	.390	0.485	0.365	0.486	0.407	0.679
z_{33}	-.913	-.870	-.924	-.868	-.908	-.713

Table C.8 Composition Eigenvectors of the two models

Only the first three composition modes of the CHE25 model are shown. The compositions are averaged over all trays adjacent to the relevant node for compatibility with the CLTEST results.

APPENDIX D

HEAT EXCHANGER NETWORK DETAILS (CHAPTER 7)

All units compatible with Tons, minutes, °F

Stream	Inlet (°F)	Outlet (°F)	Fc_p (min)	Fc_p (max)
1	572	176	0.5406	0.5690
2	392	104	0.8108	0.8535
3	104	356	0.7587	0.7966
4	284	536	1.1380	1.1949

Table D.1 Nominal steady state design problem

Operation	Stream	Inlet (°F)	Outlet (°F)	Fc_p
A	1	572	176	0.5406
A	2	392	104	0.8108
A	3	104	356	0.7966
A	4	284	536	1.1949
B	1	572	176	0.5690
B	2	392	104	0.8535
B	3	104	356	0.7587
B	4	284	536	1.1380
C	1	562	176	0.5406
C	2	382	104	0.8108
C	3	104	346	0.7966
C	4	284	526	1.1949

Table D.2 Initial and Target operating conditions




Figure D.3 Network HEX7

	A	B	C	
Duty (H1)	70.63	98.35	80.44	%
Cold Utility (U1)	64.96	71.95	62.09	%
Cold Utility (U2)	53.30	67.87	52.53	%
Hot Utility (U4)	86.38	71.17	70.10	%

Table D.4 Steady state controls for HEX7

Heat exchanger	H1	H2	H3	H4	
Heat transfer units	0.8	1.38	1.16	3.02	
Holdup	1.0	2.0	3.0	4.0	Tons
Utility exchanger		U1	U2	U4	
Utility flow limit		1.5	1.6	3.0	T/min
Utility temperature		100	60	600	°F
Utility Cost		0.2	0.2	1.0	
Holdup		1.5	1.2	2.0	Ton
Stream	1	2	3	4	
Gain	0.04	0.04	0.08	-0.03	
Reset time	1.5	0.8	1.0	1.0	min
Controlled unit	U1	U2	H1	U4	

Table D.5 Miscellaneous data for HEX7




Figure D.6 Network HEX8

	A	B	C	
Duty H1	85.70	96.07	85.24	%
Duty H3	97.22	87.67	94.81	%
Cold utility U1	69.82	88.67	69.36	%
Hot utility U4	63.19	56.01	52.26	%

Table D.7 Steady state controls for HEX8

Heat exchanger	H1	H2	H3	
Heat transfer units	0.960	1.7116	3.60	
Holdup	1.2	2.5	4.5	Tons
Utility exchanger		U2	U4	
Utility flow limit		3.5	4.0	T/min
Utility temperature		60	600	°F
Utility cost		0.2	1.0	
Holdup		2.0	2.4	Tons
Stream	1	2	3	4
Controller gain	0.03	0.05	-0.03	-0.03
Reset time	0.6	0.8	1.0	0.8
Controlled unit	H1	U2	H3	U4

Table D.8 Miscellaneous data for HEX8

APPENDIX E

DISTILLATION NETWORK DETAILS (Chapter 8)




Figure E.1 Network LENMOA

Column number	1	2
Operating pressure (Bars)	1.0	6.8
Number of trays (total)	20	20
Position of feed tray from top	15	14
Latent heat of vapourisation (GJ/Kmol)	37.0	37.0
Specific heat capacity (GJ/Kmol/K)	0.0777	0.0777
Relative volatility rectifying section	2.325	1.850
Relative volatility stripping section	5.180	4.035
Relative volatility Bubble points	3.800	3.000
Feed temperature (K)	342.45	342.45
Reboiler duty limit (GJ/min)	-	1500
Holdup, top node (Kmol)	90.0	90.0
Holdup, feed node (Kmol)	10.0	10.0
Holdup, bottom node (Kmol)	170.0	170.0

