

INTEGRATED DESIGN OF DISTRIBUTED CHEMICAL MANUFACTURING FACILITIES

by

David Anthony Rowe, MEng., A.C.G.I.

March 1997

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

Department of Chemical Engineering and Chemical Technology,
Imperial College of Science, Technology and Medicine,
Prince Consort Road,
London SW7 2BY,
United Kingdom.



Abstract

Current industrial practice for chemicals manufacturing is reviewed and the motivation for building large, integrated production sites discussed. The concept of distributed manufacturing (chemical production *in situ*) as an alternative philosophy is then introduced. Potential economic, safety, logistic and environmental benefits are considered together with circumstances when the application of such a manufacturing system may be favourable. The problems of realising a distributed plant are briefly discussed in order to highlight those design aspects of interest in a process systems methodology.

The integrated design problem is discussed in the context of a distributed plant and literature in this area is reviewed. It is concluded that any integrated design procedure for distributed plants should explicitly consider the dynamic operation of the process at the process design stage. Potential tools for the combined optimisation of the plant design and operation are reviewed and dynamic optimisation is deemed the most appropriate alternative.

A procedure for the integrated design of distributed plants is presented. The rationale for the procedure is to develop plant designs which operate feasibly in the face of frequent and potentially large changes in demand. The procedure consists of two stages. In the first stage, a candidate process structure is developed using a combination of heuristics and economic bounding properties. A dynamic optimisation formulation is then applied to determine the economically optimal cyclic operating procedure for the candidate structure. The design procedure facilitates the development of an effective design and the corresponding optimal operating strategy. The proposed design framework is demonstrated by application to three representative case studies: a continuous multiproduct plant, an integrated chemical wastewater treatment system and a distributed reactor network.

The results of the work presented are reviewed and some needs for further work identified.

*To my family
& Gaynor.*

Only Bunyan, Sir Walter Raleigh, the author of Don Quixote, and a few other people have had the best of opportunities for working, in this world. Solitary imprisonment, by compulsion, is the one perfect condition for perfect performance.... then his work becomes his pleasure, his recreation, his absorption, his uplifting and all-satisfying enthusiasm. He is miserable only when the work-day closes. And yet a man so circumstanced need never be actually miserable; for he can weave his fancies and continue his work in his head until sleep overtakes him.

Letter from Mark Twain to Mrs. Fairbanks,
6th August 1877.

Acknowledgements

I would like to express my sincere thanks to my supervisors, John Perkins and Steve Walsh. Their guidance, support and encouragement throughout the course of my research has been invaluable.

I would also like to thank my office colleagues, past and present, for their friendship over the last three years. Special thanks must go to Jae and Eric, my desk neighbours who for so long, have tolerated my leaking headphones without complaint.

I couldn't have asked for two better flatmates than Jon and Ed - they have been my friends, counsellors and technical advisors and I will always be grateful for their companionship whilst travelling this long road.

There are not enough words to thank Gaynor, my parents Norah and Tony and my sister Catherine, for their absolute faith and unceasing love which made it possible to keep going.

Contents

Abstract	2
Acknowledgements	5
List of Figures	10
List of Tables	12
Nomenclature	13
1 Introduction	18
1.1 Rationale	18
1.2 Objectives of the project	19
1.3 Overview of thesis	20
2 Distributed manufacturing and the integrated design problem	22
2.1 Introduction	22
2.2 Distributed chemical manufacturing	22
2.2.1 Current chemical processing practice	22
2.2.2 The distributed chemical plant	23
2.2.3 Centralised versus distributed chemical manufacture	24
2.2.4 Areas of research interest	26
2.2.5 Process design issues	27
2.3 Review of integrated design literature	28
2.3.1 Historical perspective on the process design problem	28
2.3.2 Approaches to the integrated design problem	33
2.3.3 Integrated process design based on controllability	34
2.3.4 Integrated process design based on dynamic operation	42
2.3.5 Conclusions	49
3 Design procedures	50
3.1 Introduction	50
3.2 Design procedure	51
3.2.1 Mixed-integer non-linear programming approach	51
3.2.2 Decomposed design procedure	53
3.3 Description of the process operation	54
3.3.1 Operating requirements of a distributed plant	55
3.3.2 Cyclic operation of processes	56

3.3.3	Re-statement of the design procedure	58
3.4	Tools for integrated process design and operation	58
3.4.1	Operating Procedure Synthesis	59
3.5	Dynamic optimisation as a tool for process design and operation	59
3.5.1	Form of the dynamic optimisation problem	60
3.5.2	Solution Methods for Dynamic Optimisation Problems	61
3.5.3	A feasible path dynamic optimiser	65
3.6	Conclusions and introduction to case studies	68
3.6.1	Multiproduct plant design	69
3.6.2	Wastewater plant design	69
3.6.3	Reactor network design	69
4	Dynamic design of a continuous multiproduct process	70
4.1	Introduction	70
4.2	Multiproduct and multipurpose processes	71
4.2.1	Motivation for multiproduct and multipurpose processes	71
4.2.2	Historical perspectives	73
4.2.3	Conclusions	76
4.3	Integrated design of continuous multiproduct processes	77
4.3.1	A multiproduct process design procedure	78
4.3.2	Optimisation based multiproduct process design	84
4.4	Multiproduct reactor case study	84
4.4.1	Problem definition	84
4.4.2	Process model	85
4.5	Dynamic design of a multiproduct process.	91
4.5.1	Framework for the multiproduct process optimisation	91
4.5.2	Objective function	93
4.5.3	Optimisation parameters	95
4.5.4	Time discretisation and control parameterisation	96
4.5.5	Constraints	97
4.5.6	Implementation	101
4.5.7	Optimisation results	101
4.6	Discussion	106
4.6.1	Comparison of methods	106
4.6.2	Consideration of secondary objectives	107
5	Integrated design of wastewater treatment systems	109
5.1	Introduction	109
5.1.1	Rationale	109
5.1.2	Chapter outline	110
5.2	Problem definition	111
5.3	Modelling the problem	114
5.3.1	pH modelling of strong-acid/strong-base systems	114
5.3.2	CSTR model	115
5.3.3	Controller model	117
5.3.4	Diversion system model	118
5.4	Design procedure	120
5.4.1	Developing a candidate structure	122

5.4.2	Formulating the optimal cyclic operation problem	123
5.4.3	Objective function	124
5.4.4	Plant parameters for optimisation	125
5.4.5	Time discretisation and control parameterisation	126
5.4.6	Constraints	127
5.4.7	Implementation	130
5.5	Optimal cyclic operation of base case treatment system	131
5.5.1	Implementation of the base case	132
5.5.2	Solution and results	132
5.6	Optimal cyclic operation of wastewater systems with diversion	135
5.6.1	Economic bounding	135
5.6.2	System One: Two CSTRs and diversion	136
5.6.3	System Two: One CSTR and diversion	140
5.6.4	Discussion of optimisation results	143
5.7	Discussion	145
6	Design of a reactor network for distributed manufacturing	147
6.1	Introduction	147
6.2	General Problem Definition	148
6.2.1	Customer requirements	148
6.2.2	Other operational constraints	148
6.3	Process building blocks	149
6.4	Modelling the Problem	150
6.4.1	Physical model of the system	150
6.4.2	Numerical robustness	152
6.5	Design procedure	153
6.6	Optimal cyclic operation	154
6.6.1	Objective function	156
6.6.2	Plant parameters for optimisation	159
6.6.3	Control parameterisation and time discretisation	160
6.6.4	Constraints	161
6.7	Development of candidate structures	164
6.7.1	Steady-state based bounding	166
6.7.2	Steady-state optimisation of the reactor network	167
6.7.3	Results from steady-state optimisation	169
6.8	Optimal cyclic operation of the candidate structure	172
6.8.1	Implementation	173
6.8.2	Results for the candidate structure	175
6.8.3	Review of bounding assumptions	177
6.8.4	Cyclic operation in the presence of non-ideality	180
6.9	Optimal cyclic optimisation under uncertainty	182
6.9.1	Bounding the uncertainty using near zero concentrations at t_f	184
6.9.2	Propagate all uncertainty into one state	189
6.9.3	Cyclic optimisation formulation incorporating a relaxed storage end-point	192
6.10	Discussion	193
6.10.1	Summary of key results	193
6.10.2	Generality of results	194

6.10.3	Alternative operating strategies	196
7	Conclusions and future work	199
7.1	Results and conclusions	199
7.1.1	Development of the design procedure	199
7.1.2	Case studies	200
7.1.3	General conclusions	201
7.2	Areas for future work	201
	References	203
A	Review of operating procedure synthesis literature	210
A.1	Valve sequencing operations	210
A.2	State transition sequences	212
A.3	Artificial intelligence methods	213
A.3.1	System decomposition	214
A.3.2	Modelling	214
A.3.3	Constraint representation	215
A.3.4	Operating procedure planning	215

List of Figures

2.1	Distributed chemical manufacturing	25
3.1	“Design funnel” to be used in process design procedure.	53
3.2	Sources of operating variability affecting distributed plant.	55
4.1	Two feasible (cyclic) schedules	74
4.2	Inventory levels for alternative schedules	75
4.3	Cycle summarising interactions between process design, product transitions and production schedule	77
4.4	Procedure for continuous multiproduct process design proposed by Cantrell and Luyben	83
4.5	Multiproduct process case study	86
4.6	Time profile for system described by cycle time model	92
4.7	Time discretisation used in multiproduct plant case study	97
4.8	Cyclic profiles for reactor composition and product storage in multiproduct plant	104
4.9	Cyclic profiles for product compositions in multiproduct plant	105
4.10	Damped reactor composition response	108
5.1	Wastewater plant design problem	113
5.2	CSTR unit for wastewater treatment	117
5.3	Functional form of diverted off-spec flow	120
5.4	Diagram showing cyclic behaviour of concentration disturbance (top) and the collapsed cycle which is implemented (bottom).	123
5.5	Required damped concentration response	129
5.6	Cyclic profiles for wastewater process base case	134
5.7	System one: profiles during operating cycle - 1	138
5.8	System one: profiles during operating cycle - 2	139
5.9	System two: profiles during operating cycle - 1	141
5.10	System two: profiles during operating cycle - 2	142
6.1	Process building blocks	149
6.2	Diagram showing original operating cycle (top) and the collapsed cycle which is implemented (bottom).	155
6.3	Screening dynamic process alternatives through steady-state optimisation	166
6.4	Variation of process economics with fraction of fresh feed to first reactor	172
6.5	Candidate process for cyclic optimisation	173
6.6	Typical state profiles from cyclic optimisation of candidate structure	178

6.7	Control profiles from cyclic optimisation of candidate structure	179
6.8	Drift in process states during idle periods for the relaxed endpoint problem	182
6.9	State profiles from solution of relaxed problem - 1	187
6.10	State profiles from solution of relaxed problem - 2	188
6.11	Decomposition of process structure	191
6.12	Annotated cyclic operation	191
6.13	Parallel arrangement for batch operating strategy	198
A.1	Goal hierarchy proposed by Rivas and Rudd	211
A.2	State transition graph proposed by Ivanov	213

List of Tables

3.1	Nominal solution for different iterative procedures	63
3.2	Feasible vs. infeasible path methods	65
4.1	Model parameters for multiproduct plant case study	88
4.2	Reactor parameters for economically optimal multiproduct plant	101
4.3	Cost breakdown for optimal multiproduct plant design	102
5.1	Optimal treatment tank sizes and costs using Walsh approximation	122
5.2	Optimal design for wastewater treatment system base-case	133
5.3	Optimal design for system one: treatment system	136
5.4	Optimal design for system one: diversion system	137
5.5	Optimal design for system two: treatment system	140
5.6	Optimal design for system two: diversion system	140
5.7	Summary of wastewater design results	143
6.1	Reactor parameters for case study	152
6.2	Optimal steady-state design for 1 CSTR	170
6.3	Optimal steady-state design for 2 CSTRs	170
6.4	Optimal steady-state design for 3 CSTRs	171
6.5	Summary of optimal reactor network costs	171
6.6	State variables at shutdown and production steady-states	180
6.7	Control values during idle period	190

Nomenclature

A_H	–	heat transfer area
A	–	amplitude of sinusoidal controller response at t_1 or t_f
A^{max}	–	maximum amplitude of controller response at t_1 or t_f
c_1	–	equality point constraints
c_2	–	inequality point constraints
C_0	–	value of concentration at time t_f
C_{A_i}	–	concentration of A in product i
C_A^{sp}	–	concentration setpoint
$C_{CSTR_{op}}$	–	operating cost of reactor
C_{CSTR}	–	capital cost of reactor
C_{base}	–	cost of feasible base case process
$C_{c_{mix}}$	–	concentration of effluent leaving perfectly mixed section of reactor model
$C_{capital}$	–	capital cost
$C_{diversion}$	–	cost of diversion tank
c_{net}	–	net concentration of acids and bases
c_{sig}	–	significant violation of path constraint
$C_{dt,i}$	–	outlet concentration of ith lag block in deadtime approximation
C_d	–	diversion tank concentration
C_{in}	–	feed concentration to unit
C_{min}	–	lower bound on cost of process
C_{op}	–	total annualised operating cost
C_{out}	–	concentration of effluent from reactor
$C_{p,J}$	–	liquid coolant heat capacity
$C_{product}$	–	product concentration from reactor network storage tank
C_p	–	liquid reactant heat capacity

C_r	–	alkali reagent concentration
C_{sd}	–	cost of shutdown phase in operating cycle
C_{store}	–	cost of storage tank
C_{su}	–	cost of startup phase in operating cycle
$C_{transition}$	–	transition cost
C_v	–	valve constant
d	–	set of design variables
f	–	residual function of DAE equations
F_c	–	coolant flowrate to reactor cooling jacket
F_c^c	–	coolant flowrate calculated by temperature controller
F_c^{range}	–	range of coolant flowrates
F_c^s	–	scaled flowrate of coolant
F_T	–	total flowrate to process
$F_{discharge}$	–	flowrate of material discharged by wastewater plant
F_d	–	diversion flowrate to wastewater treatment reactor
F_{in}	–	feed flowrate to unit
$F_{offspec}$	–	flowrate of off-spec material
F_{out}	–	exit flowrate from reactor
$F_{product}$	–	product flowrate from reactor network storage tank
F_r	–	reagent flowrate to wastewater treatment reactor
F_r^c	–	calculated flowrate of alkali reagent
F_r^{range}	–	range of alkali reagent flowrates
F_r^s	–	scaled flowrate of alkali reagent
H_{store}^{max}	–	maximum molar holdup in storage tank
ΔH_R	–	heat of reaction
$H_{reactor}$	–	molar holdup in reactor
h	–	unit inventory holding cost
$I(t)$	–	inventory level at time t
J	–	general scalar objective function
J_0	–	primary objective function in multiobjective optimisation
J_1	–	secondary objective function in multiobjective optimisation
k	–	reaction rate
k_0	–	Arrhenius rate constant
k_{bg}	–	rate of background reaction
k_v	–	valve position

k_w	—	dissociation constant of water
K_{clip}	—	controller clipping and desaturation parameter
K_c	—	proportional level controller gain
K_c^C	—	composition controller gain
K_c^T	—	temperature controller gain
N_{cycle}	—	number of plant cycles per year
n	—	number of CSTRs
pH	—	$-\log_{10}(a_{H^+})$
p_i	—	production rate for product i
$\mathbf{P}$	—	vector of production rates in cycle time model
q	—	inequality path constraints
Q	—	heat removal rate
$Q_{product}^{bg}$	—	background rate of production
Q_{mu}	—	make-up flowrate calculated by proportional level controller
r_i	—	demand rate for product i
R	—	molar gas constant
$\mathbf{R}$	—	vector of demand rates in cycle time model
s_{ij}	—	transition time from product i to product j
t	—	time
t'_d	—	delay between disturbance and control reaching the output
$t_{product}$	—	length of time of demand for product
t_0	—	initial time in cycle
t_1	—	time to first steady-state in cycle
t_{cycle}	—	cycle time
$t_{d_{mix}}$	—	mixing delay in fractional tubularity model
t_{drain}	—	drain time between $t_{restart}$ and t_0
t_d	—	deadtime period
t_f	—	final time of dynamic optimisation
t_{idle}	—	idle time between production cycles
t_i	—	production time for product i
Δt_j	—	length of jth discretisation interval
$t_{restart}$	—	time at which process is restarted after idle period
T_p	—	period of sinusoid
$\mathbf{T}$	—	vector of production times in cycle time model
T_0	—	value of temperature at time t_f

T_{in}	–	temperature of feed to reactor
T_R^{meas}	–	lagging reactor temperature measurement
T_R^{sp}	–	reactor temperature controller setpoint
$T_{c,in}$	–	temperature of cooling water entering reactor cooling jacket
$T_{c,out}$	–	temperature of cooling jacket
T_c	–	mean cooling water temperature
T_{out}	–	temperature of reactor
u	–	time varying optimisation parameter (control)
U	–	heat transfer coefficient
V_J	–	volume of reactor cooling jacket
V_{R_h}	–	volume at which reactor exit flowrate switches to linear function
V_{R_l}	–	lower limit on CSTR volume
V_R	–	volume of reactor
$V_{c_{mix}}$	–	volume of perfectly mixed section in fractional tubularity model
$V_{diversion}$	–	diversion tank volume
V_{opt}	–	optimal CSTR volume from Walsh approximation
V_{store_l}	–	low limit of volume in storage tank
V_{store}	–	level of storage in storage tank
V_{total}	–	total volume of system
$V_{offspec}$	–	accumulated volume of off-spec material
x	–	state variable
$\dot{x}$	–	time derivative of state variable
x_0	–	initial value of state
x_r	–	value of state at time $t_{restart}$
y	–	algebraic variable

Greek symbols used

ϵ	–	small number (defined more precisely where required)
ω	–	frequency
Ω	–	set of possible frequencies
ϕ_1	–	times in Φ for equality point constraints
ϕ_2	–	times in Φ for inequality point constraints
Φ	–	discrete subset of optimisation time horizon
θ	–	sum of transition times
ρ	–	liquid reactant density

ρ_J	–	liquid coolant density
ρ_m	–	molar density
ξ_C	–	error in composition controller
ξ_T	–	error in temperature controller
ξ_I	–	integral error in proportional-integral controller
τ_I	–	reset time in proportional-integral controller
τ_T	–	lag time in temperature control loop
κ	–	constant factor with value greater than one
δ_c	–	required disturbance attenuation
τ_i	–	residence time of reactor i

Subscripts used

i	–	i th storage tank or i th reactor
-----	---	---------------------------------------

Superscripts used

max	–	maximum value
-------	---	---------------

Acronyms used

CSTR	–	Continuous stirred tank reactor
DAE	–	Differential algebraic equation
DFM	–	Design for manufacture
DCMF	–	Distributed chemical manufacturing facility
DPI	–	Dynamic performance index
EPI	–	Economic performance index
HEN	–	Heat exchanger network
MILP	–	Mixed integer linear program
NLP	–	Nonlinear program
MINLP	–	Mixed integer nonlinear program
PI	–	Proportional-integral control

Chapter 1

Introduction

This chapter outlines the overall objectives of the work presented in this thesis. The motivations for initiating this work are first discussed. Following this, novel aspects of the design problem to be addressed are considered and differentiated. Finally, an overview of the thesis is presented.

1.1 Rationale

Recent economic, safety, logistic and environmental trends within the process industries are encouraging research into the design and operation of distributed chemical manufacturing facilities. Such plants manufacture chemicals for a customer *in situ* rather than transporting the same material from a centralised plant.

The implications of building a distributed chemical manufacturing facility on what is likely to be a commercial, light-industry site raise a number of important issues. Facilities for effluent disposal will be limited so the plant should produce environmentally inert effluents. Close proximity to other businesses and the general public require the plant to have a high degree of safety, or be intrinsically safe. Similarly, feedstocks to the process should be non-hazardous and inventory within the plant kept as low as possible. Plant operation in an environment away from trained personnel demands automatic operation. Finally, customers will require the plant to be capable of constantly tracking their demands. The physical realisation of such a plant would need to bring together research from diverse fields, such as process chemistry, catalysis and mechanical and materials engineering. What is particularly important and challenging in the context of process systems engineering research is the impact of the distributed plant's unique operating environment on its design and operation. This environment is likely to be highly dynamic

in nature, with the plant being subjected to frequent and potentially large changes in the demands placed upon it. Given that automatic operation is also a prerequisite, any distributed plant developed *must* have guarantees on the feasibility of its operation in this intrinsically transient environment. The focus of the work presented in this thesis will therefore be to consider the integrated design and operation of distributed chemical manufacturing facilities.

1.2 Objectives of the project

The motivations for this work have been set out above. With this in mind, a formal statement of the objective of this research is:

to establish a systematic procedure which facilitates the efficient development of economic process designs together with their corresponding optimal dynamic operation.

Although the design procedure will ultimately be applied to a restricted number of representative case studies it should be general in its approach. Experience in the area of distributed plant design is very limited. Because of this, it may therefore be difficult to clearly set out all design objectives at the problem definition stage. The design procedure developed in this work should recognise this and be sufficiently flexible to accommodate, what might initially be, an incomplete problem specification. The procedure should facilitate the creation of a tractable design problem for which robust and reliable solution methods exist. Ultimately, the feasible operation of the design in the face of operational uncertainty must be guaranteed. The design procedure should therefore provide mechanisms which enable the key dynamic effects impacting on the design to be meaningfully considered in a straightforward way at the design stage.

As stated above, there is very little knowledge in the area of distributed plant design. Few, if any, true examples of the types of plant described earlier are to be found in industry. Given this, suitable case study examples must be chosen which are relevant in the absence of an actual realisation of a distributed plant. The examples used to demonstrate the application of the integrated design procedure developed in this work should therefore possess a number of desirable characteristics:

- the case study examples chosen should be established as appropriate for use in the context of an actual distributed manufacturing system,
- the example problems should have an inherently dynamic character,

- the case studies should be of a size which is representative of “real” industrial problems.

It is hoped that by analysing the solutions obtained for the case study examples, it will be possible to highlight any important qualitative features of these designs. The aim in doing this is to try and provide a degree of engineering insight in what is a relatively new field of process systems research.

1.3 Overview of thesis

Chapter 2 discusses traditional chemical manufacturing practices and explores the key motivations for building large, centralised chemical plants. The concept of a distributed chemical manufacturing facility as an alternative production system is then introduced. Areas of research required for the realisation of a distributed plant are identified. Process design issues are highlighted as the focus of this work. The operational nature of the distributed plant is discussed and arguments for an integrated design approach are presented. A review of approaches to the integrated design problem is then presented. The conclusion drawn is that the procedure used for the design of distributed plants should explicitly consider the dynamic process operation.

Chapter 3 discusses the development of the design procedure applied to the case studies presented. Arguments are presented which lead to a procedure based on mixed integer nonlinear programming being rejected. A two-stage decomposition of the design problem is instead favoured for this work: development of a candidate structure followed by determination of the optimal design and operation of that candidate. Incorporation of the potentially complex distributed plant operation into a design procedure is considered. The simplification of cyclic operation is made to facilitate this. Tools for determining the optimal design and cyclic operation of processes are discussed and dynamic optimisation found to be most appropriate for this work. An introduction to dynamic optimisation is given and solution methods are reviewed. Feasible path control parameterisation is selected for use in this work. Finally, the case studies to be considered in the remainder of the thesis are introduced.

Chapter 4 considers the design and cyclic operation of a continuous multiproduct plant. The flexibility of these processes makes them highly appropriate for use in a distributed manufacturing system. Traditional sequential approaches to the design of continuous multiproduct plants are reviewed. Interactions between the design, scheduling and dynamic product transitions in these plants are highlighted. This motivates the use

of the integrated approach presented in chapter 3 for the design of multiproduct plants. Such plants are typically operated cyclically in industry, which enables the second stage of the proposed design procedure to be effectively demonstrated.

Chapter 5 considers the integrated design of chemical wastewater treatment facilities. The case study is motivated by the increasing development of compact treatment plants for localised wastewater treatment. Typical chemical treatment plants consist of a main treatment system and a diversion system which provides local containment as a fail-safe. This chapter aims to examine potential economic benefits from integration of the diversion system design and operation with that of the continuously operating treatment system. A feasible base case candidate process is rapidly developed by application of existing screening tools and then dynamic optimisation. The subsequent development of a simple economic bounding property enables the design search space to be dramatically reduced. The optimal process design and operating policy is found by solution of only two further dynamic optimisation problems.

Chapter 6 considers the optimal design and operation of a distributed process plant capable of instantaneous response to a customer's intermittent demand for product. The case study represents a "worst-case" product demand profile, since the process must be capable of automatic startup and shutdown. A steady-state bounding heuristic is applied which enables a candidate structure to be developed using steady-state optimisation. The optimal design and cyclic operation for this candidate are then found using dynamic optimisation. A perfect hold-state is found to exist for the problem considered which enables near instantaneous switches between the shutdown and production steady-states. Introduction of non-idealities in the process leads to a discussion on cyclic operation in the presence of uncertainty. Finally, satisfying the intermittent product demand using batch operation of a process is considered as an alternative.

Chapter 7 reviews the work presented in the thesis and draws some general conclusions. Finally, areas for future work are identified.

Chapter 2

Distributed manufacturing and the integrated design problem

2.1 Introduction

In this chapter the concept of the distributed chemical manufacturing facility (DCMF) will be introduced. The motivation for moving towards this type of plant will be considered by assessing benefits to be gained with respect to current chemical processing practice. The design challenges presented by the distributed plant will then be discussed in the context of an integrated design problem and a review of literature in this area presented.

2.2 Distributed chemical manufacturing

2.2.1 Current chemical processing practice

Chemical processing is generally carried out on large, purpose-built sites. In recent years attempts have been made to reduce the environmental effect of such facilities. However, by their very nature they remain as noisy and obtrusive features in any landscape. The poor public perception of large chemical facilities is often compounded by concerns about the safety of such sites and of their impact on the surrounding environment. Often these worries are misguided. An example is the common misconception that the plumes leaving cooling towers are composed of something more sinister than water vapour. However, incidents such as Flixborough remind us of the potential devastation that chemical plants can cause when something does go wrong. There would seem to be two key factors which have driven the manufacture of chemicals to large, single sites:

- economics and,
- safety.

The $\frac{7}{10}$ power law describing cost as a function of size is known to operate for chemical processing systems above a certain size threshold. Making use of economies of scale has probably been the single strongest motivating force in moving towards bigger and bigger chemical plants. In general a chemical manufacturing site will contain several plants each producing a different product. The grouping of many plants together onto a common site also offers advantages and is a further economic reason for moving to single large sites. It allows the sharing of utilities, feedstocks, products, and effluent treatment facilities. In addition, energy integration between the processes is made possible.

A second driver for producing chemicals on purpose-built centralised sites is that of safety. Centralised production enables the manufacturing process to be located away from major urban population centres. The use of a single site provides an environment in which the associated hazards may be more effectively managed and controlled.

Of course, when a chemical is produced at a single centralised site, it must be distributed to the customers who require it. This issue raises two questions of current practice:

- are the logistics involved in transportation economical?
- what are the safety implications when potentially hazardous materials must be transported?

An alternative production method which directly addresses these issues is to distribute the plant to the customer, rather than the product. This is shown in figure 2.1. In fact, this approach challenges more than just the logistics of centralised production as will be seen in the next section.

2.2.2 The distributed chemical plant

Recently in the literature, the concept of a miniaturised, distributed chemical production facility has been discussed (Benson and Ponton, 1993). There are instances when a customer requires only relatively small amounts of product as a feedstock for their own business activity at their site. A miniaturised plant located at this site, piping product “over the fence” to customers, may serve them better than the distribution of the same product from a central production facility. The basic concept is a simple one. What is particularly important and interesting in the context of engineering research is

the impact of the plant's operating environment on its design and operation. It is likely that the distributed plant will be operated on a commercial, light-industry site. Such a site will, presumably, possess very limited facilities for effluent disposal. This means that the plant should produce no hazardous waste. Any waste that is produced should be as environmentally inert as possible. The plant will be located in close proximity to other businesses and the general public. In view of this it should be designed with a very high degree of safety, or be intrinsically safe. The safety and environmental issues above also require that the feedstocks to the process should be non-hazardous. If potentially hazardous feedstocks must be used, they would have to be transported in a safe form and regenerated within the plant. For example, chlorine might be supplied as a salt rather than liquefied gas. Ideally, there would be no intermediate storage of either feedstocks or products. Turning now to operational aspects of the distributed plant, it is envisaged that the plant would be self-contained. The unit will be delivered, plumbed into the feed supply, and in the event of a breakdown, it would be replaced and not repaired. However, the routine servicing and replacement of equipment such as catalysts and filters, which would be expected to degrade during use, should be possible. Operation in an environment away from trained operators and engineers requires that the plant operates in a totally automatic fashion with minimal manual intervention. Economic considerations dictate that mass production manufacturing methods will be needed, and this inevitably requires the production of a standard unit. From the customer's point of view, this unit would be expected to deliver product of the required purity, at the desired rate, on demand. Further, the plant must be able to continually track changing demands. The plant should also be capable of operating feasibly under the effects of external influences, such as variability in the quality of raw materials. It is clear that achieving all of these operational and design aims presents a great challenge. It can be seen that the implications of distributing the production of chemicals force us to design a safer and more environmentally acceptable plant. Of course, it will not always be possible or sensible to move production of certain chemicals away from conventional sites. However, when distribution is appropriate it would seem to challenge centralised production in a number of areas. This will be discussed in the following section.

2.2.3 Centralised versus distributed chemical manufacture

The primary justification for building large, centralised sites for chemicals manufacturing was previously given as the application of the $\frac{7}{10}$ cost-size power law. The building of distributed chemical plants would achieve economies of scale in a different




Figure 2.1: Distributed chemical manufacturing

way. Mass production of individual, identical units merely applies economies of scale arguments to the plant manufacturing process rather than the plant itself. The “one-off” nature of conventional plants requires the use of custom-built components which are, by definition, expensive. The distributed plant would use large numbers of standard components, which would inevitably be cheaper. Capital cost comparisons between the economics of the distributed plant and the centralised plant might be favourable.

Building one-off designs can also lead to problems in guaranteeing the plant’s performance. Once the plant has been commissioned, often a time consuming process, much time may need to be spent adapting the plant in order to achieve design performance. The distributed plant would require a more thorough initial design before committing it to the production line, but building a prototype plant would be possible. Rigorous design and prototype testing would allow much better guarantees on the performance of distributed plants leaving the production line. The use of production line manufacturing may impinge on the design process. As an example, the modern car is (arguably) a more complex piece of equipment than any chemical plant. Despite this, it is assembled from thousands of individual components and driven off the production line. Car manufacture is a good example of *design for manufacture* (DFM). An important part of DFM philosophy is to involve manufacturing from an early stage in the design and to tolerate no engineering changes when the design has been fixed, i.e. to get it right first time. Conventional chemical plant ideology tends to separate design and manufacture, but the application of DFM techniques will probably be needed if distributed plants are

to become a commercial reality.

The environmental and safety benefits of the distributed plant are obvious. First, transport of hazardous material is removed and no hazardous material is discharged by the plant. Second, the size of the plant and necessarily low holdups of inventory that involves means that even a total release of material may not create a significant toxic hazard. In addition, it may be possible to contain any explosion which might occur, within the plant itself. Finally, as a result of their size, the distributed plant would be neither as noisy, nor as intrusive as conventional plants.

Where distributed chemical manufacturing facilities can be applied they would seem to offer many benefits over conventional practice. It is clear, however, that developing such plants will encompass a wide range of research areas.

2.2.4 Areas of research interest

1. Process chemistry and catalysis

A direct result of having to perform reactions at moderate temperatures and pressures will be the need to use advanced catalysts. In addition, restrictions on the types of feedstocks used may require novel reaction paths, not currently used in industry, to be developed.

2. Process unit operations

Finding the most suitable equipment to carry out reaction, separation and heat transfer duties is important. The use of novel unit operations which are not favoured on large-scale plants for reasons of capacity and economics may be far more suited to use on a distributed plant. For example, membrane technology or centrifuges may be used for separation tasks where traditional distillation would be inappropriate.

3. Materials engineering

There is much scope to use advanced materials of construction and modern construction techniques.

4. Process systems engineering

The role of process systems may be thought of here as, finding the “best” type, topology and operation of unit operations which, when combined, give the desired transformation of inputs to products. Best will usually imply the cheapest design meeting the specification.

This work will focus only on the process systems engineering aspects of the overall design problem, *i.e.* those elements related directly to the design and operation of the process. In the following section, these elements are reviewed and their impact on the process design problem discussed.

2.2.5 Process design issues

The general functionality of a distributed plant has previously been discussed. The actual realisation of such a plant was shown to incorporate a diverse range of technologies. This section focuses on the design problem from a process systems engineering viewpoint and attempts to identify those aspects which make this design problem distinctive. Consideration of the design requirements of such plants leads to an understanding of the approach which will be needed to successfully tackle this design problem.

The process systems issues may be distilled from the general functionality required from a distributed plant and summarised as follows:

1. Since the distributed plant is required to have the minimum environmental impact, it should be designed to have as small a footprint as possible. This is desirable, both to keep inventories within the plant as low as possible and to reduce the visual impact of the plant. Reducing the total volume of the system will be an aspect of the design problem.
2. The plant is required to operate in an automatic fashion since it will operate in an environment away from trained personnel.
3. The plant will be subjected to continual variability in raw material specifications and flowrates during its operating life and must be designed to operate feasibly in the presence of such variations.
4. The plant will operate under changing demands both on the purity and production rate of products. It must be able to exactly track these demands over the whole life of the plant. When product is required it must be available on demand and to specification.

The operational nature of a distributed plant is likely to be characterised by frequent changes in operating point which may lie anywhere within the plant's design envelope. In the extreme (but highly plausible) case when demand is intermittent, the plant must be capable of automatically starting up to deliver product and shutting down when demand

ceases. Given the further requirement that the plant's response to customer demands should be instantaneous, a level of stored product may be required.

In the light of this, the distributed plant must be regarded as an inherently dynamic process and any procedure developed in this work for the design of distributed plants should therefore *simultaneously consider both the time-varying operation of the process and its design*. The term integrated design has been coined to describe problems in which process operation forms an integral part of the overall design problem. The development of an integrated design procedure for this work will be considered in chapter 3. Prior to this, a review of the literature in the area of integrated design is presented to explore possible approaches to the problem.

2.3 Review of integrated design literature

This section reviews literature in the area of "integrated design", that is, process design procedures which also consider the process operation. In order to place this problem in its proper context, the process design problem is discussed from a historical perspective. Traditional approaches to the process design problem are outlined. Following this, approaches to the steady-state synthesis problem are briefly discussed. A number of diverse aspects which may be considered as operability issues are mentioned and those of interest in this work are highlighted. The initial historical review presented is intended to give sufficient background to the problem in question. However, where possible, the reader is pointed to more extensive references. The final two parts of this chapter deal exclusively with approaches to the integrated design problem itself.

2.3.1 Historical perspective on the process design problem

Traditional approaches to the design of chemical processes have been sequential in their nature. The design methods employed typically contain two distinct steps:

1. Process synthesis and design step. The aim of this step is typically to determine the economically optimal steady-state process design.
2. Process operability step. Having established the design of the process in step one the operability of the process is analysed. This is a broad term which encompasses a number of aspects relating to the operation of the plant:
 - (a) Flexibility: the ability of the process to operate successfully at steady-state under different operating conditions which arise through variations in throughput,

feedstocks, *etc.*

- (b) Controllability: the ability of the plant to reject disturbances (regulatory control) and move between operating points (servo control).
- (c) Reliability and safety verification: the ability of the plant to operate safely and reliably.

2.3.1.1 The process synthesis problem

The activity of process synthesis may be thought of as generating design alternatives and selecting the better ones based on incomplete information.

(Westerberg, 1989).

The synthesis problem is complex because of the very large number of possible alternative structures which might be considered (Douglas (1988) suggests some 10^4 to 10^9 alternatives typically exist). The challenge in the conceptual design of a chemical process is therefore to develop systematic procedures for the generation and quick screening of many process alternatives. The literature discusses two clearly distinguishable approaches to the process synthesis problem:

1. hierarchical approaches,
2. algorithmic approaches.

This vast body of work will only be summarised here, but a number of comprehensive reviews have been published, *e.g.* Westerberg (1980), Stephanopoulos (1981).

Hierarchical approaches generally use decomposition of the problem to prevent a potentially explosive growth in the search space and a series of rules to guide the search direction. At each level of the decomposition more detail is added to the problem description. Perhaps the best known of the hierarchical approaches to process synthesis is that proposed by Douglas (1988). Since it is fairly typical of this type of approach, it is summarised here as an example. The Douglas procedure is essentially a hierarchy of decisions as follows:

1. Batch *vs* continuous operation,
2. Input-output structure of the flowsheet,
3. Recycle structure of the flowsheet,
4. General structure of separation system,

5. Heat-exchanger network synthesis.

The approach enables good preliminary flowsheets to be determined for a limited (continuous vapour-liquid) class of plants. Other approaches to the hierarchical process design procedure, including attempts to automate the method, are reviewed by Mizsey and Fonyo (1990b). One appealing feature of such methods is that they are driven by the engineer. Not only does this mean that the engineer's experience is fully exploited, but it facilitates the incorporation of on-going changes in the problem definition.

Alternatively, the synthesis problem may conceptually be formulated as a discrete-continuous optimisation problem, specifically as a mixed-integer nonlinear program. Such an approach uses a process superstructure in which all possible design alternatives are embedded. A mixed-integer nonlinear programming problem describes the synthesis problem as the minimisation of some objective (usually economic), subject to a series of constraints on the process (*e.g.* maximum temperatures or vessel sizes), by optimisation both of continuous design variables (flowrates, temperatures, vessel sizes) and of the zero-one (integer) variables which determine whether a process unit appears in the flowsheet. A single-pass solution to the process synthesis problem enables structural and continuous design decisions to be optimised simultaneously, and therefore offers the potential for superior solutions when compared to hierarchical methods. However, in addition to being highly computationally intensive, such methods may suffer as a result of excluding the design engineer entirely from the synthesis procedure. Potential drawbacks of MINLP techniques for dynamic design will be discussed in more detail in section 3.2.1.

Having generated a small number of conceptual flowsheets through one of the two synthesis methods outlined above, rigorous simulation calculations may be carried out to arrive at a final design. It should be noted at this point that, although the synthesis methods described above deal with the steady-state design problem, the procedures described can be carried over to the area of dynamic design.

Having determined the process flowsheet, the next step would typically be to analyse the operability of the process and address the control needs of the process. These issues will be dealt with in the following section.

2.3.1.2 The process operability problem

Recall that process operability may be considered from many perspectives: controllability, flexibility, reliability and safety. This section briefly considers the ways in which operability aspects are addressed in traditional process design procedures.

2.3.1.2.1 Flexibility In general, any process design will be subject to *uncertainties* and *variability*, such as changes in feed specifications, measurement errors, changes in process parameters *etc.*. It is important to be sure that the design is capable of operating to specification in the presence of this potential variability. Traditional practice in chemical process design has been to incorporate design margins to cope with this type of variability. The oversizing of distillation columns is a good example. Over-design is perhaps made possible by virtue of the fact that, in general, large continuous processes tend to operate about a nominal steady-state point and often at the maximum design throughput. In an attempt to reduce over-capacity and, as a result reduce costs, a large part of process synthesis research has considered the issue of steady state process flexibility. The aim here is essentially to ensure rigorous feasible operation of the process over a bounded range of uncertain parameters. A review of work in this area has recently been carried out by Grossmann and Straub (1991). The ability of a design to operate satisfactorily despite uncertainty and disturbances has already been highlighted as important in this work in section 2.2.5. The subject of uncertainty in an integrated design setting will be re-visited later in section 2.3.4.

2.3.1.2.2 Controllability Having established the basis of a process design which is capable of operating in the desired operating envelope, a control system must be added which is capable of holding the process at the desired operating point in the face of disturbances (regulatory control), and of moving the plant to new operating points (servo control). In terms of the process operation, controllability may be defined as: *the “best” dynamic performance (setpoint following and disturbance rejection) achievable for a system under closed loop control* (Morari and Perkins, 1994). The aim of controllability analysis is to determine a control structure for the given process design which provides the best dynamic performance. Obvious analogies may be drawn between the approaches taken in this control structure selection procedure and those used for the original process synthesis. Essentially, two approaches to this problem have been discussed in the literature:

1. semi-heuristic,
2. algorithmic.

As Morari (1992) has observed, the greater part of the controllability analysis has been focussed on the “evaluation” problem, *i.e.* developing low-effort tools which enable alternative control structures to be rapidly compared and screened. This problem is comparable to the hierarchical synthesis procedure described earlier (*i.e.* to develop systematic

procedures for the generation and quick screening of many process alternatives). Over the years many indicators have been developed for the purpose of ranking alternative plants and control schemes. Examples are the relative gain array and singular value analysis. The aim of these tools is generally to determine whether there are any plant characteristics which prevent perfect control, *e.g.* inverse response, and rank options based on this. Although such tools give useful insights, they do not relate directly and unambiguously to real performance requirements. Indeed, as Morari has noted (1992):

Unfortunately, we often do not quite understand yet when and how these indicators should be used. Nevertheless they are applied widely and indiscriminately and lead to erroneous conclusions about controllability.

Following the trend of process synthesis techniques, an algorithmic approach to control structure selection has also been proposed. Narraway (1992) proposed a procedure for control structure selection based on economics. The approach recognises the ambiguity of the controllability indicators described above and uses linear dynamic economics to enable direct comparison and trade-off with the steady-state process economics. The problem is posed as a mixed integer linear programming problem (MILP) with the assumption of perfect control. Integer variables in the MILP solution determine the presence of a manipulated variable or measurement in the final control structure. One drawback of this approach is the perfect control assumption. In light of this, the basic algorithmic approach has recently been extended by Heath (1997) to include realistic proportional-integral controllers which are tuned as part of the overall control structure selection process.

2.3.1.2.3 Reliability and safety Finally, reliable and safe operation of the plant must be ensured. Plant reliability and safety will be important issues in the design of a distributed plant which must operate autonomously and away from trained operators. Safety requirements, with the exclusion of abnormal events, will be considered as one of the factors defining the operating envelope of the plant. It is considered that reliability issues can be more suitably dealt with by the interaction of the design engineer throughout the evolution of the process design. Reliability and abnormal events will therefore be neglected in this work to provide a tractable problem which captures key operability aspects.

2.3.2 Approaches to the integrated design problem

The historical background presented above has described traditional approaches to the design of chemical processes. The approach is sequential, which by definition means that design of the economically optimal process precedes any consideration of the process operation. It is becoming widely accepted that process design problems cannot be considered using an analysis based purely on steady-state operation. Despite acknowledgement of this, Morari (1992) has recently observed that, when compared to steady-state synthesis research: *far less work has been reported on process synthesis taking account of dynamic requirements*. The section which now follows will present a review of the literature which has considered the problem of including dynamic operation as an integral part of the design problem - the so-called integrated design problem.

The research presented dates back to 1975, spans over twenty years and essentially takes two viewpoints. A number of authors, *e.g.* Nishida et al. (1975), take the view that process systems are inherently dynamic and should be considered as such, for example: startup operations and changes from one set of operating conditions to another. They conclude that:

It is therefore of interest to take into account the effect of dynamic behaviour at the process synthesis stage rather than merely synthesising the system assuming steady-state operation and compensating the effect of dynamics by means of process control.

In other words the dynamic operation of the process should be considered explicitly within the process design problem.

Other authors take a somewhat different view in their approach to incorporating process operation in the design problem. The importance of the steady-state process design on the controllability of the process is stressed. Stronger dynamic interactions between process units which arise as a result of more efficient steady-state designs (*e.g.* reduction in storage tank sizes and increased integration of raw materials and energy) which can cause operational difficulties, are often cited as examples. As Lenhoff and Morari (1982) note:

Tighter design margins bring about better steady-state performance together with loss of flexibility.

This second approach typically considers the controllability (see section 2.3.1.2.2) of a given process design about some steady-state. The integrated design problem here aims to develop economic steady-state designs which also have favourable dynamic characteristics.

as judged by the application of controllability indicators at the synthesis stage. These approaches therefore inevitably revolve around systematic methods to trade-off steady-state economic performance with controllability.

The review which follows will use the two classifications described above as a basis to discuss the literature.

2.3.3 Integrated process design based on controllability

Because of disturbances and changes in operating conditions, real chemical processes are never at steady-state. However, Morari (1983a) observes that it is typically economically superior to operate at steady-state and keep deviations from it small. The view that the process is essentially at steady-state, but must be operable about this steady-state, characterises the integrated design approaches considered in this section. The review presented discusses more recent approaches to the integrated design problem which are based upon incorporating consideration of the controllability of a design at the synthesis stage. A number of extensive historical reviews discussing the development of controllability measures in the process design problem have been published (Morari and Perkins, 1994; Morari, 1992) and the reader is pointed to these for further information.

2.3.3.1 Resilient process plants

The term “*resilience*” was introduced by Lenhoff and Morari (1982) to describe the ability of the plant to move quickly and smoothly from one operating condition to another (including startup and shutdown) and to deal with disturbances. In the past, processes were generally easy to operate because significant over-design was included at the design stage (see section 2.3.1.2.1). Although resilience has always been a design factor, increased dynamic interactions resulting from steady-state process intensification required that it be more explicitly quantifiable at the design stage. Morari (1983b) stated that it was difficult to express dynamic resilience through any meaningful economic objective, such that the trade-offs between the operability of the process and its steady-state cost could be directly compared. He argued that what was of more interest to the design engineer at the early design stage was how operable (resilient) a given candidate process design would be. At that time, dynamic simulation was the only tool commonly applied to study the dynamic resiliency of a process. Such a tool was deemed by Morari to be inappropriate in the early stages of the design where there are many design alternatives. Morari showed that the dynamic resiliency of a process was determined by characteristics inherent in the system and was independent of the imposed controller structure and type.

Specifically, it was found that three system characteristics determined the resiliency of the process:

- nonminimum phase elements (*i.e.* time delays and right-half plane zeros),
- constraints on the control action,
- uncertainty (*i.e.* plant-model mismatch).

Linear frequency domain based methods for assessing the effect of each of these characteristics were developed. The aim of the tools developed was to facilitate the design of operable processes by allowing the efficient screening of the inherent dynamic characteristics of the design alternatives. Application of these methods may provide the design engineer with useful insights and help identify processes which, although economically attractive, are non-starters when judged on their operability. It should, however, be noted that these tools are based upon a process model linearised about a nominal steady-state and are not likely to be as useful when more extreme changes in operational mode must be analysed (plant startup operations, for example).

2.3.3.2 Heat exchanger network synthesis and controllability

One of the most widely studied process synthesis problems is the design of efficient heat exchanger networks (HEN). The HEN design problem may be stated as: “Given n_h hot streams to be cooled and n_c cold streams to be heated from specified supply temperatures to specified target temperatures, design the network of heat exchangers, heaters and coolers accomplishing this task at the least cost.” As discussed earlier, increased energy integration has led to stronger dynamic interactions between process units and has, in part, motivated an integrated approach to design and control. It is perhaps not surprising, therefore, to find that the design of resilient heat exchanger networks forms a large part of research in the integrated design and control area. As with the general process synthesis problem, research in the synthesis of HENs has provided both algorithmic and semi-heuristic solutions to the problem. Application of “pinch” techniques to determine “optimal” energy recovery systems is well established in industrial practice. Of all the work carried out in optimisation based process synthesis, the heat exchanger network problem has perhaps been most successful in finding its way into industrial application. A number of papers have been published which discuss the integrated design and control of HENs, based both on pinch and optimisation synthesis methods.

2.3.3.2.1 Semi-heuristic approaches In the 1980's Marselle et al (1982) considered the design and control of energy management systems. The work followed a series of papers in the area of resilient processing plants. The authors determined that it was not sufficient to concentrate on the steady-state efficient management of energy, but that the energy recovery system should also be resilient. The synthesis of a HEN for an isopropanol dehydration process was considered as a motivating problem. Pinch techniques are used to synthesise three networks whose resilience is then qualitatively analysed. The designs are found to be very sensitive to perturbations in the inlet conditions. In fact, with regard to the retrofitting of the base design required to improve its operability, the authors state that: *“the patchwork of additional heaters, coolers and auxiliary exchangers, could be obtained quite easily by an engineer without any knowledge about modern design methods”*. Their solution to the problem is to develop a semi-heuristic design technique coupled with a rigorous controller design method for resilient exchanger networks. In this case, resilience is defined by the network's ability to recover maximum energy for a specified disturbance range in the input flows and temperatures. The heuristic design method establishes a base case and develops optimal structures for four “worst cases” (*i.e.* maximum exchanger capacity, maximum cooling, maximum heating, minimum ΔT), trying to keep these designs as similar to the base case as possible. Basically, a network is designed for a small number of well selected extreme operating conditions. Given a network structure a steady-state control problem is then solved in two parts:

- regulatory control problem, *i.e.* meeting target temperatures by manipulation of valves on by-pass lines,
- network optimising control problem which minimises utility consumption.