Table E.2 LENMOA physical data

Feed composition		0.6	0.7	0.8
Feed split to column 1 %		56.6	56.4	56.6
C1	Distillate flow	11.40	13.26	15.21
	Residue flow	7.48	5.56	3.66
	Liquid flow (Rect)	19.07	19.33	19.27
	Liquid flow (Strip)	38.26	38.39	38.29
	Reflux ratio %	62.6	59.3	55.9
	Distillate composition	0.993	0.993	0.992
	Feed node composition	0.484	0.562	0.639
	Residue composition	0.001	0.001	0.001
C2	Distillate flow	8.79	10.30	11.72
	Residue flow	5.66	4.21	2.74
	Liquid flow (Rect)	22.13	22.61	22.93
	Liquid flow (Strip)	38.61	39.09	39.30
	Reflux ratio %	71.6	68.7	66.2
	Distillate composition	0.986	0.986	0.987
	Feed node composition	0.554	0.623	0.690
	Residue composition	0.001	0.001	0.001
Reboiler heat duty (GJ/min)		1231	1300	1359
Reboiler duty %		82.13	86.67	90.62
Overall distillate composition		0.99	0.99	0.99
Overall residue composition		0.001	0.001	0.001

Table E.3 LENMOA optimal operating conditions




Figure E.4 Network LENMOB

Column number	1	2
Operating pressure (Bars)	1.0	6.8
Number of trays, total	20	20
Position of Feed tray from top	13	11
Latent heat of vapourisation (GJ/Kmol)	37.0	37.0
Specific heat capacity (GJ/Kmol/K)	0.0777	0.0777
Relative volatility (rectifying section)	2.325	1.850
Relative volatility (stripping section)	4.400	4.035
Relative volatility Bubble points	3.80	3.00
Feed temperature	-	342.45
Reboiler duty limit	-	1700
Holdup,top node (Kmol)	90.0	90.0
Holdup,feed node (Kmol)	10.0	10.0
Holdup,bottom node (Kmol)	170.0	170.0

Table E.5 LENMO B physical data

Feed composition		0.6	0.7	0.8
C1	Distillate flow	20.19	23.57	26.93
	Residue flow	3.16	2.38	1.62
	Liquid flow (Rect)	15.76	15.62	14.67
	Liquid flow (Strip)	36.21	38.38	39.72
	Reflux ratio %	43.83	39.9	35.27
	Distillate composition	0.99	0.99	0.99
	Feed node composition	0.868	0.904	0.933
	Residue composition	0.00235	0.00255	0.00263
C2	Distillate flow	23.35	25.97	28.55
	Residue flow	9.98	7.38	4.78
	Liquid flow (Rect)	9.70	10.03	9.54
	Liquid flow (Strip)	47.40	47.58	46.96
	Reflux ratio %	29.35	27.90	25.05
	Distillate composition	0.856	0.899	0.934
	Feed node composition	0.697	0.775	0.850
	Residue composition	0.00057	0.00050	0.00045
Reboiler heat duty		1409	1507	1574
Reboiler duty		82.89	88.65	92.59
Overall distillate composition		0.99	0.99	0.99
Overall residue composition		0.001	0.001	0.001

Table E.6 LENMOB optimal operating conditions




Figure E.7 Network LENMOC

Column number	1	2
Operating pressure (Bars)	1.0	6.8
Number of trays, total	20	20
Position of Feed tray from top	7,15	11
Latent heat of vapourisation (GJ/Kmol)	37.0	37.0
Specific heat capacity (GJ/Kmol/K)	0.0777	0.0777
Relative volatility (rectifying section)	2.05	1.850
Relative volatility (middle section)	4.4	-
Relative volatility (stripping section)	5.18	4.035
Relative volatility Bubble points	3.80	3.00
Feed temperature (K)	-	342.45
Reboiler duty limit (GJ/min)	-	1150
Holdup, top node (Kmol)	90.0	90.0
Holdup, second node (Kmol)	10.0	10.0
Holdup, third node (Kmol)	10.0	-
Holdup, bottom node (Kmol)	140.0 7	170.0