Two sets of manipulated variables are chosen: one for regulatory control and the other for optimising control. The method does have a number of drawbacks, as Morari (1992) observed: *Marselle's approach leads to networks which are often not economical*. A further conclusion of the authors themselves is that dynamic effects in the network can be important and should therefore be incorporated into the design procedure.

2.3.3.2.2 Algorithmic approaches Georgiou and Floudas (1990) considered the problem of simultaneous process synthesis and control in heat exchanger networks. The specific issue addressed in their work was the design of heat recovery systems with minimum disturbance propagation. The authors observe that very little work has been carried out in the area of simultaneous process synthesis with controllability considerations. The

reason cited for this is a re-iteration of the observations of Morari, that the problem statement differs *significantly* when both problems are considered together. The optimality of a solution to the synthesis problem is typically measured by its cost, whereas the controllability is defined by criteria such as minimisation of loop interaction (RGA), minimisation of disturbance propagation, improvement of dynamic operability (SVD analysis) *etc.* Combining the two therefore presents issues of how to trade off economic and controllability criteria. Georgiou and Floudas make this trade off by defining the problem to be solved as follows:

Given the process synthesis goal of designing an economically optimal HEN and a process control goal to keep a set of control objectives at their set points in the presence of a given set of disturbances: determine whether a system exists that has low capital cost and features total disturbance rejection. If the cost is high, or total disturbance rejection cannot be achieved, determine the system which minimises disturbance propagation.

The problem is posed as a process superstructure which contains possible process alternatives. The total disturbance rejection criterion is incorporated into the synthesis problem using the concept of disturbability (a controlled output is disturbable if a path exists between the disturbance and this variable). The superstructure is represented as a structural matrix which shows whether streams, units and connections exist and gives cause and effect relationships between variables. Solution of a resulting mixed-integer linear program (MILP) determines a restricted superstructure which features total disturbance rejection. This structure is then optimised by solution of a nonlinear program (NLP) to find minimum investment cost. Assuming that total disturbance rejection is the most appropriate controllability criterion to be applied in this case (a point which is not addressed by the authors), few ambiguities in the approach arise when cheap designs featuring total disturbance rejection can be found. If this is not the case, however, the trade-offs between cost, energy recovery and disturbance rejection become much less clear. The approach is based upon structural controllability which make any implications for the dynamic behaviour of the system difficult to determine.

Papalexandri and Pistikopoulos (1994a) also consider the design of structurally controllable heat exchanger networks. Again, the work stems from the observation that sequential consideration of process synthesis and operability can lead to solutions which are far from optimal; for example, when retrofit design of the original process design is required to remove bottlenecks. The aims of the synthesis are similar to those of Georgiou and Floudas described above: to develop a network with the minimum annualised cost which is flexible within a range of uncertain parameters (inlet flowrates, inlet temperatures

and heat transfer coefficients *etc.*) and is structurally controllable (*i.e.* features total disturbance rejection). The solution method formulates the problem as a multiperiod hyperstructure in which all possible structural and operational alternatives are taken into account. Again, the procedure does not address the dynamic behaviour of the process. The issue of dynamics and control structure consideration is, however, discussed in a sequel paper by the same authors (Papalexandri and Pistikopoulos, 1994b). In this work the hyperstructure used previously is extended to allow exploration of control structure alternatives. Process dynamics are incorporated by the use of a variable delay matrix and variable gain matrix derived from a linear dynamic model of the process. The issue of process controllability is more thoroughly examined and several controllability criteria are incorporated into the synthesis problem as constraints to drive the MINLP solution to a network with favourable controllability characteristics. Diagonal dominance of the gain matrix, dynamic decoupling and “time-effective control” (the effect of the manipulated variable reaches the control objective before the disturbance) are explicitly incorporated within the synthesis as the required controllability criteria. The third criterion assumes that all disturbances are measured and that a feedforward control mechanism is used.

Having explored integrated approaches to the design of controllable heat exchanger networks, the following section considers more general research work on the design and control of whole chemical processes.

2.3.3.3 Integrated design of controllable chemical processes

A number of authors have recently addressed the problem of process design with simultaneous consideration of the process operability. The four pieces of work which will be discussed below share a common feature: a multi-objective approach is taken which attempts to simultaneously consider the steady-state process economics and some indicator of the controllability of the plant. Morari (1983b) suggests that this type of multi-objective approach trading steady-state performance against resilience is more appropriate than the use of a single scalar objective function.

Palazoglu and Arkun (1986) note that: “in designing a chemical plant, if only steady-state feasibility and economics are considered as design objectives, it is quite likely to end up with dynamically inefficient processes with serious operational bottlenecks.” Their solution is to incorporate the requirements on the dynamic operability into the design methodology by posing a multi-objective optimisation problem. The two objectives used are:

1. minimisation of the steady-state economics of the process,

2. minimisation of a robustness index (process condition number).

The solution of a multi-objective optimisation problem results in a set of solutions being obtained which are known as the pareto-optimal solutions, or noninferior set of solutions. Essentially, a number of solution points are obtained with the characteristic that no one objective function can be improved without a subsequent degradation in the other competing objectives. The multi-objective optimisation problem is reformulated to yield a scalar optimisation problem (referred to as the ϵ method) which can be solved using standard algorithms. The original problem is posed as:

$$\min_d J_0(d), \quad J_1(d, \omega), \quad \forall \omega \in \Omega \quad (2.1)$$

s.t. :

$$\begin{aligned} c_1(d) &= 0.0 \\ c_2(d) &\leq 0.0 \end{aligned} \quad (2.2)$$

where d are the design variables, Ω is the frequency range of interest and J_0 and J_1 are the economic and robustness objectives. This is reformulated as the scalar objective problem:

$$\min_d J_0(d) \quad (2.3)$$

s.t. :

$$\begin{aligned} J_1(d, \omega) &\leq \epsilon \quad \forall \omega \in \Omega \\ c_1(d) &\leq 0.0 \\ c_2(d) &= 0.0 \end{aligned} \quad (2.4)$$

Essentially then, the optimisation problem aims to minimise steady-state cost subject to a sufficiently small condition number (as defined by ϵ). The method is demonstrated on an example problem consisting of three stirred tank reactors in series. The approach only considers optimisation of the continuous design parameters in the process (*i.e.* reactor volumes and steady-state valve openings). Results are presented for different values of ϵ and show a difference in the steady-state economics of only five percent between the best and worse processes as judged by the dynamic operability objective. As stated earlier, controllability indicators may give useful insights, but do not relate directly and unambiguously to real performance requirements. It is not clear, therefore, how suitable values of ϵ may be meaningfully specified *a priori* in formulating the scalar optimisation problem to be solved. This point is perhaps best demonstrated by the authors themselves

in their summary of the results obtained: “...the process design that corresponds to $\epsilon = 13$ can be safely selected as the design which represents an excellent balance between the steady-state and dynamic objectives”.

Abbas et al. (1993) also consider that, because of the evolution of processes designed with tighter margins, design of the process and control system should be considered simultaneously rather than sequentially. Similarly to Palazoglu and Arkun, the requirements of obtaining a cheap process design which is controllable are posed in a multi-objective optimisation framework. In this case, however, the optimisation problem is solved using the theory of Multiple Criteria Decision Analysis which is essentially a systematic procedure for generating the noninferior set (discussed earlier) for the problem. The decision of what is the “best” solution from this set must be determined manually by the designer. The noninferior surface shows how changes in one criterion penalise the other criteria and may enable a number of possible solutions to be rejected. Ultimately, however, the decision is again a trade-off between economics and some controllability indicator. The authors demonstrate the method on a single stirred tank reactor. The choice of this unit is based on it being the most commonly used unit operation which exhibits stability and control problems. The multi-objective optimisation uses steady-state economics and the open loop damping ratio as objective functions. In the solution presented the secondary damping objective is actually included as an additional inequality constraint and the noninferior set generated by using a range of values of the secondary objective. The problem solved may therefore be considered similar to that solved by Palazoglu and Arkun, with a different objective function being used for the control requirements. There is also the additional limitation that only one stirred tank reactor is considered. Comments on the applicability of the method are as given for the work of Palazoglu and Arkun. Given the very restrictive size of the example presented, it is difficult to draw conclusions about the effectiveness of the integrated design algorithm presented.

The final work to be considered in this area is due to Luyben and Floudas (1994a). The requirements of greater energy and material integration, lower inventories and operation at multiple steady-states which increases interaction between control and design are again cited as the motivation for the research undertaken. The authors state that the aims of the synthesis approach presented are:

1. to incorporate open loop steady-state controllability measures simultaneously with the mathematical programming approach of process synthesis,
2. find a best compromise solution among the economic and control objectives by ap-

plying a multi-objective optimisation algorithm which considers quantitative trade-offs.

The process is represented as a superstructure which defines all possible alternative process structures. Also included in the superstructure are the controlled, manipulated and measured variables together with the transfer functions which relate them. The epsilon method used in the previous two papers is again used here. Recall, that only the economics are included in the objective, with the controllability criteria being added as inequality constraints in the problem formulation. Application of the epsilon method enables the design problem to be formulated as a mixed-integer nonlinear program (MINLP). Implementational difficulties prohibit analysis of the dynamic controllability of the system and therefore a number of steady-state controllability criteria (RGA, minimum singular value and steady-state condition number) are used in this work. Again the noninferior set is determined for the optimisation problem. A cutting plane algorithm is then used to determine the best-compromise solution with the trade-off information derived from the gradients of the noninferior set. In effect, the algorithm develops an implicit weighting between the multiple objectives. The validity of this weighting is unclear. The proposed algorithm is demonstrated first on a distillation column example for a fixed control structure and variable number of trays. A second example presents the design of a heat-integrated distillation column. In this case structural decisions are limited to the inclusion of a high pressure column, low pressure column or both. The control structure is again fixed. In both examples, feasibility of the designs was verified by simulating the closed loop response of the systems both to setpoint changes and disturbances. This work obviously represents a significant step forward in this area as it enables structural decisions to be included in the design and control problem. A key limitation, however, is the use of steady-state controllability indicators. Indeed Morari has noted in a recent review of combined process design and control (1994) that, “*any implications for the dynamic behaviour are tenuous at best*”. In a follow-up paper, Luyben and Floudas (1994b) consider the application of the design procedure described above to the design and control of a reactor-separator system with recycle.

2.3.3.4 Summary

The nature of the research discussed in this section may be summarised as follows. The processing plant is considered to nominally operate at steady-state, but must also possess dynamic characteristics which enable it to be operated successfully about this steady-state in the face of changing operating conditions (resilience). The common

approach taken in all this work is therefore to incorporate controllability indicators at the synthesis stage so that the resilience of the process may be assessed. In so doing, the designer is forced to trade-off an economic evaluation with some indicator which typically has no direct relationship to the “real” performance measures of interest. Such trade-offs can make the design problem difficult to define and lead to ambiguities. Although the tools described can give information about the dynamic characteristics of the plant, they are unable to guarantee that the dynamic process operation itself will be feasible in the presence of operational constraints. An illustrative example is useful here. Suppose that a process has been designed which is controllable as measured by any number of the available indicators and that a feedback control system is in place. The process is now required to make a change in operating point and is subject to one or more operational constraints. No controllability indicator can predict whether such a transition will be feasible, *i.e.* can the transition be made without violation of constraints. The only way to guarantee feasibility is to simulate the dynamic operation over time. This point is demonstrated in most of the case studies given above where the designs were ultimately verified using dynamic simulation.

A different approach to the problem is to consider the time-varying operation of the process explicitly at the design stage through the use of a dynamic process model. This will be the subject of the following review section.

2.3.4 Integrated process design based on dynamic operation

This review section considers the design of processes in which the dynamic operation forms an explicit part of the design problem. The work presented is arranged in essentially chronologically order.

Pioneering work in the area of process design with dynamic considerations was carried out by Nishida and Ichikawa (1975). They observe that process systems are inherently dynamic (*e.g.* startup, batch or semi-batch operation, changes from one set of operating conditions to another and perturbations which develop as process conditions fluctuate). They therefore determine that the effect of dynamic behaviour should be considered at the process synthesis stage rather than synthesising the system assuming steady-state operation and compensating the effect of dynamics by means of process control. The aim of their work is to develop a method of synthesising a process system where the dynamic operation is inevitable or preferable.

They describe the dynamics of the process system using a set of ordinary differential equations. The synthesis problem is then posed as an optimal control problem.

Optimisation of a set of time-varying piecewise continuous controls together with a set of design variables (*e.g.* vessel volumes) enables some objective function to be minimised. Typically the objective will be to minimise the total cost of the system. In addition to the set of design variables a number of “structural” optimisation parameters are included. These may be considered as continuous analogs to the structural integer variables used in mixed-integer programming approaches to process synthesis. The choice of values for these parameters establishes the flowsheet for the process. The authors develop a necessary condition for an optimal dynamic system using variational calculus. The resulting optimal control problem is solved using a gradient method which determines the optimal design and control parameter values for a restricted structure. The values of the structural parameters are then updated and the system is re-optimised. This iterative procedure terminates when the change in objective function values for consecutive iterations is sufficiently small.

The approach is demonstrated on the synthesis of two semi-batch reactor-separator networks, one with the irreversible reaction $A \rightarrow B$, the other with the reversible reaction $A \rightleftharpoons B$. The dynamics of the reactor are assumed to dominate such that an ideal separator model is used in which dynamics are neglected. A restricted “superstructure” is used with the reactor holdups as design variables and temperature profiles as the controls.

Nishida et al. (1976) extended this early work by considering the synthesis of dynamic process systems with uncertainty. The approach taken is similar to the previous paper. In this case, however, the design and operation is found which minimises the effect of the worst variations in the uncertain parameters chosen. The problem is posed as a minmax optimisation. The example considered is the semi-batch operation of a reactor-separator system with two exothermic reactions. Two useful observations stated in this work are:

1. The uncertain parameters have a significant influence over the dynamics of the process and a different structure will result compared to solution with the nominal values of the uncertain parameters.
2. Engineering experience and insight can be usefully applied to restrict the initial set of process structure alternatives in the optimisation.

Given the limitations of the computational power available at that time, the achievements of this work cannot be overstated. Indeed, the optimal design of dynamic systems under uncertainty will be the subject of the final paper in this review section. As will be seen, the problem Nishida et al. considered in 1976 still presents an immense

challenge some twenty years later.

Lenhoff and Morari (1982) considered the problem of process design under consideration of dynamic aspects. This work is founded on the observation that steady-state processes should be designed to be resilient. Because of this, the approach taken features a number of characteristics of the work reviewed earlier in section 2.3.3. It is, however, included here as the procedure employed requires integration and optimisation of a dynamic process model. The authors consider that the method of Nishida et al. described previously is too complex to be applicable to larger systems. Recall that this work considers process operations which are inherently dynamic. The problem considered by Lenhoff and Morari is that of process resiliency about a steady-state and for this purpose the method of Nishida may indeed be less appropriate.

In keeping with the approaches discussed in section 2.3.3 Lenhoff and Morari conclude that assessment of alternative designs should be considered on two bases:

1. steady-state aspects (economic performance index, EPI)
2. operational aspects (dynamic performance index, DPI)

The authors observe that these objectives are non-commensurable and therefore propose the use of a two valued objective function with the search for the optimal design being represented as a multiple objective problem:

$$\min_{k \in K} \{EPI(k), DPI(k)\} \quad (2.5)$$

where K is the set of all possible designs.

The separation of a binary mixture in two heat-integrated distillation columns is used as an example since the system has strong dynamic interactions. A procedure-based bounding approach with rigorous lower bounds is used to screen alternative designs. The lower bounds in the problem are determined through the use of dynamic optimisation. Decomposition of the process into subsystems allows the original complexity of this problem to be reduced. In the cases presented the DPI used is the weighted integral square error for the system as it moves from an initial deviation in the compositions back to steady-state. The authors present some 47 alternative designs with a range of flowsheet structures, control systems and process parameters and compare their respective EPI and DPI indices. The procedure proposed allows the number of attractive systems to be considerably reduced. Although the flowsheet structures and column designs are generated “manually”, the procedure is able to provide the engineer with insight into the effect of the steady-state design on the process dynamics. This is partly possible because the

design engineer can more directly relate the dynamic performance measure used (integral square error of the response) to actual performance requirements.

Howell (1984) also considered the effect of energy integration in distillation columns on the dynamics of the process. Specifically, Howell investigated the switchability of these processes, that is: *how readily a process switches between different operating states*. Howell's work considers the dynamic behaviour of systems which experience major disturbances in operating conditions. Given this objective, the author stresses that control and operating constraints, as well as the nonlinear characteristics of the system, must be considered explicitly. The problem of optimal switching is therefore posed as a dynamic optimisation where the index used to measure the switchability is the extra cost in moving from one steady-state to another. The use of this objective gives an economically meaningful measure of the process dynamics. This work, however, considers the process controls as the only design variables. Extensions to this work which consider the effect of the process design on switchability have been carried out by White (1994).

White presents two case studies in which design aspects are included in the switchability analysis. Again, dynamic optimisation is the tool used for this purpose. Results show that design changes can improve switchability as measured by an integral square error (ISE) objective. White observed that columns operating under higher reflux ratio, but with fewer stages were more "switchable". Reducing tray weir heights reduced the ISE switchability objective. The reactor-separator case study originally presented by Nishida and Ichikawa (1975) was also revisited. White's analysis incorporates a more rigorous dynamic separator model which enables optimisation of the number of stages through the use of a lumped column model. The operating cost during a switch in product grade is used as the objective function.

Narraway's (1992) work on control structure selection based on economics was reviewed earlier in section 2.3.1.2.2. Recall that the formulation presented determines the economically optimal selection of control structures based on perfect controllers. The original MILP form of the problem was extended to an MINLP formulation capable of selecting the optimal multi-loop PI control structure. The method explicitly considers the dynamic response of the process to a single disturbance over a finite time horizon. Solution of a dynamic optimisation subproblem enables optimal, feasible close-loop operation of the process in the face of this disturbance to be guaranteed. Since the focus of this work was control structure selection techniques, reference to issues of process design are limited.

Walsh (1993) considers the integrated design of chemical wastewater treatment systems. A number of reasons are cited for the wastewater treatment problem being a

suitable study for integrated design: the system is highly nonlinear with complex reaction kinetics; performance constraints are demanding; there is significant uncertainty in the process parameters; process disturbances can be very harsh. The stated aim of the design procedure is to:

maximise return on investment while ensuring that performance constraints are satisfied for all possible plant parameters and disturbances.

In other words, can a particular design or structure be made to satisfy its operating constraints despite uncertainty and disturbances *and* what is the impact of this on the economics of the processing system considered? Walsh observes that these questions do not fit comfortably within any of the conventional categories of operability analysis, such as controllability, switchability, flexibility and robust control (ability to maintain stability and performance of a control system despite perturbations in the characteristics on the controlled system). The approach taken is therefore to consider an integrated approach where operability issues are on a par with economic issues. This allows design and synthesis decisions to be made within a common framework. An important implication of this approach is therefore that:

performance requirements can be represented either by constraints or by a contribution to the economic objective.

This is in contrast to a number of works discussed previously where the use of multiple objectives was deemed necessary because of the difficulty in “measuring dynamic resilience with a scalar objective function bearing economic significance”.

The solution approach taken in Walsh’s work is hierarchical in nature and acknowledges both the value of contributions from the design engineer and from optimisation techniques. The procedure aims to narrow the search domain by efficient elimination of large numbers of inviable options and deepen the search by picking good candidate processes for more detailed analysis. Efficient screening of process structures is achieved through the use of heuristics and bounding properties. As viable candidate processes are found, more sophisticated optimisation techniques using refined models are employed for full performance evaluation of the design. At the highest level, a worst-case dynamic optimisation of the process ensures feasible dynamic operation in the face of disturbances and uncertainty.

Dimitriadis and Pistikopoulos (1995) acknowledge that chemical plants are complex dynamic systems that must operate successfully throughout time and that methods for assessing resilience are limited to a linearisation about some steady-state operating

point. They note that uncertainties are generally dynamic (e.g. quality and quantities of feeds, product demands *etc.*) and proceed to define a dynamic flexibility index. The dynamic flexibility index is a dynamic analogy to the steady-state flexibility index proposed by Swaney and Grossmann (1985). Dimitriadis and Pistikopoulos make the point that when the variation of uncertain parameters over time is known, the flexibility problem transforms into a standard dynamic optimisation problem that can be solved by established techniques. Operation of a batch reactor and a heat exchanger network are presented as case studies. In the examples given, the dynamic flexibility index is found for a known time variation in the “uncertain parameter”. Design aspects are considered, but limited to the retrofit of the exchanger network to improve the flexibility index. More interestingly, an extension to the case when the time-varying uncertainty is unknown is proposed. The method uses orthogonal collocation on finite elements of the states and controls and determines the maximum deviation from feasible operation under the best achievable parameterised control action. The calculation of the flexibility index for a system consisting of a single tank with one feed and one product is given as an example. A discretisation of ten finite elements, each with two collocation points results in a problem with 333 continuous variables, 143 binary variables and 423 constraints. Applicability to larger problems must, at present, therefore be considered limited.

A capacity-based economic approach to assess process controllability during the design stage has recently been published by Elliot and Luyben (1995). The method uses linear models and closed loop frequency response techniques to determine worst case disturbance conditions. A sinusoidal disturbance at the peak condition is then applied to the process model. Using the linear time response of the process and bounds on the desired specification range, the amount of off-spec material generated is determined. Given that a cost is associated with re-processing the off-spec, a calculation of the annualised profit may then be made. The authors conclude that the approach: *has the potential to be a viable screening tool for a rapid assessment of both steady-state economics and process controllability*. However, the method presented is flawed on at least two counts:

- The procedure requires the process design, control structure and controller tuning to be completed before any analysis can take place. Given that this point must be reached in the design process, the method does not represent a particularly valuable screening tool for conceptual designs.
- The method is applied to a process *and* its tuned control structure and therefore does not give an analysis of the controllability of the *process* as suggested.

Recently, the design of rapid-transition continuous multiproduct processes has been discussed in the literature. In such processes, the time required to switch between operating points can impact significantly on the overall cost by changing cycle times and inventory profiles. Cantrell and Luyben (1996) proposed a design algorithm which considers the dynamics of transitions in the design of such processes. The method used elements of cycle time modelling from operational research, process control theory and heuristic process design to develop a multiproduct plant and cyclic schedule which more rigorously incorporates the dynamics of transitions. Since the process is inherently dynamic, feasibility of the process operation over time is examined through the use of dynamic simulation within the design algorithm. This work will be considered in more detail in chapter 4.

The last work to be reviewed in this chapter was recently published in the literature by Mohideen et al. (1996). The paper considers the optimal design of dynamic systems under uncertainty. A unified process design framework is described which facilitates the integration of process and control system design. The processes developed in this framework are economically optimal and can cope both with parametric uncertainty and process disturbances. The problem description is formed from:

- a dynamic mathematical model describing the process which includes all operational constraints,
- a model describing the uncertain parameters and time-varying disturbances,
- a set of process design and control alternatives.

Based on this, the design problem is posed as a mixed-integer optimal control formulation. Solution of the problem is achieved through iteration of a two stage decomposition algorithm:

1. solution of a multiperiod design subproblem to determine the process structure and a suitable control structure to satisfy a set of critical parameters over time,
2. a time-varying feasibility analysis step, which identifies a new set of critical parameters for a fixed design and control structure.

The procedure is demonstrated on a ternary distillation column which is subject to uncertainty in the feed rate and composition and in the cooling water temperature. A superstructure tray-by-tray model is used for the column, thus allowing the number of stages to be included in the optimisation. Solution using the proposed decomposition algorithm gives a 15% improvement in capital costs over a column designed in a more traditional sequential manner, but operating costs increased by 7%.

The approach presented provides the promising basis of a framework for integrated process and control structure design under consideration of operational uncertainties. Clearly, as improvements in computational power continue, approaches such as this may become more practical. Whether the use of single pass “black-box” optimisation as the only tool for design, where the design engineer is totally excluded, is appropriate is another question. As Walsh has observed, “the challenge to integrated design is to find a good division of labour between human and computer and an appropriate way of interfacing the two”. This question will be addressed in further detail in chapter 3 when the design procedure to be used in this work is discussed.

2.3.4.1 Summary

This review section has considered a class of integrated design problems in which the time-varying operation of the process is considered explicitly. An advantage of this approach over design with controllability considerations, is that the operating envelope of the process is not restricted to deviations about a nominal steady-state. As a result, more extreme changes in operating mode may be considered - for example, startup operations and batch production. This is clear from the example problems studied, which ranged from startup of semi-batch processes to continuous wastewater treatment processes where consideration of dynamic effects is of crucial importance. The inherently dynamic nature of a distributed plant was highlighted earlier in section 2.2.5 and an assessment of the process operability about steady-state (using some controllability measure) would therefore seem to be inadequate for use in this work. It is therefore evident that any integrated design approach which is developed here should enable explicit consideration of the process operation at the design stage.

2.3.5 Conclusions

This chapter has introduced the concept of a distributed chemical manufacturing facility. The design requirements of such a plant were discussed, first generally and then in the specific context of the process design problem. The inherently dynamic nature of the distributed plant was highlighted by consideration of its operational requirements. It was concluded that an integrated design approach was appropriate in this work. A review of the integrated design literature was then presented. Based on this review and on the expected operational nature of the distributed plant, it was concluded that the design procedure used in this work should enable explicit consideration of the process operation at the design stage. This design procedure will be discussed in the following chapter.



Chapter 3

Design procedures

3.1 Introduction

In the previous chapter, the concept of a distributed chemical manufacturing facility (DCMF) was introduced and the functional requirements were outlined. The realisation of a DCMF in practice was shown to require the combination of technologies from a wide range of scientific and engineering disciplines. Design of DCMFs in the context of this work was restricted only to the area of process systems engineering. As stated previously this may be thought of as: finding the “best” type, topology and operation of units which, when combined, give the desired transformation of feeds to products. It was further argued that because the distributed plant will be inherently dynamic and that feasible operation must be ensured at all times, an integrated approach which simultaneously considers the process operation is highly appropriate to the design of such plants.

This chapter describes the development of an integrated design approach which will be used in tackling the case studies to be presented later in this thesis. Starting from a general description of the problem which is to be addressed in section 3.2, two solution methods are explored; one is a wholly optimisation based approach, the other a hierarchical procedure. A hierarchical approach is preferred for this work. In the procedure presented in section 3.2.2, the design problem is made more tractable by decomposition. A two-stage procedure is proposed in which a viable structure is first established. The corresponding operation and design for this candidate structure are then determined. The concept of cyclic operation is introduced in section 3.3 which facilitates incorporation of the process operation into the design problem.

Tools for determining the optimal design and operation of the candidate pro-

cess are considered in section 3.4. Of the methods presented, only dynamic optimisation enables both the design and operation of the process to be simultaneously considered. Section 3.5.2 gives a summary of methods for the solution of dynamic optimisation problems. A feasible path control parameterisation technique is selected for use in this work. Implementation details of the specific code used are then given in section 3.5.3.

Finally, section 3.6 gives an introductory overview of the case studies to be presented in the remainder of this thesis.

3.2 Design procedure

Recall that the aim of the integrated design problem, as stated by Walsh is, *to obtain minimum cost designs while ensuring that safety and performance constraints are satisfied for all possible plant parameters and disturbances*. This is generally achieved by determining the economically optimal combination of:

1. unit type and topology,
2. unit design,
3. process operation.

Two possible approaches to the solution of such a problem are considered in the section which follows.

3.2.1 Mixed-integer non-linear programming approach

In determining the optimal design and operation of a distributed plant, three distinct sets of decision variables will generally be involved. The choice of these variables will determine the ultimate design and operation of the process:

1. discrete decisions (*i.e.* does a unit or controller appear in the final flowsheet) which determine the process structure,
2. continuous, time-invariant parameters (*e.g.* controller tuning, vessel sizes) which determine the unit and controller design,
3. continuous, time-varying control profiles which may determine the process operation.

A mixed-integer non-linear program (MINLP) enables all of these decision variables to be simultaneously manipulated to optimise some objective function (usually economic)

subject to the constraints imposed on the problem. In other words, mixed-integer non-linear optimisation enables the process structure, design and operation to be determined at one level. However, for this work such an approach suffers a number of drawbacks and was rejected for the following reasons:

1. In an MINLP approach a complete *a priori* specification of the design problem (constraints, objectives *etc.*) is required before solution may begin.
2. Process design is typically an on-going and highly interactive activity in which the effects of external influences may require revision and refinement of the problem definition. In an entirely optimisation based approach, any changes in problem definition during the design process would require the MINLP to be reposed and resolved.
3. Generation of an appropriate process superstructure which covers all possible design options can be both difficult and time consuming, especially where the system may include numerous units of different types.
4. The design engineer can apply judgement and experience in order to decide which set of design decisions to explore next. An MINLP algorithm typically uses a linearisation about the NLP solution for certain sets of the discrete decision variables to capture the results for that structure. The design engineer has the potential to extract much more information from the current design than this.
5. Solution of MINLPs is very numerically intensive, particularly where solution of the NLP is under-pinned by a numerical integration of the dynamic process model. This will be the case here, since consideration of the process dynamics is a key element of this work.
6. Based on the experience of other workers (Narraway, 1992), the available solution technology for MINLPs was not considered satisfactorily robust to form the basis of a procedure for the design of distributed plants. Recent work by Mohideen et al (1996) in the area of optimal dynamic design (discussed earlier in section 2.3) is encouraging and offers some hope that wholly optimisation based approaches to the process design problem may be a more realistic prospect in the future.

Having excluded a design procedure based entirely on optimisation, the following section considers an alternative decomposition strategy which was ultimately adopted for this work.

3.2.2 Decomposed design procedure

The design procedure used in this work does not attempt to simultaneously solve the combined structural and operational problem described above. The alternative approach taken here is to decompose the problem into two stages which separates determination of the process structure from the process unit design and operation. The two stages of the proposed procedure are as follows:

1. develop a candidate process structure,
2. determine the corresponding (optimal) operation and design for that candidate.

The underlying philosophy of the design procedure may be described as a “design funnel” and is shown figure 3.1. In applying the procedure, a tradeoff is sought between the cost




Figure 3.1: “Design funnel” to be used in process design procedure.

of using a given design tool and the level of input it requires, and the discriminating power of the tool. At the higher levels where many design options exist, discrimination between structures may be entirely heuristic. For example, at the top level nominal unit operations to be used in the process may be specified based upon past engineering experience. The objective of the preliminary design stage, which corresponds to the upper levels in the design funnel, is to produce a viable candidate design for more detailed analysis. It is anticipated that the required level of data for this level may be significantly less than is needed for the complete design. Where the candidate structure is feasible, it provides an upper bound on the achievable economic performance of the process. Only when a viable candidate is in place are more computationally intensive design tools employed to determine the optimal design and operation of the process. Solution of this second

stage in the procedure is a much more tractable problem than the MINLP formulation discussed previously. Indeed, robust and well-tested numerical solution methods do exist and will be discussed in section 3.4. Having outlined the overall approach to be taken in the design procedure, the stage one screening of alternative process structures will now be discussed.

3.2.2.1 Developing a candidate structure

In this work the application of a heuristic-based search will be used which enables process structures to be efficiently eliminated. The search applies a bounding strategy to provide information on the achievable performance of a given design. Using such an approach in this work would seem sensible given there is limited experience in the design and operation of distributed plants. Note that it is these same limits which prevent design objectives and options being set out clearly in a form amenable to MINLP solution. Predictor-based bounding has been applied successfully by a number of authors as a design tool. Mizsey and Fonyo (1990b; 1990a) use a combination of heuristics and rigorous bounding properties to rapidly synthesise energy integrated total flowsheets. Walsh (1993) used a similar predictor-based bounding approach in his work on the integrated design of chemical wastewater treatment processes. The tools developed by Walsh enabled dynamic as well as steady-state screening of process alternatives. An integral part of the case study work presented in this thesis will therefore be to develop suitable bounding properties and heuristics to facilitate the efficient initial screening of design alternatives at the preliminary level. Although these properties are likely to be case specific, it is intended to show through the cases presented that they are often straightforward to establish and have the potential to rapidly reduce the search space of viable alternatives.

The determination of candidate structures will be covered in further detail later in this thesis when the solution of specific case study examples is presented.

3.3 Description of the process operation

In the previous section a procedure for the design of distributed plants was proposed. The procedure aims to make the design problem more tractable by decomposition into two stages. The first stage uses the hierarchical application of increasingly sophisticated design tools to develop a viable candidate process structure. Once this structure is in place, the second stage of the procedure is to determine the optimal operation and design of the fixed structure. This section re-considers the nature of the distributed plant

operation, first discussed in section 2.2.5, with a view to incorporating this operation into the second stage of the design procedure.

3.3.1 Operating requirements of a distributed plant

The potential operating requirements of a distributed plant discussed in section 2.2.5 are briefly reviewed here. Recall that the distributed plant must operate in a totally automatic fashion. The design must be capable of exactly tracking customer demands on both the quality and quantity of product (within the limits of the plant design) over a time horizon defined by the design life of the plant itself. The plant must similarly be capable of continued feasible operation in the face of variability in the feeds entering the plant. An example of this would be a treatment plant where concentration and flowrate loads to the process are variable and limits on the discharged effluent must be met. Again, these variations occur over an indefinite time horizon. The potential variability affecting the process is summarised in figure 3.2.




Figure 3.2: Sources of operating variability affecting distributed plant.

The generality of the uncertainty described above presents significant difficulties in determining the design and operation of the process. Not only must the design be feasible for all operating points within the specified product space (demand and quality) and the defined variability limits on the feedstocks (flowrate and quality), but it must be capable of *feasibly moving between* any two arbitrary operating points contained within this space. Clearly, a more restricted description of the variability in the feed and product profiles is needed in this work if the process operation is to be successfully incorporated into the second stage of the design procedure. The key property required from a simplified process operation description is that it should enable the process behaviour over an “infinite time” horizon to be captured by a finite horizon optimisation problem which incorporates the important dynamic effects impacting upon the design.

3.3.2 Cyclic operation of processes

3.3.2.1 Motivation for cyclic operation

This work will make the simplifying requirement that the process operates in a cyclical manner. The use of a cycle provides a way of collapsing the infinite operating horizon into a finite, but infinitely repeatable operation. Examples of the use of cyclic operations in industrial processes can be cited. Multiproduct plants, for example, often use a cyclic production schedule to give easily implementable, repeatable operating strategies (Sahinidis and Grossmann, 1991). It is intended that cyclic operation will give similar benefits in terms of incorporation into the second stage of the design procedure. The example of a cyclically operating multiproduct process will be used as an introductory case study in this work. Although this process operates under conditions of constant feed and product specifications (*i.e.* no external disturbances), the production schedule requires frequent changes in operating point which in turn requires that the process dynamics are considered at the design stage. Cyclic operation of batch distillation columns, in which three phases of operation (filling-up, total reflux and dumping) are continuously repeated, has been discussed by Sørensen and Skogestad (1994). The authors show that such operation is beneficial where difficult separations must be carried out. The nature of the cycles to be implemented in this work will be discussed further in section 3.3.2.3 below and in case-specific detail in the relevant case study chapters.

3.3.2.2 Disturbance description

As stated in section 3.3.1 above, the distributed plant will be subjected to a number of external disturbances in the feed and/or product specifications. The nature of these disturbances partially determines the resultant dynamic behaviour of the system. A further simplification in this work is that only one external disturbance is considered at a time. With regard to this, the cases presented later in this thesis individually consider:

1. changes in the feed concentration load,
2. changes in the product demand rate.

Recall that the process operation which is implemented is required to capture the key dynamics which affect the process design. For simplicity, the external disturbances considered in this work are therefore limited to step changes between two specified levels. The step which is used is chosen to represent a perceived worst-case. Note that in the multiproduct plant case study presented first, the sequencing of the production schedule

(which may be considered as a form of “internal” process disturbance) is actually a design parameter. This contrasts with the second and third case studies where the disturbance profile is imposed on the system.

3.3.2.3 Form of the operating cycle

It has already been stated that the process operation will be required to be cyclic, the motivation for this being that such a cycle is repeatable *ad infinitum*. The general form of this cycle is now explained in more detail.

At time zero, the process exists in some state which is not specified *a priori*, but is determined as an integral part of the solution to the design problem. The specified sequence of disturbances is then applied to the process, which must operate feasibly throughout. A cycle is formed by requiring that the process returns to its initial state by the end of the cycle, when the sequence of disturbances is completed. It is important to note that this condition does not in itself guarantee that the cycle will be repeatable. Operation in a cyclical manner only guarantees repeatability if the cycle is repeated without a delay. This is the case with the multiproduct plant example presented as the first case study in this thesis. A further condition must be imposed on the operating cycle in the case where there is a delay between the end of one cycle and the start of the next. The condition applied is that the process must reach a steady state at the final time in the cycle. Given that the state of the process at the final time in the cycle exactly matches that at the initial time *and* that the final states are steady, the process can be held in this state indefinitely until the sequence of disturbances is again applied and the cycle repeated.

As will be seen later in section 6.8.4, a more general specification for cyclic operation may be formulated which defines a possible set of process initial conditions rather than a single operating point. The process is then required to return to a state contained within this possible initial set. The formulation enables the exact steady-state endpoints required for optimal open-loop cyclic operation to be relaxed. Though less general, the original form of the cyclic operation posed, in which the initial and end values of the states are defined by a single point, is still applicable to a wide range of problems. This will be seen in the case study chapters which follow.

This discussion has provided a generalisation of the form of the process operating cycle. The specific requirements needed to accurately capture a given problem definition are necessarily case dependent. These requirements will therefore be described in detail in the respective case study chapters.

3.3.3 Re-statement of the design procedure

In the light of the discussion above, the design procedure presented in section 3.2.2 is re-stated as follows:

1. develop a viable candidate process structure,
2. determine the corresponding (optimal) cyclic operation and design for that candidate.

Clearly, the design procedure stated above does not set out to find the optimal design and process operation, but rather to determine efficient designs and corresponding feasible operating strategies.

The hierarchical design procedure to be used in this work has been discussed previously. Recall from section 3.2.2 that the philosophy of the approach is to efficiently screen design alternatives such that more intensive numerical methods are applied only when a viable candidate process structure has been established. The concept of cyclic operation was then introduced to enable the process operation to be more easily incorporated into the second stage design problem. By introducing the above simplifications and problem decomposition, the overall design problem has been reduced from a single pass optimisation to a number of smaller, more tractable, sub-problems. The biggest of these sub-problems is the simultaneous design and cyclic operation of the candidate process. A robust tool for undertaking this stage of the design procedure is therefore required for this work to be successful. Potential alternatives for this purpose will be considered in the next section.

3.4 Tools for integrated process design and operation

As stated above, a tool for developing cyclic operating policies within an integrated design framework is required for this work. Section 2.3.4 reviewed research which has considered the problem of integrating the process operation and design, as is the aim here. In all the work discussed, dynamic optimisation was used exclusively as the tool to achieve this. Given that the aim here is to optimise a set of process design parameters in relation to performance criteria based on a nonlinear dynamic model, dynamic optimisation would appear appropriate for this work. Before adopting dynamic optimisation as the second stage design tool for this work, a review of alternative methods for determining the process operation was carried out.

3.4.1 Operating Procedure Synthesis

The synthesis of operating procedures is described by Fusillo and Powers (1987) as, *finding a set of actions which would bring the [chemical processing] system from some initial state to a goal state, subject to constraints dictated by consideration of process chemistry, process equipment, safety and environmental consequences*. Over the past twenty years a number of authors have developed formal methods and tools to solve the problem of operating procedure synthesis for chemical processes. These tools fall into three main categories:

- valve sequencing operations,
- state transition sequences, and
- artificial intelligence techniques.

All of the formal methods investigated were found to apply the operating procedure synthesis to a fixed process design, with the methods themselves being largely based upon steady-state or quasi steady-state models. Since these methods allow no consideration of the process design, they cannot be considered appropriate for use in this work. A full review of these methods may be found in appendix A. The use of dynamic optimisation as an appropriate tool for use in this work is considered in further detail in the sections which follow.

3.5 Dynamic optimisation as a tool for process design and operation

The dynamic optimisation problem as stated by Kirk (1970) is:

determine the control signals that will cause a process to satisfy the physical constraints and at the same time minimise (or maximise) some performance criterion.

Dynamic optimisation techniques have been widely applied to a range of problems where the time varying behaviour of the process is of importance. For example, in the case of batch reactors, it has provided a means of finding operating policies which maximise the production of valuable intermediate reaction products (Vassiliadis *et al.*, 1994; Palanki *et al.*, 1994). Similarly, dynamic optimisation strategies have been implemented in bio-processes in order to maximise protein production (Lee and Ramirez, 1994). Also in the biochemical engineering area, Van Impe *et al* (1994) found dynamic optimisation to be

a useful tool for model discrimination purposes. Of particular relevance to this work, a number of authors have used dynamic optimisation to simultaneously optimise both process design variables and the time-varying process operation. For example, White (1994) included the design of a distillation column as a degree of freedom when determining the optimal switching trajectories of the column. Dynamic optimisation offers a method by which both the process design and operation may be simultaneously considered. The application of such a technique therefore seems appropriate in this work. The general form of the dynamic optimisation problem which will be solved in the case studies is presented below in section 3.5. The specific form of the optimisation problem to be solved is case dependent. This will be considered in detail when determination of the optimal design and operation of the each of the three case studies is introduced.

It should be stressed at this point that dynamic optimisation is not a purely theoretical tool and the resulting strategies can be implemented in real time controllers. Leversund et al (1994) maximise the profitability of a reactive batch distillation system and design a PI controller to implement the optimal profiles generated. Earlier work by Sargent and Sullivan (1979) led to the implementation of sub-optimal control laws for refinery unit feed changeovers which captured the form of previously determined optimal control trajectories.

In the sections which follow the general mathematical form of the dynamic optimisation problem will be stated and its strengths in the context of this work considered. Following this, solution methods will be discussed along with the tools available for solution of such problems.

3.5.1 Form of the dynamic optimisation problem

The general dynamic optimisation problem to be solved can be written mathematically as,

$$\begin{aligned}
 & \min_{u(t), d} J(x(t_f), y(t_f), u(t), d) \\
 & s.t. \quad f(\dot{x}(t), x(t), y(t), u(t), d, t) = 0 \\
 & \quad \quad q(x(t), y(t), u(t), d, t) \leq 0 \\
 & c_1(\dot{x}(\phi_1), x(\phi_1), y(\phi_1), u(\phi_1), d, \phi_1) = 0 \quad \forall \phi_1 \in \Phi \\
 & c_2(\dot{x}(\phi_2), x(\phi_2), y(\phi_2), u(\phi_2), d, \phi_2) \leq 0 \quad \forall \phi_2 \in \Phi \\
 & \quad \quad 0 \leq t \leq t_f \\
 & \quad \quad \Phi \subset t
 \end{aligned}$$

where, $x(t)$ is the vector of state variables and $\dot{x}(t)$ is the vector of state time derivatives, $y(t)$ is the vector of algebraic variables, $u(t)$ is the set of time varying optimi-

sation variables (controls) and d are design variables. The dynamic optimisation problem is therefore to find a control vector ($u^*(t)$) and design variables (d) which cause the system ($f(t)$) to follow a trajectory ($x^*(t)$) that does not violate any of the constraints (q , c_1 or c_2) and optimises the performance measure (J). $u^*(t)$ is called an optimal control and $x^*(t)$ is an optimal trajectory. The relevance of the constraints is as follows,

1. f is a set of differential algebraic equations (DAEs) which form a model of the system,
2. q is the set of path constraints which must be satisfied for all times from 0 to t_f ,
3. c_1 and c_2 are equality and inequality constraints which must be satisfied at a discrete set of times (point constraints).

In the case of the optimal cyclic operation of a distributed plant, q might represent operational constraints such as maximum allowable temperatures, or purity requirements. c_1 and c_2 will include constraints such as those which ensure that the operation is cyclic, and/or steady at the end of a cycle. The objective function, J , will typically be an economic measure. Note that in the cyclic operation problem the set of design parameters d is taken to include the initial state of the process at the beginning of a cycle.

3.5.2 Solution Methods for Dynamic Optimisation Problems

Methods for the solution of dynamic optimisation problems may be placed in one of four groups described in the literature:

1. dynamic programming (flooding),
2. iterative solution of the two-point boundary value problem,
3. mathematical programming approaches, which can be subdivided further as,
 - simultaneous or infeasible techniques,
 - sequential or feasible techniques.

The background to each of these techniques will now be briefly discussed.

3.5.2.1 Dynamic programming and iterative solution techniques

Formulation of the dynamic optimisation problem using the calculus of variations will, in general, lead to a non-linear two-point boundary problem. Although it is a relatively straightforward exercise to solve an initial-value problem, extension to solving

two-point boundary problems is not trivial and *must necessarily involve either flooding or iterative procedures* (Bryson and Ho, 1975). The problem is to find the n state and adjoint variables and m control variables which simultaneously satisfy the n system and adjoint differential equations, the m optimality conditions and initial and final boundary conditions.

One method of solving this problem is that of dynamic programming. The idea here is to systematically choose values for the unspecified initial or final conditions and compute the corresponding optimal trajectory. The remaining unspecified boundary conditions are used as parameters. If the range of specified boundary conditions at one end are chosen correctly an optimal trajectory passing through the desired boundary conditions at the other end will be found. If the optimal solution does not pass exactly through the boundary conditions the solution may be interpolated using acquired knowledge of the state space. Storage requirements for this technique are large, even for small problems, and quickly become prohibitive.

An alternative method for the solution of the two-point boundary value problem is to use an iterative procedure. Bryson and Ho describe three such procedures which have been widely used:

1. neighbouring extremal methods,
2. gradient methods,
3. quasilinearisation methods.

All of the above methods aim to improve a nominal solution¹ by successive linearisation until the true solution is obtained. Each iterative procedure differs in the nominal solution it uses and hence the means by which the solution is obtained. The nominal solutions used are summarised in table 3.1.

The method of neighbouring extremals uses an iteration to improve the estimate of the unspecified boundary conditions so as to satisfy the specified boundary conditions. The weakness of this approach stems from the sensitivity of the optimal trajectory to the value of the unspecified boundary conditions. An initial guess for the unspecified boundary conditions which gives a solution close to the specified conditions at the other end is needed and often difficult to provide.

In order to overcome the problem of providing an initial guess for the neighbouring extremal method first-order gradient algorithms have been developed. These

¹a solution which satisfies none, one, two, or three of either the system equations, the adjoint equations, optimality conditions or boundary conditions.

Iterative procedure	Nominal solution satisfies			
	System equations	Adjoint equations	Optimality conditions	Boundary conditions
Neighbouring extremal method	yes	yes	yes	no
Gradient method	yes	yes	no	no
Quasilinearisation method	no	no	yes	yes or no

Table 3.1: Nominal solution for different iterative procedures

algorithms attempt to improve the estimates of the control trajectories in order to satisfy the optimality and boundary conditions. This method gives excellent initial convergence, but poor convergence as the solution is approached. The poor convergence characteristics of the first-order gradient method are tackled by introducing a second-order gradient method. However, significant restrictions apply as to when these second-order methods can be used (Bryson and Ho, 1975).

The final iterative approach to be considered is that of quasilinearisation. Here, the control and state trajectories for the nominal solution are chosen such that the optimality conditions and as many of the boundary conditions as possible are satisfied. The system and adjoint equations are then linearised about the nominal solution. The solution is modified by solving a series of linear two-point boundary problems until the system and adjoint equations are satisfied to the required accuracy. Providing an initial guess for the state trajectory is fairly straightforward and convergence is rapid near the solution.

In their simplest form the algorithms described above do not allow inequality constraints on either state or control variables. Constraints of this type may be approximated by augmenting the objective function with an integral penalty function.

The discussion above is intended to give a summary of iterative methods for solution of dynamic optimisation problems. For specific details of algorithms the reader is referred to chapter 7 of Bryson and Ho (1975).

3.5.2.2 Mathematical programming techniques

Solution techniques in this area fall into two categories. A summary of each of these is presented below. A more comprehensive review may be found in Vassiliadis (1993).

The common feature of both categories is that they reduce the infinite dimen-

sional dynamic optimisation problem to a finite dimensional problem which may be solved using a straightforward nonlinear program (NLP). The infinite dimensional nature of the problem arises from the infinite number of points which lie on the time horizon for the optimisation. The specifics of each approach will now be considered in turn.