Table E.8 LENMO C physical data

Feed composition		0.6	0.7	0.8
C1	Distillate flow	20.19	23.57	26.93
	Residue flow	13.14	9.76	6.40
	Liquid flow (Rect)	6.06	3.42	0.734
	Liquid flow (Middle)	14.65	12.91	11.59
	Liquid flow (Strip)	36.24	33.68	31.02
	Reflux ratio %	23.08	12.70	2.65
	Distillate composition	0.99	0.99	0.99
	Upper feed node composition	0.971	0.975	0.979
	Lower feed node composition	0.250	0.346	0.468
	Residue composition	0.001	0.001	0.001
C2	Distillate flow	9.50	10.50	12.00
	Residue flow	23.83	22.83	21.33
	Liquid flow (Rect)	13.75	13.58	12.77
	Liquid flow (Strip)	50.98	50.64	49.69
	Reflux ratio %	59.14	56.40	51.56
	Distillate composition	0.964	0.973	0.977
	Feed node composition	0.692	0.770	0.847
	Residue composition	0.455	0.574	0.700
Reboiler heat duty		1020	1040	1056
Reboiler duty %		88.70	90.43	91.82

Table E.9 LENMOC optimal operating conditions

APPENDIX F

Heat Integrated Distillation Networks (Chapter 9)




Figure F.1 General arrangement of system




Figure F.2 Schematic of column model

Section/Node	Number of trays	Holdup (Tons)
D	7	40.0
2	5	88.0
3	5	85.0
4	4	80.0
B	-	40.0
L	4	8.0
H	4	5.0

Table F.3 Column mechanical data

Specific heat	0.0028 GJ/Ton
Latent heat of vapourisation	0.2 GJ/Ton
Column pressure	1.9 Bars
Antoine coefficients (heavy component)	7.0192, 1917.0, 133.5
Relative volatilities for bubble points	60, 5, 2, 1
Relative volatilities for separation	20, 6, 2, 1
Stream values expressed relative to light crude;	
Distillate	1.20
LGO	1.11
HGO	1.05
Residue	0.84
Light crude	1.0
Heavy crude	0.968
Fuel oil to vapourise 1 ton crude	0.00662
Steam to vapourise 1 ton crude	0.0203

Table F.4 Miscellaneous data

Type of crude	Heavy (A)	Light (B)
Feed flow Tons/min	19.16	18.81
Feed temperature K	719.16	693.79
Feed composition x_1	0.26939	0.335076
Feed composition x_2	0.10240	0.1192331
Feed composition x_3	0.08351	0.075666
Distillate flow Tons/min	5.19	6.25
Distillate temperature K	480.5	480.5
Distillate pump around Tons/min	11.29	8.11
Distillate pump around temperature K	409.75	401.49
Node 2 temperature K	520.23	514.37
Node 2 pump around Tons/min	11.95	9.07
Node 2 pump around temperature K	453.34	452.72
LGO flow Tons/min	2.06	2.95
LGO temperature K	531.5	531.5
LGO stripping steam Tons/min	0.600	1.200
Node 3 temperature K	579.75	573.75
Node 3 pump around Tons/min	6.71	6.68
Node 3 pump around temperature K	515.80	506.10
HGO flow Tons/min	1.57	0.88
HGO temperature K	607.5	607.5
HGO stripping steam Tons/min	0.29717	0.121793
Residue flow Tons/min	10.35	8.73
Residue temperature K	631.0	631.0
Residue stripping steam Tons/min	2.0	2.0