3.5.2.2.1 Simultaneous or infeasible methods In this method the control profiles and DAE system are approximated with a suitable polynomial and rewritten as a set of algebraic equations. During solution, the DAE system and optimisation are converged together, hence the term simultaneous. Until the converged solution is reached the solutions at intermediate iterations will generally be meaningless, thus the term infeasible is sometimes used. Much of the work in this area is due to Biegler et al. at Carnegie Mellon University (Biegler, 1988; Vasantharajan and Biegler, 1990; Logsdon and Biegler, 1993). The main advantage of this approach is the potential to lower computational cost by reducing the objective function and constraint violation together. A second argument for this method is that the optimiser can “see” the constraint violations at all discretisation points along the path of the optimisation which is not always the case with feasible path methods.

3.5.2.2.2 Sequential or feasible methods In this method it is only the control parameters which are discretised. The control variables are approximated by a piecewise polynomial function. Starting from an initial guess for the control profiles, integration of the system state and sensitivity equations is carried out. An outer optimisation loop then computes the required update for the control profiles. In this way a two level sequential optimisation is carried out with intermediate solutions of the DAE system always being feasible. Several authors have presented algorithms for the solution of dynamic optimisation problems using this technique dating back as far as 1970 (Pollard and Sargent, 1970).

The choice of a suitable technique to use for this work will now be addressed. Current research has focused mainly on the two mathematical programming approaches described above. Walsh has discussed the “best” technique to use based on four factors. He concluded that, at that time, a feasible path method was the preferred option. These factors are summarised in table 3.2.

Given that:

1. the factors favouring feasible path methods still largely apply,

	Feasible path	Infeasible path
Ease of solution	Requires integration of DAEs followed by solution of moderate sized NLPs. Well established software for both.	Requires use of specialised optimisation codes and new error control strategies - difficulties with large problems.
Efficiency of solution	comparison is inconclusive.	comparison is inconclusive.
Generality of method	Hybrid approach uses path constraint converted to terminal constraint and interior point constraints.	Path constraints translated to point constraints at the discretisation point.
Usefulness of intermediate results	In the event of solution failure it will give a meaningful intermediate solution	In event of failure "solution" may not satisfy DAEs.

Table 3.2: Feasible vs. infeasible path methods

2. the feasible path control parameterisation method has been successfully applied to numerous problems of the type which this work is likely to address,

a tool based on a feasible path method is preferred for this work.

3.5.3 A feasible path dynamic optimiser

The general form of the dynamic optimisation problem to be solved was presented earlier in section 3.5. The method by which this problem will be solved in this work was outlined above in section 3.5.2.2.2. Recall that a feasible path control parameterisation approach will be used. The specific implementation which is used here was developed by Walsh in his work on the "Integrated Design of Chemical Wastewater Treatment Systems". A brief overview of the implementation is presented below. More specific details of the code may be found in Walsh (1993).

3.5.3.1 Time discretisation and control parameterisation

Recall that the time varying optimisation parameters in dynamic optimisation problems are discretised in order to create a finite non-linear programming problem. The time horizon of the problem is split into a finite number of elements. In each time element the control variables are defined as low-order polynomial functions. The desired control profile is achieved in the optimisation by manipulating both the time interval lengths and the parameters which describe the form of the polynomial. Two key factors influence the order of polynomial which is used:

1. Piecewise linear controls give a better approximation to smoothly varying controls than piecewise constant controls and can also approximate sharp discontinuities.
2. Both piecewise linear and piecewise constant controls have the following property. *If a control parameter is interpolated between two breakpoints² and its values at the two breakpoints are within the bounds on that control, the interpolated values of the control are guaranteed also to lie within those bounds.* This property cannot be guaranteed for higher order polynomial functions.

Continuous piecewise linear controls are therefore used throughout this work to give the highest precision while still ensuring the control variables remain correctly bounded using just simple bounds on the break points.

The issue of choosing a suitable number of discretisation intervals in the solution of dynamic optimisation problems by control parameterisation has received little attention in the literature. White (1994) states that the number of intervals to use is arbitrary and suggests a trial and error approach. In general it is considered that the application of successive quadratic programming algorithms to problems containing more than two hundred degrees of freedom is likely to result in very poor performance of the algorithm. With this in mind the number of intervals used will be chosen to keep the total number of optimisation parameters below this figure.

3.5.3.2 Solution method

The solution of the discretised problem obtained above follows a two stage iterative procedure:

1. Integration of the process system equations for fixed values of the optimisation parameters to obtain constraint residual information and/or gradients,
2. Solution of a non-linear programming problem to determine new values for the optimisation parameters which improve feasibility and/or the objective function.

The dynamic optimisation code implemented by Walsh uses DASOLV (Jarvis and Pantelides, 1992) to integrate the DAE system and calculate the sensitivity values, which may be done in parallel with the forward integration. The non-linear program solver SRQPD (Chen and Macchietto, 1991) is used for solution of the finite optimisation problem. It should be noted that other methods for obtaining gradient information include finite differences or backward integration of the adjoint system. However, both Walsh

²a breakpoint occurs when one time interval ends and the next begins

and Vassiliadis (1993) argue that the sensitivity approach is preferable for the classes of problems considered.

3.5.3.3 Implementation

This section details the input requirements of the optimisation code to be used. Specific issues relating to the implementation of path constraints are also discussed.

3.5.3.3.1 Input requirements The input file to the optimiser is a FORTRAN subroutine which is called by the driver program controlling the optimisation process. In addition to linking the problem description to the NLP solver and integrator, the driver provides a number of systems which enable the integrity of the problem description to be systematically checked.

The open format of the input file gives the user great flexibility in the types of problem which may be posed. Essentially, the information which must be included in the input file is as follows:

- a DASOLV formatted model of the dynamic system, *i.e.*:
 - model residual equations,
 - model jacobian equations in sparse form and sparsity pattern,
 - model discontinuity flags.
- definition of the objective function,
- definitions of problem constraints (equality constraints followed by inequalities),
- gradients of objective function and constraints with respect to the optimisation parameters,
- specification of optimisation parameters, *i.e.*:
 - pointers from the optimisation parameters to variables in the DASOLV model,
 - the set of parameters which are available to be optimised,
 - optimisation parameter bounds.

Note that in the case of time varying controls, it is their breakpoint values which are optimised. Additional subroutines are then used to determine the interpolated values for the controls together with their respective gradients.

In all of the case studies presented, the system of DAEs describing the process is first developed in the gPROMS (Barton, 1992) modelling environment. After validation, the gPROMS “CODE” facility is used to generate the required DASOLV formatted FORTRAN code required for the optimisation input file.

3.5.3.3.2 Path constraint representation. Recall from section 3.5 that the general formulation of a dynamic optimisation problem includes a set of inequality path constraints which must be satisfied for all times in the optimisation

$$q(x(t), y(t), u(t), d, t) \leq 0 \quad \forall t \in [0, t_f] \quad (3.1)$$

This infinite constraint is transformed into the single-valued differentiable constraint:

$$\int_0^{t_f} \max(0, q(x, y, u, d, t))^2 dt \leq c_{sig} \quad (3.2)$$

where c_{sig} represents a slight relaxation of the original constraint and is used to improve the numerical properties of the optimisation problem. The set of path constraints q may now be considered as part of the set of inequality point constraints c_2 . It should be noted that the gPROMS “CODE” facility does not generate analytical jacobian elements for functions involving the *max* operator. The path constraints are, however, differentiable and these elements are added manually.

3.6 Conclusions and introduction to case studies

This chapter has considered two possible approaches for the integrated design and operation of distributed chemical manufacturing facilities.

1. a single pass MINLP optimisation based method,
2. a decomposed problem which first develops a viable process structure and then establishes the design and operation of this structure.

The concept of cyclic operation was then introduced to enable the nominally uncertain operating profile of the process to be included in the design problem. Possible tools capable of simultaneously determining the design and operation of the process were discussed and dynamic optimisation shown to be the most appropriate technique. Following a general discussion on dynamic optimisation methods, the specific implementation details of the feasible path dynamic optimisation code to be used in this work were described.

The remainder of this thesis will consider the application of the design procedure and optimisation methods discussed earlier to a number of case study examples. The

case studies have been chosen to bring out the important aspects of a distributed plant design as discussed in section 2.2.5. Issues such as responsiveness, intermittent product demand and the use of process intensification to reduce the process footprint are all considered. The three cases presented share one important common feature: the dynamics and operation of the process are an integral part of the design problem. In all three cases, steady-state design alone is insufficient to produce the optimal, feasible design. This may only be found by consideration of the process operation. Recall that in this work a cyclic operation of the process is used to incorporate the process operation into the design problem. The specific form of the cycle to be used, together with its mapping to a dynamic optimisation formulation will be introduced at the relevant point for the case study being considered.

3.6.1 Multiproduct plant design

Chapter 4 considers the dynamic design of a multiproduct process which must manufacture four products in a repeated sequence. The example provides a “real world” example of the use of cyclic operation to give easily implementable and repeatable production schedules. The case study is based upon a fixed process structure which enables the second stage of the design procedure to be effectively demonstrated.

3.6.2 Wastewater plant design

New environmental regulations and limited construction space on existing sites have led to an interest in the design of compact wastewater treatment plants. Traditional chemical wastewater treatment plants typically include a main chemical treatment system and an emergency diversion system. Chapter 5 considers the design of a plant which integrates these two systems in an attempt to gain capital and space reductions.

3.6.3 Reactor network design

Chapter 6 applies the proposed design procedure to the optimal design and operation of a distributed process plant which is capable of instantaneous response to a customer’s intermittent demand for product. This case study is intended to represent a “worst-case” demand profile, since the process must be capable of automatic startup and shutdown. The case study illustrates a potential situation where environmental regulations on chemical storage make production *in situ* using a distributed plant the only viable means to satisfy a customer’s raw material requirements.

Chapter 4

Dynamic design of a continuous multiproduct process

4.1 Introduction

The flexibility of multiproduct and multipurpose plants is well known. They have the potential to produce many different products, typically in fairly small quantities, using the same basic equipment. Where restrictive constraints on cost and space exist, this degree of functionality is highly desirable. Benson and Ponton (1993) have rightly suggested that the ability to manufacture multiple products in a single set of equipment, simply by reconfiguring it, would be a highly desirable feature of a distributed plant. The multiproduct plant seems appropriate for use in a distributed manufacturing system; the integrated design and operation of such a process is therefore used as a motivating example for this thesis.

Multiproduct plants are typically scheduled in an operating cycle, with a single product being manufactured at any one time. Such an operation provides a “real-world” example of the use of cyclic operation to give repeatable and easily implementable operating strategies. Recall that this type of operating strategy was discussed in section 3.3.2 as a simplification enabling the operation of the process to be more easily incorporated into the design problem. This chapter is intended to introduce, by demonstration on a case study, the form of the dynamic optimisation problem used in the second stage of the design procedure proposed in section 3.3.3. Recall that this is to: *determine the optimal cyclic operation and design for a candidate process structure*. In this motivating case study, it is assumed that a candidate process structure is already in place, *i.e.* that stage one of the design procedure has been completed.

Section 4.2 provides an introduction to multiproduct and multipurpose plants. Two different modes of operation for these plants are then discussed. A historical perspective on work in the design and scheduling of multiproduct plants is then presented in section 4.2.2. The importance of incorporating product transitions into the scheduling problem is discussed in section 4.2.2.1 and transition dynamics are also considered. Design and operation of continuous multiproduct plants is shown to be a highly interactive, inherently dynamic problem. Section 4.3.1 describes a recently published algorithm which considers the dynamics of product transitions in the design and scheduling of the process. An alternative optimisation based approach is proposed and demonstrated in section 4.5. Results are presented and discussed in section 4.5.7.

4.2 Multiproduct and multipurpose processes

In this section the terms *multiproduct* and *multipurpose* with regard to chemical processes will be explained and distinguished. The motivation for designing and operating processes of this type will be discussed. A brief overview of this diverse field of research is then given. The review presented is by no means exhaustive, but is intended to be sufficiently extensive to provide a background which places the problem considered in this chapter into a suitable context.

4.2.1 Motivation for multiproduct and multipurpose processes

A sensible starting point is to define and differentiate the terms *multiproduct* and *multipurpose*. Suitably general descriptions of each type of process are as follows:

“In a **multiproduct plant** all the products follow essentially the same path and use the same equipment. Normally such plants produce one product at a time, a range of products being achieved by carrying out production runs or campaigns for each product in turn. A multiproduct plant is usually designed for a fixed set of products.

A **multipurpose plant** is rather like an engineering “job shop”, in which jobs are done on machines as they become available on completion of earlier jobs. Thus, there is no common pattern of movement for the products; indeed successive batches of the same product may take different routes through the process. In such plants two or more products may be produced simultaneously. Also the product set for a multipurpose plant is normally not fixed, new products are added, and old products either die off or pass on to new plants.”

(Sparrow *et al.*, 1975)

Based on these definitions one might consider multiproduct plants to be specific instances of multipurpose plants. In the literature, research relating to each class of process is clearly separable. The two types of process described above may be further subdivided according to their respective mode of operation, which may be *batch* or *continuous*. Previous research has considered the design and operation both of batch multipurpose and batch multiproduct processes. Continuous operation has more typically only been considered for multiproduct plants. Features of batch and continuous operation of multiproduct and multipurpose plants will now briefly be discussed.

Batch plants typically use fairly standardised equipment which is easily adapted and reconfigured to produce many different products. They are particularly suitable when a large number of high added value products must be manufactured in relatively small amounts in the same facility (*e.g.* pharmaceuticals and fine chemicals). If the plant is sufficiently flexible, different products may be manufactured either successively or simultaneously giving potentially low production lead times with rapid adaptation to market trends. A significant fraction of turnover in the chemical manufacturing industries is attributable to batch processes, *e.g.* 80% of pharmaceuticals and 65% of food and beverages (Reeve, 1992).

Continuous (multiproduct) plants arise frequently in the chemical process industries when a number of products with different specifications must be manufactured by a continuously operating system. When it is uneconomic to purchase a processing system for each product, multiple products must be scheduled on a much smaller set of processing units. Typical industrial examples of this type of process are: petroleum refineries which have to process different crudes, polymer and speciality chemicals plants, paper production and cement milling. Pinto and Grossmann (1993) have identified a current trend in the chemical industries towards continuous flexible multiproduct plants that can respond more quickly to demand changes and to the processing of a variety of products.

The two types of process (multiproduct and multipurpose) which are extensively used in industry have been identified and discussed. These types of process were further segregated by the nature of their operation (batch or continuous). In summary then, research in this area has typically focussed on the design and operation of:

1. continuous multiproduct processes,
2. multiproduct batch processes,
3. multipurpose batch processes.

This work represents a very large and diverse body of research indeed. In addition to the fundamental work on the design and operation of such processes, many extensions to these basic problems have been discussed in the literature, *e.g.*: incorporation of uncertainty (Shah and Pantelides, 1992), energy integration (Lee and Reklaitis, 1995a; Lee and Reklaitis, 1995b), combined maintenance and product scheduling (Dedopoulos and Shah, 1995) and even the management of entire supply chains (Wilkinson *et al.*, 1996), as well as solutions to the simultaneous design and scheduling problem (Voudouris and Grossmann, 1993). This chapter is concerned only with the design and operation of continuous multiproduct processes and it is therefore this literature which will primarily be discussed in the following section. Research in the area of multiproduct and multipurpose batch processes will be included where it is considered relevant. For reference, more comprehensive reviews of research in the area of multiproduct and multipurpose batch processes may be found elsewhere (Rippin, 1983; Reklaitis, 1990; Rippin, 1993).

4.2.2 Historical perspectives

The body of process systems work in the area of multiproduct continuous plants takes three primary perspectives:

1. process design,
2. product scheduling,
3. product transition dynamics and control.

Cantrell and Luyben (1996) give a useful insight into the part played by each of these areas in the design and operation of multiproduct plants. They state that the typical method for designing such plants is to cast the problem in terms of optimising multiple steady-states. Once the optimal steady-state plant has been designed, a control system is added which is both capable of moving the plant between the operating points required for each product and holding the plant at those operating points. Finally, a production schedule (*i.e.* the order in which products are generated) may be determined to meet supply requirements whilst minimising inventory and transition costs (the cost incurred while changing operating points). The schedule is typically a cyclic operation which is generally preferred by production managers as it is easy to implement.

As Cantrell observes: *the now obvious shortcoming with this organisation is performing tasks in sequence, rather than parallel.* As discussed in section 2.3 of this thesis, there is a growing realisation that simultaneous consideration of both the design

and operation of continuous processes is important. Similarly, a number of authors in the area of multipurpose batch plant design have stressed the importance of considering product scheduling as an integral part of the design problem (Király *et al.*, 1989; Shah and Pantelides, 1992). However, the three research areas in continuous multiproduct plants listed above have remained fairly isolated. Perhaps the only interaction which has been considered, albeit in a limited way, is that between the product scheduling problem and transitions between product grades.

4.2.2.1 Incorporating transitions into the scheduling problem

Sahinidis and Grossmann (1991) pose the problem of cyclic multiproduct scheduling on continuous parallel lines as a mixed integer non-linear program. The authors state that: *...product transitions and transition times are of critical importance for this (scheduling) problem.* They also observe that, since demand for the products will be simultaneous and that not all products can be manufactured simultaneously: *...proper modelling should take into account not only transitions, but inventory considerations as well.* They present a simple example to illustrate these points which it is useful to repeat below.

Two products A and B are to be produced on a continuous line. Figure 4.1 shows two feasible operating cycles: case (a) and case (b). In case (a) product A




Figure 4.1: Two feasible (cyclic) schedules

is produced initially, then production is changed to B and a transition time is incurred. Later, production will again be changed to A and another transition time will be incurred. This scheme (A-transition-B-transition) is repeated in time (cyclical operation). A second possible schedule is shown in figure 4.1 (case (b)). In case (b), the amounts of A and B produced are the same as in case (a), but the transitions between A and B occur more frequently. Schedule (a) might initially seem to be superior, since less penalty is paid for frequent transitions. However, the cost of inventory must also be considered. The

inventory varies with time and will have an associated cost, which in this example is proportional to the integral of the inventory level over time:

$$\text{inventory cost} = h \int_0^{t_{\text{cycle}}} I(t) dt \quad (4.1)$$

where h is the unit inventory holding cost, t_{cycle} is the processing time and I is the inventory level at time t . The inventory levels (I) are shown in figure 4.2 for cases (a) and (b). For simplicity, the transition times are not shown. In case (a) production of A




Figure 4.2: Inventory levels for alternative schedules

begins at time 0 and proceeds until time t , during which time inventory of A increases. At time t production of A stops and demand for A is supplied from inventory alone until time H , when the inventory level reaches zero. The inventory cost is proportional to the area of the triangle. In case (b), the same amount of A is produced, but over two identical cycles. The reduced cost of inventory is proportional to the shaded area shown. It can be seen that as the cycles become more frequent, the cost of inventory falls. However, the increased cycle frequency results in more transitions and increased transition costs.

Based on this example, the authors conclude that there must exist an *optimal cycle time which is determined largely by the inventory and transition costs*. In previous work, simplifications to the problem were made which involved either ignoring product transitions and changeovers, or by assuming them to be sequence independent. Sahinidis and Grossmann “rigorously” incorporate the costs associated with transition periods by applying a fixed charge term wherever a transition occurs. Because the design of the process is fixed, the transition times are known *a priori* and may be considered in this way.

4.2.2.2 Dynamics and control of product transitions

The importance of product transitions in the scheduling of multiproduct plants is clear. Making such transitions is, however, a *dynamic* problem which has, in turn, been addressed by several authors. Early work in the field by Sargent and Sullivan (1979) considered changing the product profile of a petroleum refinery. They discuss a sub-optimal control system design which uses open-loop dynamic optimisation to determine optimal control trajectories and the implementation of sub-optimal feedback control laws which retain the dominant features of these trajectories. More recently, McAuley and MacGregor (1992) have studied the optimisation of product changeovers in a polyethylene reactor in which three grades of material are manufactured. In this work an objective function is chosen which simultaneously penalises the production of off-spec material during the grade change transition (which has an associated cost) and forces the system to a pre-determined steady-state at some future time. A time-optimal strategy for distillation column startup and product changeover has recently been published by Kruse et al. (1996). The authors make the point that rapid transition is desirable because: *...dynamic transitions are always non-productive procedures*. The research described above shares important common ground: it is all based upon a *fixed* process design. Good (rapid) transitions are achieved purely by determination of appropriate control trajectories. White *et al.* (1994) have applied optimal control techniques to investigate the “switchability” of chemical processes. White (1994) qualitatively defines switchability as “*a plant’s ability to move between operating points*”. An important extension in this work was to allow the design of the process as well as the control trajectories to be optimised to give the optimal transition. Significantly, White showed that the design of the process could have a considerable effect on its switchability.

4.2.3 Conclusions

The discussion presented in section 4.2.2 above can be summarised in diagrammatic form as shown in figure 4.3. The figure shows the way in which the process design and control system interact with the product schedule and product transitions. It is very clear from this diagram that a sequential approach to the design, control and scheduling of continuous multiproduct plants is far from adequate. Because of the many economic trade-offs which interact in this problem (one of which is cost of inventory versus cost of transitions, another which has not yet been mentioned is the cost of the plant itself) a more integrated methodology is desirable. Recently, a methodology for continuous multiproduct plant design has been published which incorporates the dynamics of transitions.

The next section will review this work and examine potential benefits and shortcomings. Following this, an alternative approach based on the optimal cyclic operation discussed earlier is presented.




Figure 4.3: Cycle summarising interactions between process design, product transitions and production schedule

4.3 Integrated design of continuous multiproduct processes

In the preceding section the intrinsically dynamic nature of continuous multiproduct processes was discussed. Research in the area of scheduling these plants has noted that product transitions are of *critical* importance in determining an economically optimal schedule *and* has incorporated fixed penalty terms to reflect this in their economic objective functions. Since the nature of the product transitions is determined largely by

the design and control of the process, development of a truly optimal process operation should consider the process design as a degree of freedom. In other words, the process design and operation (*i.e.* schedule) should be considered together.

The shortcomings of a sequential approach to the design and operation of multiproduct processes have recently been addressed in the literature by Cantrell and Luyben (1996) and will now be discussed.

4.3.1 A multiproduct process design procedure

In a recently published paper, Cantrell and Luyben recognised that: “*time to transition (sic)... can significantly impact on the overall operating cost by changing the cycle time and inventory profiles*”. Their solution is a design procedure which uses elements of operational research, process control theory and heuristic process design to develop a multiproduct plant and cyclic production schedule which more rigorously incorporates the dynamics of transitions.

Before outlining the design procedure it is first necessary to introduce some of background theory used in this procedure.

4.3.1.1 Simple multiproduct cycle time model

The cycle time model is used to determine the amount of time needed to generate each of the products in the production cycle. In order to do this, the demand for each product, production rate of each product and transition times between products must be known. The production times partly determine the total cycle time of the process and will therefore have an impact on the overall cost. In the procedure proposed by Cantrell and Luyben the demand rates form part of the problem specification, the production rates are determined as part of the design problem and the transition times are a function of the process and controller design and are found using dynamic simulation.

The cycle time calculation presented by Cantrell and Luyben is a form of the economic lot scheduling problem from operations research (Hanssmann, 1962) and makes three assumptions:

1. there is no slack time,
2. there is a common cycle for all products¹,
3. transitions between products are sequence dependent.

¹each product is produced once per cycle

Thus, a process producing n products requires n steady-state periods connected by product transitions. When production of the n th product ends the cycle is repeated.

Based on these assumptions, an expression for the lengths of the production periods may be derived using a balance on the rates of demand and production of each product over a complete cycle:

$$p_i t_i = r_i \left(\sum_{i=1}^n t_i + \theta \right) \quad i = 1, \dots, n \quad (4.2)$$

where p_i is the rate of production of product i , t_i is the length of the production period for product i , r_i is the rate of demand for product i and θ is the sum of the transition times for a specified production schedule:

$$\theta = \sum_{n_{trans}} s_{ij} \quad (4.3)$$

where n_{trans} is the number of transitions and s_{ij} is the transition time from product i to the next product, j , in the specified schedule. Note that the total cycle time is given by the following expression:

$$t_{cycle} = \sum_{i=1}^n t_i + \theta \quad (4.4)$$

Equation 4.2 may be rearranged to give an explicit expression, in matrix form, for the n production times (t_i):

$$\mathbf{T} = \theta [\mathbf{P}^{diag} - \mathbf{RL}]^{-1} \mathbf{R} \quad (4.5)$$

In equation 4.5, $\mathbf{T}$ is an $n \times 1$ matrix of the production times (t_i), $\mathbf{P}^{diag}$ is an $n \times n$ matrix with the production rates (p_i) on the diagonal and zeros elsewhere. $\mathbf{R}$ is an $n \times 1$ matrix of the product demand rates (r_i) and $\mathbf{L}$ is a $1 \times n$ matrix of ones.

Given the n demand and production rates, the following necessary condition must be satisfied in order for the cycle time to be positive and finite:

$$\sum_{i=1}^n \frac{r_i}{p_i} < 1 \quad (4.6)$$

In the case where the right-hand side of equation 4.6 is exactly equal to one, the cycle time is infinite and the bracketed term in equation 4.5 becomes singular.

4.3.1.2 Inventory models

Recall that demand for a given product must be met from stored inventory during the periods when that product is not generated by the process. The cost of maintaining this inventory can be substantial and the authors describe a number of models which can be used to account for this cost:

- A time integral triangle model (see equation 4.1),
- a working capital model based upon the expected return on the production cost of stored product,
- storage capacity capital cost.

In their work Cantrell and Luyben use the last of these models which determines inventory costs from the maximum inventory level over a cycle.

The background to the problem has been introduced together with the relevant theory and assumptions made. The procedure in which the cycle time and inventory models are applied will now be discussed.

4.3.1.3 Design procedure

The basic procedure for the design of continuous multiproduct processes proposed by Cantrell and Luyben is shown in figure 4.4. Each stage of this procedure will now be considered in turn.

- Stage one

The proposed procedure requires a number of *a priori* decisions to be made at stage one which fall into two categories:

1. specifications on the final design,
2. simplifications to make the problem tractable.

In the first category are decisions such as the product demand rates and the product specifications, which are required to form the problem definition. The decisions made in the second category essentially fix both the process structure and the control structure. Cantrell and Luyben do make the point that the control structure may be included as a design choice if it is not known prior to design. They do not, however, provide guidance on how this should be incorporated.

- Stage two

Cantrell and Luyben give limited guidance on the selection of design parameters. They suggest that parameters with the most effect on the transitional dynamics are those which change control loop interactions and also the dominant time constants (*e.g.* liquid holdups). Optimising the controller settings is not mentioned by the authors, but in the light of the above advice would seem to be important at this stage. Since the procedure proposed has an outer iterative loop around the selection

of design parameters it is necessary to keep the number of such parameters to a minimum in order to avoid excessive computational effort.

- Stage three

The authors suggest a suitable initial guess for $\mathbf{P}$ as:

$$\mathbf{P}_{guess} = \frac{n+1}{n} \mathbf{P}_{min} \quad (4.7)$$

where $\mathbf{P}_{min}$ is a scaling factor on which the production rate profile depends (*e.g.* feed rate to the process) and n is the number of products. The requirement that $\mathbf{P}$ should be chosen to satisfy the necessary condition given in equation 4.6 is assumed to apply here, but is not discussed by the authors. On later iterations, $\mathbf{P}$ is found by optimisation. The authors state that the overall cost is convex with respect to the production rate, which enables the optimal production rate to be located via a simple search technique. Given that a production rate profile must be determined for each product it is not clear that this procedure is as straightforward as the authors seem to suggest.

- Stage four

As the authors observe, the simulation load can become large as the number of possible transitions is $n(n-1)$.

- Stage five

If the transition is sequence dependent there are $(n-1)!$ possible schedules to enumerate. In the case study presented a heuristic approach is taken to eliminate a number of potential schedules.

- Stages six and seven

The overall cost is straightforward to calculate for the current design and production schedule.

- Stage eight

No guidance on how the design parameter values should be updated is given, or indeed if the set of design parameters should be changed for the next iteration in the light of results for the current iteration.

The authors themselves identify that three key limitations of the procedure presented are: a high simulation load to determine transition times; a potentially large set of possible schedules to enumerate; optimisation of the design variables. The last observation is a

significant limitation, as one would expect effective optimisation of the design variables to be fundamental to the success of any procedure which aims to determine optimal process designs.




Figure 4.4: Procedure for continuous multiproduct process design proposed by Cantrell and Luyben

4.3.2 Optimisation based multiproduct process design

The purpose of the procedure described above is essentially to determine the economically optimal operating cycle of a (multiproduct) process by finding the best combination of design parameters within that fixed process and control structure. Previously in section 3.3.2, operation of a process in a cyclical manner was presented as a simplification which enabled the process operation to be considered at the process design level. It was shown that dynamic optimisation techniques could be used to determine the optimal cyclic operation and design of a given process. The dynamic optimisation problem discussed formed the second stage of a two-stage design procedure. Recall that in the first stage a viable candidate process structure is developed for further analysis using numerical optimisation. The dynamic optimisation problem discussed in chapter 3 has similar aims to the procedure of Cantrell and Luyben, but may be considered more generally applicable: *given a process structure, determine its optimal design and cyclic operation*. The similarity in objectives is not surprising given that in both cases the motivation stems from a need to consider process operation and design simultaneously.

In the remainder of this chapter, solution of the second stage of the proposed design procedure will be demonstrated by application to the dynamic design of a multiproduct plant. The case study used by Cantrell and Luyben to demonstrate their design procedure will be considered.

4.4 Multiproduct reactor case study

In this section dynamic optimisation will be used to determine the optimal design and cyclic operation of a multiproduct process. One important difference between the optimisation based approach presented below and the methodology used by Cantrell and Luyben is that here the dynamics of *both* the reactor *and* the storage vessels are considered. This contrasts with the previous work which dealt only with the reactor dynamics and appended the storage behaviour based on a simplified cycle time model. The optimisation based method may therefore be considered a more integrated approach to this design problem.

4.4.1 Problem definition

A continuous stirred tank reactor (CSTR) is to be used to manufacture four grades of product B through the reaction scheme:



Specifications on the four products ($B_1 \rightarrow B_4$) are summarised as follows. Purity requirements are given in terms of concentration of A :

- B_1 : 25 lbmol.hr^{-1} continuous demand with purity $0.0500 \pm 1 \times 10^{-4} \text{ lbmol.ft}^{-3}$.
- B_2 : 25 lbmol.hr^{-1} continuous demand with purity $0.0375 \pm 1 \times 10^{-4} \text{ lbmol.ft}^{-3}$.
- B_3 : 25 lbmol.hr^{-1} continuous demand with purity $0.0250 \pm 1 \times 10^{-4} \text{ lbmol.ft}^{-3}$.
- B_4 : 25 lbmol.hr^{-1} continuous demand with purity $0.0125 \pm 1 \times 10^{-4} \text{ lbmol.ft}^{-3}$.

The CSTR produces one of these grades at a given time and so the continuous demand for all four products must be met by storing a quantity of each product. The operation of the process is cyclic and each product is manufactured only once in a single cycle. Any off-specification material produced during the cycle has an associated cost implication. The full process to be considered incorporating reactor, control system and storage is shown in figure 4.5. Control of the product composition within the reactor is achieved using a primary composition controller which adjusts the setpoint on a secondary temperature controller. The desired temperature is achieved by manipulating the cooling water flow through the CSTR jacket. A proportional-integral controller is used in the composition control loop. This loop has an associated five minute deadtime. A proportional controller is used for the secondary temperature control loop. This loop has an associated one minute lag. The maximum allowable cooling water flow is set at four times the steady-state value. The CSTR is assumed to have constant volume and the production rate of each product is therefore determined by the feed rate of A to the process. The problem is essentially to meet the product supply requirements with the lowest annualised cost by finding the optimal process design, controller settings and cyclic scheduling of the products.

4.4.2 Process model

Having outlined the definition of the problem, this section discusses the modelling of the system. The model of the system is considered in three parts:

1. reactor modelling,
2. controller modelling,
3. storage modelling.

Each of these modelling aspects will now be dealt with in turn.




Figure 4.5: Multiproduct process case study

4.4.2.1 Dynamic reactor model

The dynamic behaviour of the CSTR is represented by a set of differential and algebraic equations based on dynamic mass and energy balances around the unit.

1. feed component material balance,

$$\frac{dC_{out}}{dt} = \frac{F_{in}}{V_R}(C_{in} - C_{out}) - kC_{out} \quad (4.9)$$

where C_{out} is concentration of A in the reactor ($lbmol.ft^{-3}$), F_{in} is the volumetric feed rate ($ft^3.hr^{-1}$), V_R is the reactor volume (ft^3), C_{in} is the concentration of A in the feed and k is the reaction rate (hr^{-1}).

2. reactor energy balance,

$$\frac{dT_{out}}{dt} = \frac{F_{in}}{V_R}(T_{in} - T_{out}) - \frac{\Delta H_R k C_{out}}{\rho C_p} + \frac{U A_H}{\rho V_R C_p}(T_{c,out} - T_{out}) \quad (4.10)$$

where T_{out} is the reactor temperature ($^{\circ}F$), T_{in} is the temperature of the feed, $T_{c,out}$ is the jacket temperature ($^{\circ}F$).

3. jacket energy balance,

$$\frac{dT_{c,out}}{dt} = \frac{F_c}{V_J}(T_{c,in} - T_{c,out}) - \frac{U A_H}{\rho_J V_J C_{p,J}}(T_{c,out} - T_{out}) \quad (4.11)$$

where V_J is the jacket volume, F_c is the volumetric feed rate to the jacket and $T_{c,in}$ is the cooling water feed temperature.

The reaction rate, k , is defined by:

$$k = k_0 e^{\left\{ \frac{-E_a}{R(460+T_{out})} \right\}} \quad (4.12)$$

values for the reaction rate parameters and physical constants used in the above equations are given in table 4.1. The heat transfer area (in ft^2) between jacket and reactor, A_H , is defined by:

$$A_H = 4.65 V_R^{\frac{2}{3}} \quad (4.13)$$

The jacket volume (in ft^3) is defined in terms of the reactor volume as follows:

$$V_J = 1.55 V_R^{\frac{2}{3}} \quad (4.14)$$

Parameter	Value	Parameter	Value
k_0	$8.16 \times 10^{10} \text{ hr}^{-1}$	E_a	30000 BTU/lbmol
ΔH_R	-30000 BTU/lbmol	R	$1.99 \text{ BTU/(lbmol } ^\circ\text{R)}$
C_{in}	0.50 lbmol/ft^3	U	$150 \text{ BTU/(hr ft}^2 \text{ } ^\circ\text{R)}$
T_{in}	$70 ^\circ\text{F}$	$T_{c,in}$	$70 ^\circ\text{F}$
C_p	$0.75 \text{ BTU/(lb}_m \text{ } ^\circ\text{R)}$	$C_{p,J}$	$1.0 \text{ BTU/(lb}_m \text{ } ^\circ\text{R)}$
ρ	$50 \text{ lb}_m/\text{ft}^3$	ρ_J	$62.3 \text{ lb}_m/\text{ft}^3$

Table 4.1: Model parameters for multiproduct plant case study

4.4.2.2 Controller modelling

As stated earlier, control over the reactor composition is achieved through two control loops: a primary proportional-integral composition controller and a secondary proportional temperature controller.

4.4.2.2.1 Composition controller Recall that there is a deadtime associated with the composition controller. In the model this is approximated by a number, n , of first order lags in series:

$$\begin{aligned} \left(\frac{t_d}{n}\right) \frac{dC_{dt,1}}{dt} &= (C_{out} - C_{dt,1}) \\ \left(\frac{t_d}{n}\right) \frac{dC_{dt,i}}{dt} &= (C_{dt,i-1} - C_{dt,i}) \quad i = 2, \dots, n \end{aligned} \quad (4.15)$$

where t_d is the deadtime and $C_{dt,i}$ is the outlet concentration of the i th lag block. The proportional integral controller is then defined by the equations:

$$\xi_C = C_A^{sp} - C_{dt,n} \quad (4.16)$$

$$\tau_I \frac{d\xi_I}{dt} = K_c^C \xi_C \quad (4.17)$$

$$T_R^{sp} = K_c^C \xi_C + \xi_I \quad (4.18)$$

where ξ_C is the composition controller error, C_A^{sp} is the composition setpoint, ξ_I is the integral error, τ_I is the controller reset time, T_R^{sp} is the setpoint on the temperature controller and K_c^C is the composition controller gain. In the work which follows a value for n of twenty was used to give a reasonable approximation to the deadtime whilst avoiding an excessive increase in the size of the model.

4.4.2.2.2 Temperature controller The 1 minute lag in the temperature control loop is incorporated by the simple first order lag model:

$$\tau_T \frac{dT_R^{meas}}{dt} = (T_{out} - T_R^{meas}) \quad (4.19)$$

where τ_T is the lag time in the control loop and T_R^{meas} is the lagging temperature measurement. The cooling water flow rate F_c^c is calculated as follows:

$$\xi_T = T_R^{sp} - T_R^{meas} \quad (4.20)$$

$$F_c^c = K_c^T \xi_T \quad (4.21)$$

where ξ_T is the temperature controller error. Recall that the coolant flow is bounded between zero and four times the steady-state flow. The temperature controller output must therefore be clipped to ensure that it remains within these bounds. In doing this it is important to avoid introducing significant discontinuities into the model which can cause numerical problems during integration of the model, such as failures when evaluating sensitivities or cycling across the discontinuity. Walsh (1993) presents integral freezing in a controller as an example of this. Smoothed clipping functions with continuous first derivatives (based upon those suggested by Walsh (1993) for achieving smooth integral desaturation) are therefore implemented:

$$K_{clip} = \begin{cases} 1 & \text{if } \epsilon < F_c^s < 1 - \epsilon \\ \frac{3}{\epsilon^2}(F_c^s)^2 - \frac{2}{\epsilon^3}(F_c^s)^3 & \text{if } 0 \leq F_c^s \leq \epsilon \\ \frac{3}{\epsilon^2}(1 - F_c^s)^2 - \frac{2}{\epsilon^3}(1 - F_c^s)^3 & \text{if } 1 - \epsilon \leq F_c^s \leq 1 \\ 0 & \text{otherwise} \end{cases}$$

where ϵ typically between 0.01 and 0.05 and F_c^s is the scaled controller output :

$$F_c^s = \frac{F_c^c}{F_c^{range}} \quad (4.22)$$

F_c^{range} is the cooling water flow range (four times the steady-state value). The actual cooling water flow to the process is therefore given by:

$$F_c = \begin{cases} F_c^{range} F_c^s K_{clip} & \text{if } F_c^s < 0.5 \\ F_c^{range} (1 - (1 - F_c^s) K_{clip}) & \text{otherwise} \end{cases} \quad (4.23)$$

4.4.2.3 Dynamic storage model

As mentioned previously, because the multiproduct plant problem is to be solved within a dynamic optimisation framework, the dynamics of the storage vessels must be

considered explicitly in the model of the process. Modelling of the storage will be discussed in this section.

The storage tanks are assumed to be well mixed, with no transfer of energy transfer taking place. Storage tank i is therefore described by the following dynamic overall and component mass balances:

$$\left. \begin{aligned} \frac{dV_{store_i}}{dt} &= p_i - r_i \\ V_{store_i} \frac{dC_{A_i}}{dt} + C_{A_i} \frac{dV_{store_i}}{dt} &= C_{out}p_i - C_{A_i}r_i \end{aligned} \right\} \quad i = 1 \text{ to } 4 \quad (4.24)$$

where V_{store_i} is the volume of storage tank i , p_i is the volumetric flow of feed to storage tank i from the reactor, r_i is the volumetric flow of product from storage tank i , C_{out} is the composition of product leaving the reactor and C_{A_i} is the composition of product leaving storage tank i .

The flowrate of off specification material, $F_{offspec}$, may then be calculated from an overall material balance on the reactor:

$$F_{offspec} = F_{in} - \sum_{i=1}^4 p_i \quad (4.25)$$

The amount of off specification material generated during a cycle, $V_{offspec}$, is then given by:

$$V_{offspec} = \int_0^{t_{cycle}} F_{offspec}(t) dt \quad (4.26)$$

where t_{cycle} is the cycle time of the process. It should be noted that for the purposes of this case study, the physical properties of the feed and product streams are assumed to be equal.

4.4.2.3.1 Addition of a make-up controller for numerical robustness As discussed in chapter 3, the major tool for carrying out the integrated design and operation in this work will be dynamic optimisation. A model which accurately represents the physical system and is robust during the course of a dynamic simulation will not necessarily be numerically robust in the regions that may be entered during solution of the dynamic optimisation problem. It is therefore important to ensure that process models to be used for optimisation purposes are as numerically reliable as possible. As an aside, it should be noted that placing “soft constraints” (*e.g.* inequality constraints) on some variables such that numerically problematic regions lie outside the feasible region will not overcome the problem since such constraints are satisfied *only* at the solution. The storage model described above must therefore be reconsidered in this light.

As has already been stated, the profile over time for the exit flow is fixed and cannot be manipulated. During the course of a dynamic optimisation it is possible for a

situation to arise in which the storage tank empties *and* product continues to be drawn at a rate which is greater than the flow of feed to the storage tank. This will cause the holdup to violate variable bounds and the equations describing the storage tank concentration to become physically meaningless. The dynamic storage may be made robust by the addition of a proportional level controller on the tank as follows.

$$Q_{mu} = \begin{cases} 0.0 & \text{if } V_{store_i} > V_{store_{i,l}} \\ K_c(V_{store_{i,l}} - V_{store_i}) & \text{otherwise} \end{cases} \quad (4.27)$$

where K_c is a controller gain and $V_{store_{i,l}}$ is a lower limit on the storage tank holdup. The level controller ensures that the holdup is at least at $V_{store_{i,l}}$ during the optimisation by using a “dummy” make-up stream (Q_{mu}) with a concentration equal to that of the tank itself. A path constraint on the tank holdup ensures that the proportional controller does not operate at the solution point (see 4.5.5.2.1).

4.5 Dynamic design of a multiproduct process.

In the preceding sections the problem to be addressed has been defined. The basic structure of the process together with its control system was presented and a mathematical description given in terms of a set of differential-algebraic equations. A recently published procedure for the design of multiproduct processes which allows integration of the process design, process control and cyclic process scheduling was reviewed. A parallel was drawn between this procedure and the cyclic optimisation framework used for this work. The next section will show how the dynamic design and operation of a multiproduct process may be achieved within the general cyclic optimisation framework discussed previously.

4.5.1 Framework for the multiproduct process optimisation


In this section the form of the process operating cycle to be used in this case study is defined. Simplifications and assumptions relating to the problem will also be explained.

Recall from section 4.4.1, where the problem definition for this case study was set out, that:


- the operation of the process is cyclic,
- each product is manufactured once per cycle (primal cycle),

- at a given time in the cycle the process is either in a transition period or a production period.

Graph (a) in figure 4.6 shows a typical time profile for the four product case study as described by the cycle time model used by Cantrell and Luyben (see section 4.3.1.1). The notation used is the same as in section 4.3.1.1: t_i indicates the length of production period i and s_{ij} indicates the transition period from product i to j . In the design procedure



(a) process cycle showing product transitions



(b) process cycle showing underlying setpoint changes

Figure 4.6: Time profile for system described by cycle time model

proposed by Cantrell and Luyben, the values of s_{ij} are determined by simulation of each possible transition from i to j . These are then used in the cycle time model for a chosen schedule to compute t_i . The capital and transition costs may then be determined.

Also shown in graph (b) of figure 4.6 is the profile of the composition setpoint for this case study which corresponds to the transitions above. Each of the products to be manufactured has an appropriate setpoint value. Since four products are to be produced in a primal cycle, the operating cycle may be split into four periods, each with the setpoint required to make one of the products. In this case study it is therefore the *sequence in which the composition setpoint changes which determines the process schedule*. By in-

cluding the setpoint profile for the cycle as a time varying optimisation variable (control), the optimal production schedule may be integrated into the optimisation problem.

Optimisation of the reactor setpoint profile determines the sequence in which the products will be manufactured by the reactor. The optimisation framework must also be flexible enough to enable:

- on-spec products to be fed to storage for any production schedule,
- off-spec material produced during transitions to be diverted as waste.

This is achieved by subdividing each of the four periods shown in graph (b) figure 4.6 into a transition period and a production period (as shown in graph (a) of the same figure). Manipulating the product flows (p_i) allows the storage tanks to be filled with on-spec product during the production periods. Off spec product produced in the transitions may be prevented from entering the storage tanks by setting the product flows to zero. This time discretisation is considered flexible enough to achieve the aims set out above, whilst being straightforward to implement.

The framework set out above will form the basis of the dynamic optimisation problem to be solved. The constituent parts of the optimisation problem will be discussed in detail in the sections which follow.

4.5.2 Objective function

The objective function for this case study is to minimise the annualised cost of the process. This cost is made up of three constituent parts:

1. capital cost of the reactor,
2. capital cost of the storage,
3. transition costs.

4.5.2.1 Capital costs

The costs of the reactor and the storage vessels are related to the total molar holdups in the respective vessels according to the following relationships:

$$C_{CSTR} = 2523 (H_{reactor})^{0.62} \quad (4.28)$$

$$C_{store_i} = 2523 (H_{store_i}^{max})^{0.62} \quad i = 1 \text{ to } 4 \quad (4.29)$$

where C_{CSTR} is the reactor cost in dollars, C_{store_i} is the cost in dollars of the i th storage tank, $H_{reactor}$ and $H_{store_i}^{max}$ are the maximum molar holdups in the reactor and i th storage vessel respectively. The total annualised capital cost of the process is therefore given by:

$$C_{capital} = \frac{1}{3} \left\{ \sum_{i=1}^4 C_{store_i} + C_{CSTR} \right\} \quad (4.30)$$

The $\frac{1}{3}$ premultiplier is a capital charge factor (CCF) which approximates “the time value of money” (Douglas, 1988), and converts the total capital cost to an annualised equivalent.

4.5.2.2 Transition costs

In addition to the capital cost of the processing equipment and storage vessels, there is a cost incurred for any material produced during the operating cycle which is not to specification. This cost represents a devaluation for being off-spec or a cost for extra processing required to achieve the desired product. The off-spec is charged at 0.1 \$/lbmol⁻¹. The amount of off-spec produced per cycle, $V_{offspec}$, was defined earlier in equation 4.26. If it is assumed that the process operates continuously, in a given year there will be:

$$N_{cycle} = \frac{365 \times 24}{t_{cycle}} \quad (4.31)$$

cycles, where t_{cycle} is the cycle time of the process in hours. The total cost per year for off-spec is therefore equal to:

$$C_{transition} = 0.1 V_{offspec} \rho_m \left(\frac{8760}{t_{cycle}} \right) \quad (4.32)$$

where ρ_m is the molar density of the product.

The total annualised cost for the process is the sum of transition and capital costs and the objective function is given by:

$$\min_{u(t), d} \{C_{capital} + C_{transition}\} \quad (4.33)$$

4.5.2.3 Economic trade-offs

Having introduced the component parts of the annualised cost objective function it is useful to briefly consider how these costs interact and the economic trade-offs that may be expected to arise in this process.

The most obvious economic conflict in the above economics is between the storage costs and the transition cost. Equation 4.32 shows that the transition cost is inversely proportional to the time taken to complete a cycle. However, the storage required during

a cycle is directly proportional to the cycle time. In other words, the storage cost is reduced by many cycles of short duration, but this leads to more off-spec product being generated and thus higher waste costs. Of course, the cycle time itself is ultimately determined by the design of the reactor. For example, a more expensive reactor may offer an overall economic benefit by allowing smooth, rapid transitions which generate less off-spec material and enable frequent cycling.

4.5.3 Optimisation parameters

The aim in solving this optimisation problem is to determine the optimal process design and cyclic production schedule. This is achieved by manipulation of a number of degrees of freedom in the process:

1. Time varying parameters (controls).

It was stated earlier that the product scheduling is determined by the profile of the reactor setpoint with time. This setpoint must therefore be included as an optimisation parameter in the process. The remaining control variables to be optimised are the flowrates of each of the four streams feeding the storage tanks. The nature of these control profiles will be considered further when the time discretisation and constraints are discussed.

2. Process initial state.

The form of the dynamic optimisation problem described earlier in chapter 3 includes the process initial condition as a degree of freedom in the optimisation.

In this case study, only the initial state of the storage system (*i.e.* initial volume and product concentrations) are optimised as the reactor initial state is determined by the composition setpoint at t_0 . This will be discussed further in section 4.5.5.2.2.