Table F.5 Column steady states

	Minimum	Maximum
LGO stripping steam (Tons/min)	0.1	1.25
HGO stripping steam (Tons/min)	0.1	1.25
Top pump around (Tons/min)	8	12
Mid pump around (Tons/min)	8	12
LGO draw-off (Tons/min)	1	3
HGO draw-off (Tons/min)	0.5	2
Feed preheater duty (percent)	80	100

Table F.6 Column operating limitations




Figure F.7 Heat Recovery Network 1 - BP

Heat exchanger	Heat transfer coefficient - area product	Holdup Tons	Duty GJ/min
1	9.7	33	-
2	-	-	0.5
3	6.5	35	-
4	12.1	8.1	-
5	-	-	1.479
6	5	5.4	-
7	3	2.9	-
8	-	-	0.5 (max)
9	10	11.6	-
10	-	-	0.368
11	10	14	-
12	12	11.6	-
13	-	-	1.008
14	12	11.3	-
15	5	6.7	-
16	-	-	10 (max)

Table F.8 Network 1 mechanical data

Steady state	Heavy (A)	Light (B)
Duty: unit 2 (percent)	90	60
Duty: unit 8 (percent)	90	0
Duty: unit 16 (percent)	95.867	84.862
Operating cost (equivalent Tons light crude/min)	-0.16812	-0.08912

Table F.9 Network 1 steady states




Figure F.10 Heat Recovery Network 2 - Hanson

Heat exchanger	Heat transfer coefficient-area product	Holdup Tons	Duty GJ/min
1	13.9	40.0	-
2	4.25	18.0	-
3	8.66 (max)	17.0	-
4	7.68	15.0	-
5	7.0 (max)	14.0	-
6	15.8	16.0	-
7	8.6	17.0	-
8	11.2	22.0	-
9	-	-	0.368
10	2.7	3.0	-
11	-	-	1.008
12	10.78	21.0	-
13	-	-	10 (max)

Table F.11 Network 2 mechanical data

Steady state	Heavy (A)	Light (B)
Duty: unit 3 (percent)	90.0	53.64
unit 5	0.0	100.0
unit 13	93.573	87.990
Split fraction to unit 7 (percent)	39.0	59.02
Operating cost (equivalent tons light crude/min)	-0.16964	-0.08726

Table F.12 Network 2 steady states




Figure F.13 Heat Recovery Network 3 - Minimal

Heat exchanger	Heat transfer coefficient-area product	Holdup Tons	Duty GJ/min
1	4.0	10.0	-
2	8.0	12.0	-
3	8.0	12.0	-
4	36.0	60.0	-
5	-	-	15 (max)

Table F.14 Network 3 mechanical data

Steady state	Heavy (A)	Light (B)
Cooling flow: unit 1 (Tons/min)	1.812	1.401
unit 2	5.013	3.786
unit 3	8.881	5.422
Duty: unit 5 (percent)	95.592	90.820
Operating cost (equivalent tons light crude/min)	-0.13666	-0.05533

Table F.15 Network 2 steady states

UNECA65
84/09/21
RABA

APPENDIX G - PROBLEM LISTINGS

01

01

UNECA65
84/09/21
RABA

APPENDIX G - FORTRAN PROGRAMS

01

01

UNECA65
84/09/21
RABA

APPENDIX G - FORTRAN PROGRAMS

02

02

UNECA65
84/09/21
RABA

APPENDIX G - FORTRAN PROGRAMS

03

03

UNECA65
84/09/21
RABA

APPENDIX G - FORTRAN PROGRAMS

04

04

UNECA65
84/09/21
RABA

APPENDIX G - FORTRAN PROGRAMS

05

05

UNECA65
84/09/21
RABA

APPENDIX G - FORTRAN PROGRAMS

06

06

UNECA65
84/09/21
RABA

APPENDIX G - FORTRAN PROGRAMS

07

07