3. Time invariant optimisation parameters

There are a number of design parameters within the reactor which may be optimised and which are fixed over time. In this case study, these time invariant parameters are: the feed to the reactor, the reactor residence time and the controller tuning parameters (gain and reset time for the composition controller and gain for the temperature controller).

4. Equipment sizes.

As stated earlier, the optimisation based approach used in this work considers the dynamic behaviour of the storage vessels as an integral part of the design problem. The sizes of these vessels (which appear in the objective function) are therefore optimised explicitly within this framework.

It is worth pausing at this point to compare the optimisation parameters used here to those considered by Cantrell and Luyben. The basic approach taken by Cantrell and Luyben in the solution of this case study is:

- fix the reactor residence time,
- determine the controller settings for this residence time using Ziegler-Nichols settings and a nominal feed rate,
- choose a production schedule,
- determine the feed rate for this system and cyclic schedule which minimises cost,
- repeat for different values of reactor residence time.

The impact of simultaneous optimisation of all of these degrees of freedom, which a wholly optimisation based method allows, will be discussed in the light of results from both methods.

4.5.4 Time discretisation and control parameterisation

Earlier in section 4.5.1 the form of the operating cycle to be used in the optimisation was discussed. The implementation of this time discretisation in practice will now be described together with the functional form of the controls to be used.

The cycle time is discretised into sixteen intervals as shown in figure 4.7. This figure also depicts typical trajectories for the setpoint profile and each of the four flowrates to the storage tanks. The discretisation essentially consists of the eight variable length time intervals required for the transition-production form of the cycle described in section 4.5.1, each of which is separated by a fixed interval of length 10^{-3} hours. The shorter intervals allow rapid ramping of piecewise linear setpoint and flow trajectories such that they approximate a piecewise constant profile. Referring back to section 4.5.1, intervals 2,6,10 and 14 in figure 4.7 correspond to transition periods whilst intervals 4,8,12 and 16 correspond to production periods. In order to obtain the form of the setpoint and flowrate trajectories shown in the figure, a number of constraints must be applied in the problem description. These will be described in more detail in the following section which deals with constraints in this case study.




Figure 4.7: Time discretisation used in multiproduct plant case study

4.5.5 Constraints

The optimal design and operation of the multiproduct plant is subject to a number of constraints which arise either from operational requirements on the process or as a result of the optimisation framework itself.

4.5.5.1 Operational constraints

In section 4.4.1 the operating constraints on the process were defined. These consisted of limits on the purities of each of the products. The purity constraints must be enforced for all times in the cycle and are therefore implemented as inequality path constraints. It is important to note that these constraints are applied to the products flowing *from the storage tanks*. In their approach, Cantrell and Luyben apply these constraints to the reactor effluent because the storage dynamics are not considered. This

may prove to be conservative as potentially beneficial product buffering resulting from the storage holdup is not taken into account. The path constraints to be implemented are:

$$\begin{aligned} 0.0499 &\leq C_{A_1} \leq 0.0501 \\ 0.0374 &\leq C_{A_2} \leq 0.0376 \\ 0.0249 &\leq C_{A_3} \leq 0.0251 \\ 0.0124 &\leq C_{A_4} \leq 0.0126 \end{aligned} \tag{4.34}$$

4.5.5.2 Non-operational constraints

In addition to the operational constraints described above, a number of constraints are required by the optimisation framework. These might, for example, be used to enforce cyclic operation. The non-operational constraints may be either path constraints or equality point constraints.

4.5.5.2.1 Path constraints The maximum holdup in the storage vessels over the operating cycle must be determined so that the capital cost of the storage tanks can be calculated. This is achieved by applying the following path constraint:

$$V_{store_i} \leq V_{store_i}^{max} \quad i = 1 \text{ to } 4 \tag{4.35}$$

where $V_{store_i}^{max}$ is a time invariant optimisation parameter.

Recall that equation 4.25 in section 4.4.2.3 was used to define the flow of off-spec material from the process. Since the feed flow to the process and flows from the reactor to the storage vessels are included as degrees of freedom in the optimisation, a path constraint is required to ensure that equation 4.25 is meaningful by constraining the flow of off-spec to be greater than or equal to zero at all times:

$$0 \leq F_{offspec} \tag{4.36}$$

Previously in section 4.4.2.3.1, numerical robustness in the storage tank model was achieved through a proportionally controlled make-up stream described by equation 4.27. This was added to prevent negative holdup in this tank. At the solution point the make-up stream must not be active and a further path constraint must be added to ensure this. Referring back to equation 4.27, it can be seen that the make-up stream becomes active when the volume in the storage vessel falls to $V_{store_i,l}$. The following path

constraint requires V_{store_i} to be a factor of κ times the volume at which the make up controller is triggered. κ is a constant whose value must be greater than one.

$$\kappa V_{store_{i,l}} \leq V_{store_i} \quad (4.37)$$

For this work a value for κ of 2 was used.

4.5.5.2.2 Equality point constraints The remaining constraints used in this formulation are equality point constraints. These are required for two purposes:

- to enforce the cyclic operation,
- to give the correct form for the control trajectories as discussed in section 4.5.4.

These two sets of constraints will now be discussed in turn.

1. Cyclic constraints.

As has already been stated, the aim of this optimisation formulation is to determine the design and *cyclic schedule* of a multiproduct process. A suitable set of endpoint constraints must therefore be applied at the final time in the cycle to enforce cyclical operation of the process. Since all the initial values of the states in the storage vessels are optimised, a set of endpoint constraints is required to ensure that these states are cyclic:

$$\left. \begin{aligned} C_{A_i}(t_0) &= C_{A_i}(t_{cycle}) \\ V_{store_i}(t_0) &= V_{store_i}(t_{cycle}) \end{aligned} \right\} \quad i = 1 \text{ to } 4$$

The optimisation formulation must also ensure that the reactor state is the same at the final time as at the initial time. The reactor model contains some twenty six states because of the need to model deadtime and lag blocks in the controllers. To avoid introducing unnecessary complexity into the problem by specifying cyclic endpoints on each of these states, the simplification is made that the reactor is at steady-state at time zero in the cycle. The assumption here is that the controllers implemented settle to steady state, such that at t_{cycle} the plant is approximately at steady-state. Given that very tight product specifications are imposed and that only one production period is allocated for the manufacture of each product, this assumption is not unreasonable. The initial state of the reactor is now defined explicitly by the composition setpoint and a cyclic constraint may be applied as follows:

$$C_A^{sp}(t_0) = C_A^{sp}(t_{cycle}) \quad (4.38)$$

2. Constraints on control trajectories

As discussed in section 4.5.4, a number of constraints are required to ensure that the setpoint and product flow trajectories have the correct form as shown in figure 4.7. Recall that this time discretisation allows any primal production schedule, such that the reactor setpoint trajectory can take only four values during a cycle. The setpoint must remain constant across consecutive transition and production periods. To achieve this, the sixteen breakpoint values for the reactor setpoint are grouped into four sets, each containing four consecutive breakpoint values. Within each set the breakpoint values are constrained such that their values must be identical, as follows:

$$\begin{aligned} C_A^{sp^1} &= C_A^{sp^2} = C_A^{sp^3} = C_A^{sp^4} \\ C_A^{sp^5} &= C_A^{sp^6} = C_A^{sp^7} = C_A^{sp^8} \\ C_A^{sp^9} &= C_A^{sp^{10}} = C_A^{sp^{11}} = C_A^{sp^{12}} \\ C_A^{sp^{13}} &= C_A^{sp^{14}} = C_A^{sp^{15}} = C_A^{sp^{16}} \end{aligned} \quad (4.39)$$

where the superscript sp^j refers to the j th breakpoint value of the control.

The following equality constraints ensure constant product flows during production periods:

$$\left. \begin{aligned} p_i^3 &= p_i^4 \\ p_i^7 &= p_i^8 \\ p_i^{11} &= p_i^{12} \\ p_i^{15} &= p_i^{16} \end{aligned} \right\} \quad i = 1 \text{ to } 4$$

where the subscript i refers to the i th storage feed flow and the superscript j refers to the j th breakpoint value. The assumption is made that during a transition period the reactor product will be off-spec. A simplification to the problem may then be made by setting all flows to the storage vessels to zero during these periods.

A final equality constraint synchronises the final time of the cycle with the sum of the discretisation intervals:

$$t_{cycle} = \sum_{j=1}^{16} \Delta t_j \quad (4.40)$$

where Δt_j is the j th discretisation interval length.

It should be noted that in the case of a multiproduct process the operating cycle is repeated continuously. For this reason, no steady-state endpoint constraints on the cycle are required as is the case with the other design examples presented in this thesis.

4.5.6 Implementation

4.5.6.1 FORTRAN model generation

The gPROMS modelling environment was used to develop a model of the process based on the units described in section 4.4.2. In addition to the underlying equations describing specific process units in the model, a number of ancillary equations were included. The additional equations provide information to be used by the optimiser in calculating values for the objective function and constraints.

The path constraints required for the optimisation, which were described in the optimisation problem formulation are also added in the form described earlier in section 3.5.3. Following implementation and testing, FORTRAN code for the model (in DASOLV format) was generated using the gPROMS “CODE” facility.

4.5.7 Optimisation results

Solution of the dynamic optimisation problem described in this chapter required 562 CP minutes on a Sparcstation10. The values obtained for the design parameters in the reactor and its control systems are shown in table 4.2. Figure 4.8 shows the profiles over time of the reactor concentration and product storage volumes. Figure 4.9 shows the variation of the four product compositions over time. An economic cost breakdown for the process is shown in table 4.3.

<i>Parameter</i>	<i>Value</i>
K_c^T	$-192.05\text{ }^\circ F/(lbmol.ft^{-3})$
τ_I	$3.66 \times 10^{-2} hr$
K_c^C	$-23.78 ft^3/(hr.^\circ F)$
F_{in}	$42.41 ft^3.hr^{-1}$
τ_R	$1.23 hr$

Table 4.2: Reactor parameters for economically optimal multiproduct plant

The optimum annualised cost for the system was found to be \$ **173167**. The cycle time for the schedule was 14.4 *hours*. The product schedule chosen, shown in figure 4.8, was $B_2 \rightarrow B_4 \rightarrow B_3 \rightarrow B_1$. Note that this schedule was excluded by Cantrell and Luyben based on the heuristic that it will, “...require moving back and forth through the same operating region more than once.” The schedule chosen does, however, avoid the large jump from operating point B_4 to B_1 required by the Cantrell and Luyben schedule.

	<i>Annualised Cost (\$)</i>
Product transitions	23991
Reactor capital	19289
Storage capital	129887
TOTAL	173167

Table 4.3: Cost breakdown for optimal multiproduct plant design

Unfortunately, exact numerical comparisons of the solution obtained with that found by Cantrell and Luyben are not possible as the data provided in their paper appears incomplete. Repeated attempts to obtain a missing molar density from the authors were unfruitful. A value of $3.0 \text{ lbmol.ft}^{-3}$ was used which enabled reasonable agreement between simulation results given in the Cantrell and Luyben paper with those for the model used here. Assuming the molar density to be approximately correct, the cost obtained through optimisation is some 30% lower than that found by the sequential procedure of Cantrell and Luyben.

The annualised cost found is dominated by the cost of product storage which contributes some 75% to the total. The weighting of the objective function largely explains the form of the solution obtained. Recall that low storage costs are favoured by short cycles. This has been achieved in this case by using a reactor with a small residence time which can make rapid transitions between products. As a result of this, the responses obtained are highly oscillatory. Indeed, in the case of production of product B_3 , the reactor response is both oscillatory and unstable. This raises a very important qualitative contrast in the solutions obtained by the two methods. Recall from the optimisation formulation presented earlier that the product purity constraints are applied on the products flowing from the storage tanks. The storage tanks provide an attenuation of the reactor response, so even though the reactor composition oscillates and violates composition bounds as shown in figure 4.8, this attenuation ensures that the composition of the actual products, shown in figure 4.9, remain within their specified bounds. Applying the composition bounds to the *reactor* product in Cantrell and Luyben's approach is highly conservative as it requires a reactor with much larger residence time (6 *hours* *vs.* 1.23 *hours*) which makes smooth transitions to the setpoint. In their form of the problem, charging of a given product storage tank may only begin when the reactor product reaches *and* remains within the specified purity bounds. In the optimisation solution, the

minimum value of stored product for all four products (which occurs when each of tanks is about to be re-charged from the reactor) is somewhat greater than zero. Although it might initially seem that this would lead to larger storage volumes, much shorter transition times are possible because of the additional attenuation of deviations in the reactor composition. The overall result of this is to reduce the total operating cycle length and enable smaller storage tanks to be used. Notice that the largest minimum storage values correspond to products B_1 and B_2 . This is consistent since the reactor composition response exhibits the greatest oscillation about the setpoint for these products. The responses of the product compositions in figure 4.9 clearly show the times at which material of the required specification from the reactor is fed to the corresponding storage tank. Interestingly, a number of linear changes in product composition can also be seen. These are caused by the addition of small amounts of reactor product ($\approx 0.5\%$ of the total flow) to the tank to correct the state profile such that it is cyclic.




Figure 4.8: Cyclic profiles for reactor composition and product storage in multiproduct plant




Figure 4.9: Cyclic profiles for product compositions in multiproduct plant

Results obtained from the optimisation have been presented and discussed. The final section of this chapter considers extensions both to the optimisation framework and to its objectives.

4.6 Discussion

4.6.1 Comparison of methods

The optimisation formulation for cyclic operation developed in section 4.5.1 has been successfully applied. A number of key qualitative differences between the solution obtained by optimisation and by Cantrell and Luyben were discussed. The most important of these was that consideration only of the reactor dynamics and subsequent application of purity constraints on the reactor product is highly conservative. In this instance, the production phase for a given product can only begin when the reactor concentration reaches and remains within the specifications given for that product. By comparing the responses of reactors with different residence times, Cantrell and Luyben determined that a larger reactor which made smooth transitions to the setpoint was most suitable. Imposing this form on the reactor response would seem to greatly reduce the scope for further optimisation of the design. A problem formulation which simultaneously considered both the storage and reactor dynamics has enabled economically superior solutions to be found. Consideration and exploitation of the many more degrees of freedom resulting from such a formulation was made possible through the use of dynamic optimisation.

The computational load of the optimisation based method is not excessive when compared to the Cantrell and Luyben procedure. Each pass of their procedure requires $n(n - 1)$ simulations to be carried out. Each of these simulation results needs to be inspected manually to determine transition times. Even for the small number of optimisation variables considered (residence time and production rate) the manual effort required to locate the “optimum” is very substantial. The optimisation uses the same dynamic model as the procedural approach and yet is also able to exploit many more degrees of freedom. As the number of optimisation parameters and model complexity increases, it is difficult to see how the Cantrell and Luyben procedure could still be successfully and efficiently applied.

Many extensions to the original optimisation framework can easily be accommodated, such as:

- gain scheduling of the controllers to further improve performance,

- time-varying feed and product profiles.

4.6.2 Consideration of secondary objectives

The results presented in section 4.5.7 are a feasible and economic solution to the problem posed and fulfil in all ways the objectives and constraints applied. Recall that the responses obtained were oscillatory and in the case of product B_3 , unstable. Suppose now that in the light of the initial result obtained, a more highly damped reactor response is required for implementation on the actual plant. Achieving such a response will have an associated impact on the overall economics, since the solution found is an economic optimum. The plant manager is prepared to suffer a fraction ϵ increase in the economics to get a better damped plant response. A second optimisation problem may now be formed with a new objective:

$$\min_{u(t),d} \{J_1\} \quad (4.41)$$

$$s.t. \quad (C_{capital} + C_{transition}) \leq (1 + \epsilon) J_0 \quad (4.42)$$

where, J_1 is a new objective which is some measure of the damping in the reactor response and J_0 is the minimum cost solution. $C_{capital}$ and $C_{transition}$ are defined by equations 4.30 and 4.32 respectively.

Solution of this problem for an ϵ of 0.05, *i.e.* 5% increase in cost, with J_1 as the integral square error of the reactor response *summed over the production phases*, gives the response shown in figure 4.10. Although the response is still oscillatory, it is damped and has an ISE which is a factor of ten smaller than that for the lowest cost system. Examination of the optimisation parameters used reveals that this results from a 33% increase in the reactor residence time, reduction in both controller gains and an increase in the reset time of the composition loop.




Figure 4.10: Damped reactor composition response

Chapter 5

Integrated design of wastewater treatment systems

5.1 Introduction

5.1.1 Rationale

The Urban Wastewater Treatment directive of 1991 requires that, by the year 2000, towns with a population in excess of fifteen thousand must have wastewater treated to a specified minimum standard. The general perception of wastewater treatment facilities is that they are odorous, large and intrusive. Construction of the new treatment facilities necessary to meet the 1991 directive, especially for municipalities and industries which are situated in sensitive areas (such as coastal regions), is therefore considered by the public to be undesirable and unacceptable. Pressures to build plants on constricted, sensitive sites under increasingly stringent environmental law, have recently led to the development of compact (or small foot-printed) plants. The provision of small plants which provide localised treatment of wastewater is an example of distributed chemical manufacturing - although in this case the product is “clean” water.

The proceedings from a conference on “Advances in Water and Effluent Treatment” held in 1993 (White, 1993) dedicated an entire chapter to progress in the area of compact treatment plants. Much of this work focuses on developing efficient new technologies which enable a reduction in equipment sizes. However, a paper presented by Hannon and King discusses the concept of *process intensification* as a means of reducing equipment volumes. They state that the main driver for process intensification is reduction in plant capital cost possible through “intensifying” the plant. A number of ways of achieving this intensification are also discussed and are listed here:

- process auditing,
- equipment intensification - replacing an old process design with a new one; incorporating improved equipment,
- intensification of the chemical operating conditions (*e.g.* using oxygen instead of air in an activated sludge process),
- mechanical process intensification - using a single equipment item for more than one purpose.

This case study considers the design and operation of chemical wastewater treatment facilities, specifically for the neutralisation of an acid wastewater. Typically, treatment of chemical wastewater is carried out on a continuous basis in a series of stirred tank reactors. Some buffering capacity is often included either to damp out oscillations in the effluent from the treatment system, or to provide local containment facilities (diversion system) in the event of process abnormalities arising (*e.g.* equipment failure). The aim of this work is to:

determine potential economic benefits from integrating the diversion system design and operation with that of the continuously operating treatment system.

This may be considered as an intensification of the process in two ways:

- by improving the overall process design,
- by using the diversion system as an additional degree of freedom during normal operation, in addition to it being a fail-safe mechanism.

5.1.2 Chapter outline

After introducing the specific design problem to be tackled in section 5.2, section 5.3 details the model of the system which is used. This includes a background to the system chemistry together with approximations and assumptions made as well as issues of numerical robustness. A design procedure is then discussed in section 5.4 which enables rapid screening of potential treatment system structures to obtain a candidate structure. Section 5.4.2 presents the specific form of the dynamic optimisation problem which is used to determine the optimal design and cyclic operation of the process. An economic bounding property is developed in section 5.6.1 which enables the solution search space to be reduced. Optimisation results for two wastewater plants with integrated diversion systems are then presented and a discussion given.

5.2 Problem definition

An acidic wastewater stream from an upstream process must be neutralised before discharge to the environment. Environmental regulations on the pH of any material which leaves the treatment facility require:

$$5.0 \leq pH \leq 9.0 \quad (5.1)$$

The wastewater to be treated has a flowrate of $3 \text{ m}^3.\text{min}^{-1}$ and has a concentration of 0.20 M acid. The wastewater is subject to concentration load changes. These load changes are in the form of step changes up to a maximum loading of 0.25 M acid. The duration of the load change is of variable length, but typically lasts in excess of 30 minutes, after which time the load returns to the nominal value. It is assumed in this case study that the response of the process is monotonic with respect to the input loading. In other words, if the process is feasible for the worst-case load, it will be feasible for all loads between the worst-case and nominal loads. The problem definition will therefore assume a worst-case concentration load of 0.25 M acid.

A wastewater treatment facility must be designed which is capable of meeting the discharge limits on the effluent in the face of disturbances in the concentration load described above. There is limited space available on site for the construction of the new treatment facility and a tight constraint has therefore been imposed on the total volume of the system:

$$V_{total} \leq 35.0 \text{ m}^3 \quad (5.2)$$

As discussed earlier, treatment of chemical wastewater is typically carried out on a continuous basis in a series of stirred tank reactors. Some buffering capacity may be included, either to damp out oscillations in the effluent from the treatment system, or to provide local containment facilities (diversion system) in the event of process abnormalities arising (*e.g.* equipment failure). A process consisting of a chemical treatment system and diversion system will therefore form the basis of this design problem. The general design problem to be considered is shown diagrammatically in figure 5.1.

The key feature of the design procedure presented here is that the design and operation of the diversion system will be treated as an integral part of the overall problem. This is in contrast to conventional practice in which the wastewater treatment system is designed and a diversion system is added only as a fail safe for the main treatment process.

Referring now to figure 5.1, the overall design problem specification (*i.e.* feed specification, constraints and disturbance description) is represented by the outer box

(shown with dashed lines). The two inner boxes show the treatment system and diversion system, which are the design degrees of freedom by which the specification is to be met. The specific details of each of these two systems are as follows:

1. Treatment system.

The treatment system consists of a number of continuous stirred tank reactors (CSTRs). The number of reactors used essentially depends on the required degree of attenuation between the input and discharge pH . Each CSTR has a composition control system comprising a proportional-integral controller which manipulates an alkali reagent stream. The reagent available for neutralisation of the acid waste is a 6 molar solution of sodium hydroxide. The control response should recover to near-steady-state within finite time. To ensure adequate mixing, the minimum residence time allowable in each CSTR is 1 minute (Walsh, 1993).

2. Diversion system.

The treated effluent from the final CSTR in the treatment system is either discharged to the environment or diverted to a diversion storage tank depending on its pH . Material held in the diversion tank can be returned to any of the CSTRs in the treatment system for further processing. The most severe deviations in reactor concentration typically occur at the rising and falling edges of a concentration load change. A traditional wastewater treatment process which discharges “on-spec” effluent continuously must be capable of treating this large deviation in the concentration. The inclusion of an integrated diversion system gives the potential for storage of off-spec effluent which can then be bled back into the treatment system when the initial impact of the load disturbance has passed.

The aim of the design process is therefore to determine the cheapest wastewater plant design which is capable of meeting effluent discharge regulations, in the face of disturbances in the concentration load, while satisfying the stringent limit on the plant’s maximum volume. The novel approach in solving this problem is to include, as a degree of freedom, the design and operation of the diversion system in parallel with the continuous wastewater treatment system.




Figure 5.1: Wastewater plant design problem

5.3 Modelling the problem

Having outlined the definition of the problem, this section discusses the system models. Recall that the process will consist of a number of continuous stirred tank reactors (CSTRs) in which the wastewater is treated, as well as a diversion system. In the following sections, the modelling of the CSTRs and diversion system will be considered in turn. In the discussion given, relevant background information regarding the system chemistry will be briefly outlined, modelling approximations and assumptions will be detailed and issues of model robustness will be covered.

5.3.1 pH modelling of strong-acid/strong-base systems

As discussed earlier in section 5.2, for the process design to be feasible it must discharge effluent which meets strict limits on the pH of any effluent which is discharged (namely $5 \leq pH \leq 9$). Since the model of the wastewater treatment system will be in terms of concentrations, a means of modelling the relationship between pH and concentration is clearly necessary.

5.3.1.1 Definitions

Before detailing the pH -concentration model, some relevant terminology must be defined which is used both here and in the process model.

1. Strong-acids and strong-bases.

The terms strong-acid and strong-base imply that the H^+ and OH^- groups completely dissociate from the other components of the acid or base.

2. Normals.

The concentrations used are in normals (N). A *normal* concentration is defined as available OH^- (M), where M is *molar* concentration (gmoles/litre). So, for example, a 1M acid solution can be described as a -1N solution.

3. Net concentration.

The net concentration c_{net} is used throughout this case study and is defined as the net concentration of acids and bases in normals:

$$c_{net} = [OH^-] - [H^+] \quad (5.3)$$

where $[H^+]$ and $[OH^-]$ are hydrogen and hydroxide ion concentrations respectively.

4. pH .

The pH of a solution is given by the theoretical relationship:

$$pH = -\log_{10}([H^+]) \quad (5.4)$$

5.3.1.2 pH - concentration relationship

For a strong-acid/strong-base system such as that considered in this case study, a simplified thermodynamic model of the electrolyte solution yields the equilibrium equations:

$$\begin{aligned} cnet &= [OH^-] - [H^+] \\ pH &= -\log_{10}([H^+]) \\ [OH^-][H^+] &= k_w \approx 10^{-14} \end{aligned} \quad (5.5)$$

where $cnet$ and pH are as defined previously and k_w is the dissociation constant of water. Explicit solution of equations 5.5 gives the relationship:

$$pH = \begin{cases} -\log_{10} \left(\frac{-cnet + \sqrt{cnet^2 + 4k_w}}{2} \right) & cnet < 10^{-4} \\ -\log_{10} \left(\frac{k_w}{cnet} \right) & cnet \geq 10^{-4} \end{cases} \quad (5.5a)$$

The conditional in equation 5.5a is used to give a numerically robust expression without significant discontinuity (Walsh, 1993).

5.3.2 CSTR model

Neutralisation of the incoming wastewater (which is assumed to be an instantaneous reaction) is achieved by treatment in a series of one or more continuous stirred tank reactors (CSTRs). The general CSTR unit to be used in this process incorporates two inlets: one for connection to the preceding upstream unit, the other to receive material from the diversion system. A single exit stream is provided. The tank concentration is controlled through the use of a proportional integral (PI) composition controller, which manipulates the level of alkali reagent (NaOH) addition. The CSTR is assumed to have a constant volume.

Having outlined the general form of the CSTR unit to be modelled, more specific modelling issues will now be addressed beginning with the type of mixing model to be used.

5.3.2.1 Mixing model

In this work the treatment vessels will be modelled using the “fractional tubularity” (delay-lag) model. Walsh (1993) has observed that “this is the most commonly used model of a mixed vessel” and showed it to give good agreement with more rigorous computational fluid dynamics models. In the fractional tubularity model, part of the vessel is taken as exhibiting plug-flow conditions which contributes a delay ($t_{d_{mix}}$) with the remainder being perfectly mixed and contributing a first order lag ($t_{c_{mix}}$).

5.3.2.1.1 Choice of $t_{d_{mix}}$. A constant value for $t_{d_{mix}}$ of 10 seconds is implemented in the CSTR models used for this case study. This is considered by Walsh to be a realistic design specification in the absence of detailed mixing analysis.

5.3.2.1.2 Model approximation for $t_{d_{mix}}$. The deadtime $t_{d_{mix}}$ can be approximated by a number of first order lags in series. The models implemented here use a total of twenty lag blocks in series to give a good approximation to the deadtime, whilst avoiding excessive increases in the model size.

5.3.2.2 CSTR dynamic model

The CSTR unit to be modelled is shown in figure 5.2. The figure incorporates the mixing model described above in diagrammatic form. The CSTR model may be described by the following equations:

1. Overall material balance,

$$F_{out} = F_{in} + F_d + F_r \quad (5.6)$$

where F_{in} is the feed flow, F_r is the reagent flow, F_d is the flow from the diversion system and F_{out} is the output flow. Note that all flowrates used in the case study are $m^3.min^{-1}$.

2. Concentration material balance (perfectly mixed section) ,

$$V_{c_{mix}} \frac{dC_{c_{mix}}}{dt} = C_{in}F_{in} + C_dF_d + C_rF_r - C_{c_{mix}}F_{out} \quad (5.7)$$

where $V_{c_{mix}}$ is the volume of the perfectly mixed section, $C_{c_{mix}}$ is the concentration in the perfectly mixed section, C_{in} , C_d and C_r are the concentrations of the feed, diversion tank and alkali reagent respectively.




Figure 5.2: CSTR unit for wastewater treatment

3. Deadtime section

$$\begin{aligned}
 \left(\frac{t_d}{n}\right) \frac{dC_{dt,1}}{dt} &= (C_{mix} - C_{dt,1}) \\
 \left(\frac{t_d}{n}\right) \frac{dC_{dt,i}}{dt} &= (C_{dt,i-1} - C_{dt,i}) \quad i = 2, \dots, n-1 \\
 \left(\frac{t_d}{n}\right) \frac{dC_{out}}{dt} &= (C_{dt,n-1} - C_{out})
 \end{aligned} \tag{5.8}$$

where n is the number of first order lags used in the deadtime approximation, t_d is the total deadtime, $C_{dt,i}$ is the concentration in the i th lag block and C_{out} is the concentration of product leaving the last lag block (*i.e.* the CSTR effluent concentration).

5.3.3 Controller model

As stated earlier, control of the CSTR composition is achieved through a PI composition controller which manipulates the amount of alkali reagent added to the reactor. A model of the PI controller which is numerically robust during optimisation must:

- incorporate a mechanism to prevent integral wind-up,

- clip the controller output to the desired bounds.

In achieving this it is important to avoid introducing significant discontinuities into the model. These can cause numerical problems during integration of the model, such as failures when evaluating sensitivities or cycling across the discontinuity. Walsh (1993) suggests functions with continuous first derivatives that produce smooth integral desaturation and output clipping. These are incorporated into the PI model given below:

$$\xi_C = C^{sp} - C_{out} \quad (5.9)$$

$$\tau_I \frac{d\xi_I}{dt} = K_{clip} K_c^C \xi_C \quad (5.10)$$

$$F_r^c = K_c \xi_C + \xi_I \quad (5.11)$$

where ξ_C is the composition controller error, C^{sp} is the composition setpoint, ξ_I is the integral error, τ_I is the controller reset time, K_c^C is the controller gain and F_r^c is the calculated reagent flowrate.

In the equations above the value of K_{clip} is determined using the following conditionals:

$$K_{clip} = \begin{cases} 1 & \text{if } \epsilon < F_r^s < 1 - \epsilon \\ \frac{3}{\epsilon^2}(F_r^s)^2 - \frac{2}{\epsilon^3}(F_r^s)^3 & \text{if } 0 \leq F_r^s \leq \epsilon \\ \frac{3}{\epsilon^2}(1 - F_r^s)^2 - \frac{2}{\epsilon^3}(1 - F_r^s)^3 & \text{if } 1 - \epsilon \leq F_r^s \leq 1 \\ 0 & \text{otherwise} \end{cases}$$

where ϵ is typically between 0.01 and 0.05, F_r^s is the scaled controller output:

$$F_r^s = \frac{F_r^c}{F_r^{range}} \quad (5.12)$$

and F_r^{range} is the reagent flow range. The reagent flow to the CSTR is therefore given by:

$$F_r = \begin{cases} F_r^{range} F_r^s K_{clip} & \text{if } F_r^s < 0.5 \\ F_r^{range} (1 - (1 - F_r^s) K_{clip}) & \text{otherwise} \end{cases} \quad (5.13)$$

5.3.4 Diversion system model

As stated in the problem definition, all effluent leaving the process must lie in a pH range of 5 to 9. In a traditional wastewater treatment process flow through the plant is continuous. This discharge requirement must therefore be met by ensuring that the effluent is adequately neutralised within the CSTRs. Typically, the treatment plant will also include diversion facilities which provide containment for wastewater in the event of equipment failure in the treatment plant. This case study considers the

integration of diversion facilities with the chemical treatment tanks. The availability of a diversion system as an additional degree of freedom enables discharge constraints to be met by diverting off-specification material to the diversion system. The diversion system essentially consists of mechanism to divert the treatment plant effluent, a mixed tank to hold diverted material and connections from the diversion tank to the treatment CSTRs.

5.3.4.1 Diversion mechanism

This section describes the modelling of the mechanism by which the off-spec material is diverted.

A set of conditional equations is used to force any effluent from the wastewater treatment process which falls outside the specified pH limits to be diverted into the diversion tank. As discussed earlier in section 5.3.3, it is important to avoid introducing unnecessary discontinuities into the model of the system. For this reason, a set of smoothed conditional equations with continuous first derivatives with respect to time are used. Figure 5.3 shows the form of the switching function. If the outlet pH (calculated using equation 5.6) lies between 5.1 and 8.9 pH the flow to diversion is zero. Outside these limits, the flow smoothly increases until at pH 5 and below, or pH 9 and above, the flow to the diversion system equals the exit flow from the treatment system. The complete model for the diversion mechanism is as follows:

$$F_{offspec} = \begin{cases} 0.0 & \text{if } 5.1 < pH < 8.9 \\ F_{out} [300(5.1 - pH)^2 - 2000(5.1 - pH)^3] & \text{if } 5.0 \leq pH \leq 5.1 \\ F_{out} [300(pH - 8.9)^2 - 2000(pH - 8.9)^3] & \text{if } 8.9 \leq pH \leq 9.0 \\ F_{out} & \text{otherwise} \end{cases}$$

where F_{out} is the exit flow from the treatment process and $F_{offspec}$ is the flow of material to the diversion system. A material balance around the diversion unit defines the flow of material actually discharged by the plant:

$$F_{discharge} = F_{out} - F_{offspec} \quad (5.14)$$

where $F_{discharge}$ is the discharge flow rate.

5.3.4.2 Diversion tank model

The final model required is that of the diversion tank itself. This tank is assumed to be perfectly mixed, is fed by the flow $F_{offspec}$ and has n product streams, one for each of




Figure 5.3: Functional form of diverted off-spec flow

the CSTRs in the treatment system. The dynamics of the diversion system are described by overall and component material balances:

$$\begin{aligned} \frac{dV_{diversion}}{dt} &= F_{offspec} - \sum_{i=1}^n F_{d_i} \\ V_{diversion} \frac{dC_d}{dt} + C_d \frac{dV_{diversion}}{dt} &= C_{out} F_{offspec} - C_d \sum_{i=1}^n F_{d_i} \end{aligned} \quad (5.15)$$

where, F_{d_i} is the diversion flow to CSTR i in the treatment system, $V_{diversion}$ is the volume of the diversion tank and C_{out} is the concentration of the effluent leaving the treatment plant.

In order to ensure that the diversion tank is numerically robust during the optimisation, a proportionally controlled make-up stream is added to the dynamic model to maintain the level in the vessel above a pre-specified lower limit. A detailed discussion on this subject and the form of the equations implemented can be found in section 4.4.2.3.1. A path constraint on the tank holdup ensures that the proportional controller does not operate at the solution point (see 5.4.6.2.1).

5.4 Design procedure

The design problem to be solved has been defined earlier in section 5.2. Recall that the complete wastewater plant may be considered in two parts,

1. a continuous chemical treatment system,

2. a diversion system,

and that in traditional systems it is the former that is the focus of the design problem. Once the reagent to be used for treatment has been selected, a typical wastewater treatment system design problem is essentially,

- determine the number of treatment tanks,
- determine the required sizes of these tanks,
- determine the appropriate PI settings of composition controllers,

to ensure feasible operation of the plant in the face of variable concentration loading. The solution of such problems has been considered extensively and a number of procedures to aid in the design of wastewater treatment systems exist.

This work investigates potential benefits to be gained from integrating the operation of the diversion system with a continuous wastewater treatment system as described above. The new design problem must therefore integrate the original design problem described above (number and size of CSTRs, PI settings) with an optimisation of the design and operation of the diversion system.

As discussed earlier, the general design procedure used in this thesis aims to solve the combined structural and operational problem by decomposing the design problem into two stages:

1. develop a candidate process structure,
2. apply a dynamic optimisation formulation to determine the corresponding optimal operating cycle and design for that candidate.

By decomposing the problem into two levels, potentially expensive dynamic optimisation techniques are applied only when a viable process structure is in place. Given that the structural decisions to be made relate to the wastewater treatment system (*i.e.* number of CSTRs), the development of a candidate structure should be focussed initially on finding a feasible structure for the treatment system alone. The design procedure presented in the following section (which is based on these two stages) uses established design heuristics to develop a feasible base case candidate structure. The optimal cyclic operation formulation to be applied at the second stage of the case study is then introduced and its elements are discussed in detail. Recall that the use of an operating cycle provides a means of incorporating the process operation into the overall design problem. Cyclic optimisation of the candidate structure is then carried out to provide an economic upper bound on the

cost of the process. In order to avoid the enumeration of many further structures using the dynamic optimisation formulation, an economic bounding property is introduced which enables much of the solution search space to be eliminated.

5.4.1 Developing a candidate structure

This section considers the development of a candidate structure to carry forward to the second stage dynamic optimisation.

Recall that in section 5.4 the “traditional” wastewater design problem was discussed and the existence of design guidelines for this type of problem was mentioned. One such guideline which enables rapid determination of the economically optimal number and sizing of CSTRs in series has been developed by Walsh (1993) based on his observations using optimally tuned PI feedback control. Given a treatment system with n equally sized CSTRs in series (the industrial norm), the optimal volume of each CSTR is given by:

$$V_{opt} \approx F_T t'_d \left(\frac{1.5}{\delta_c} \right)^{\frac{1}{n}} \quad (5.16)$$

where F_T is the total flow to the process, t'_d is the delay between a disturbance reaching the output and the controller response reaching the output (see section 2.3) and δ_c is the required disturbance attenuation, *i.e.* the allowable outlet concentration variation divided by the inlet concentration change.

Equation 5.16 may be used to efficiently and quickly search for the economically optimal structure of the wastewater treatment system. The feasibility of this structure can then be verified by dynamic optimisation and provides a sensible base case on which to consider the integration of a diversion system. Applying equation 5.16 to the problem specification given earlier in section 5.2 gives the results shown in table 5.1. The costs presented (in pounds) are determined from the correlation,

$$C_{CSTR} = 20000 + 2000V_{opt}^{0.7} \quad (5.17)$$

which will be discussed further in section 5.4.3. The cheapest structure from table 5.1

n	<i>CSTR volume</i>	Total cost
2	43.3	0.96×10^5
3	9.8	0.90×10^5
4	4.7	1.00×10^5

Table 5.1: Optimal treatment tank sizes and costs using Walsh approximation

corresponds to a 3 CSTR system. Note that the system with 2 CSTRs violates the constraint on total volume of $35 m^3$.

An “optimal” structure for the wastewater treatment system has been established using equation 5.16. This 3 CSTR structure is now taken as the candidate to be carried forward. In section 5.5 the feasibility of this structure will be validated through the use of a dynamic optimisation formulation for cyclic operation. Optimisation will also enable the design of this “base case” system to be fully optimised. The form of the dynamic optimisation problem is presented in the following section.

5.4.2 Formulating the optimal cyclic operation problem

Previous sections in this chapter have dealt with the specification of the design problem, modelling of the unit operations forming the process and have discussed the design procedure to be used. This section discusses the way in which the design problem defined in section 5.2 may be solved using a dynamic optimisation formulation for cyclic operation. Figure 5.4 shows how the concentration load disturbance description maps to




Figure 5.4: Diagram showing cyclic behaviour of concentration disturbance (top) and the collapsed cycle which is implemented (bottom).

the cyclic operation framework. Also shown is a typical response of the controlled CSTR concentration to the load change. The plant is operating at steady state at time t_0 at which time the inlet concentration load steps from $-0.2 N$ to $-0.25 N$ (the worst-case

disturbance). This new load remains in effect until some arbitrary time t_2 . At t_1 the process has recovered to steady-state. The concentration load returns to its original value at time t_2 and by time t_3 the process has again recovered from this disturbance and is steady. At time t_{cycle} , after operating for an arbitrary period at the new steady-state, the cycle is repeated. The operating cycle described above can effectively be collapsed as shown in figure 6.2 to avoid unnecessary and time consuming integration through the steady state period during solution. This is achieved simply by switching the concentration load back to the original level as soon as the process has recovered to steady-state. The two cycles are identical for the purposes of determining the optimal cyclic operation and process design.

As previously discussed, the wastewater treatment process to be designed comprises a series of generic CSTR units (in which chemical treatment of the wastewater occurs) and a diversion system. The cyclic optimisation problem may therefore be posed for a general system consisting of n CSTRs and a diversion system. This problem formulation is developed in the following section.

5.4.3 Objective function

The objective function to be used in this case study is the minimisation of the annualised cost of the process. This cost is made up of:

- cost of sodium hydroxide reagent,
- capital cost of stirred tank reactors,
- capital cost of diversion tank,
- capital and operating cost of sodium hydroxide feed system,
- capital and operating cost of pH control system.

Walsh (1993) provides correlations for these costs based upon costing graphs presented by Cushnie (1984). However, a number of these costs may be excluded from the optimisation objective function.

- The output concentration of the treated wastewater is tightly bounded. The total reagent usage will therefore be approximately constant for all systems meeting the design specifications.

- The cost (operating and capital) of the reagent addition system is a function of the total reagent feed capacity. Similarly to the cost of reagent, this will be approximately constant for all designs meeting the specifications.
- The pH control system cost (operating and capital) is given as a constant fixed cost.

In other words, only the capital costs of the stirred tank reactors and diversion tank need be considered in the economic objective function for this optimisation. The costs of these units are functions of the CSTR volumes and maximum diversion tank volume:

$$\begin{aligned} C_{CSTR} &= 20000 + 2000 V_{CSTR}^{0.7} \\ C_{diversion} &= 20000 + 2000 (V_{diversion}^{max})^{0.7} \end{aligned} \quad (5.18)$$

The annualised capital cost for a process containing n CSTRs is therefore:

$$C_{capital} = \frac{1}{3} \left\{ \sum_{i=1}^n C_{CSTR}^i + C_{diversion} \right\} \quad (5.19)$$

where the $\frac{1}{3}$ pre-multiplier is a capital charge factor which approximates “*the time value of money*” (Douglas, 1988), and converts the total capital cost to an annualised equivalent. The objective function for the optimisation is therefore given by:

$$\min_{u(t), d} \{C_{capital}\} \quad (5.20)$$

5.4.4 Plant parameters for optimisation

Given a fixed structure containing n CSTRs in the wastewater treatment system, the aim of the optimisation problem described here is to find the best combination of continuous treatment system design and cyclic operation of the diversion system. The degrees of freedom to be optimised in each of these systems are now considered in turn.

5.4.4.1 Treatment system optimisation parameters

Each CSTR included in the treatment system has associated with it a number of degrees of freedom. These design parameters are all time invariant (*i.e.* fixed over time) and are as follows:

1. PI composition controller gain,
2. PI composition controller reset time,
3. Composition setpoint,
4. CSTR volume.

5.4.4.2 Diversion system optimisation parameters

Integration of the diversion system with the treatment system is achieved by means of a number of process streams which enable the material held in the diversion tank to be bled back into the treatment tanks for re-processing. The flows in each of these return streams during the cycle will therefore be determined by optimisation. Since there are n CSTRs there will correspondingly be n such return streams. The form of these controls will be considered further in section 5.4.5. The diversion system is described by two state variables, volume ($V_{diversion}$) and concentration (C_d). In the cyclic operation framework proposed the initial values for both these states are included as optimisation parameters. Finally, the maximum size of the diversion tank is added as an optimisation parameter as this appears in the objective function.

It should be noted that since the system is at steady-state at time t_0 , the initial state of the treatment system is defined by the controller setpoints.

5.4.5 Time discretisation and control parameterisation

Recall from chapter 3 that a feasible path control parameterisation method will be used to solve the dynamic optimisation problem. This section considers the choice of both the functional form of the controls and the manner in which the time horizon is discretised.

5.4.5.1 Control parameterisation

As discussed in section 3.5.3.1 piecewise linear controls will be implemented as these give the highest precision whilst still ensuring the control variables remain correctly bounded using just simple bounds on the break points.

5.4.5.2 Time discretisation

The collapsed cycle shown in figure 5.4 can be divided into five periods:

1. steady operation at inlet concentration of $-0.2 N$,
2. load change to $-0.25 N$,
3. recovery to steady-state,
4. load change to $-0.20 N$,
5. recovery to original steady-state.

The time discretisation used for this case study is based upon these five periods. Thirteen time discretisation intervals are used which are allocated as follows:

1. one fixed time interval of 1 *minute* (steady operation),
2. one fixed time interval of 1×10^{-3} *minutes* (first ramp in concentration load),
3. five variable length time intervals (recovery to steady-state),
4. one fixed time interval of 1×10^{-3} *minutes* (second ramp in concentration load),
5. five variable length time intervals (recovery to original steady-state),

The lower bound on the variable time intervals was set at 0.25 minutes, the upper bound was set at 15 minutes.

5.4.6 Constraints

The optimal design and operation of the plant is subject to a number of constraints which arise either from operational requirements on the process or as a result of the optimisation framework itself.

5.4.6.1 Operational constraints

In section 5.2 the operating constraints on the process were defined. These constraints were: strict limits on the *pH* of the effluent, a limit on the maximum total volume of the system and a minimum residence time in each CSTR of one minute. The first of these constraints is met implicitly by the presence of the off-spec diversion mechanism discussed in section 5.3.4.1. The second operational constraint, that the total volume be less than 35 cubic metres, is applied as the following inequality constraint:

$$V_{total} \leq 35.0 \quad (5.21)$$

Finally, the minimum residence time requirement is added as follows:

$$\tau_i \geq 1.0 \quad i = 1, ..n \quad (5.22)$$

where τ_i is the residence time in CSTR *i*.

5.4.6.2 Non-operational constraints

In addition to the operational constraints described above, a number of constraints are required by the optimisation framework, for example to enforce cyclic operation. The non-operational constraints may be either path constraints or equality point constraints.

5.4.6.2.1 Path constraints The maximum holdup in the storage vessels over the operating cycle must be determined so that the capital cost of the storage tanks can be calculated. This is achieved by applying the following path constraint:

$$V_{diversion} \leq V_{diversion}^{max} \quad (5.23)$$

where $V_{diversion}^{max}$ is a time invariant optimisation parameter.

Recall that numerical robustness in the storage tank model was achieved through a proportionally controlled make-up stream (see section 4.4.2.3.1 for discussion). This was added to prevent negative holdup in this tank. At the solution point the make-up stream must not be active and a further path constraint must be added to ensure this. The form of this constraint was discussed previously in section 4.5.5.2.1.

5.4.6.2.2 Equality and inequality point constraints The remaining constraints used in this formulation are equality and inequality point constraints. These are required for two purposes:

- to ensure that steady-state is reached at t_1 and t_f ,
- to enforce the cyclic operation as discussed in chapter 3.

These two sets of constraints will now be discussed in turn.

1. Steady-state constraints

The cyclic operating cycle shown in figure 5.4 shows the operation passing through two steady-state points: one at t_1 and a second at t_f . In order to be feasible, the process must be able to recover to these steady-states after a disturbance has entered the system. Constraints are therefore required to ensure that steady operation is achieved at these points in the cycle.

The trajectories of the states in the diversion tank are determined by optimisation of the open loop trajectories of the control parameters. Strict equality constraints must be applied to these states, as follows, to ensure that their values will remain constant during the steady-state periods in the operating cycle:

$$\left. \begin{array}{l} \dot{C}_d(t) = 0 \\ \dot{V}_{diversion}(t) = 0 \end{array} \right\} t = [t_1, t_f]$$

where a dot indicates the time derivative of a variable.

Because the concentration states in the wastewater treatment CSTRs are controlled in closed loop, the requirement imposed at t_1 and t_f is that the controlled concentration responses should recover to near steady-state. This is sufficient to ensure that they remain bounded and decay towards the exact steady-state. The damped behaviour required is shown in figure 5.5. At the points t_1 and t_f in the cycle, the maximum amplitude of the concentration response must lie within the limits $\pm A^{max}$. One possible way of enforcing this behaviour would be to apply two




Figure 5.5: Required damped concentration response

time-varying path constraints, depicted by the bold solid lines in figure 5.5. The implementation of such a constraint is not trivial since the locations of t_1 and t_f are not fixed. The required response may be implemented in a more straightforward manner using point constraints at t_1 and t_f . The closed loop response of the controller error in CSTR i (ξ_{C_i}) will approximate a decaying sinusoidal signal with period T_p and amplitude A_i at time t_1 or t_f . The period is approximately given by $4t_d$, where t_d is the CSTR mixing delay (specified as 10 seconds). The requirement to be imposed in CSTR i is therefore that:

$$|A_i| \leq |A_i^{max}| \quad i = 1, \dots, n \quad (5.24)$$

The sample time on the sinusoid at which the constraint is applied does not necessarily correspond to the maximum in ξ_{C_i} . It is therefore not sufficient just to apply the constraint:

$$\xi_{C_i} \leq A_i^{max} \quad (5.25)$$

Instead, two constraints must be applied to each CSTR. When both these constraints are simultaneously satisfied the maximum amplitude of the error response is guaranteed to be less than the specified A_i^{max} :

$$\left. \begin{aligned} -\epsilon A_i^{max} &\leq \xi_{C_i}(t) \leq \epsilon A_i^{max} \\ -\omega A_i^{max} &\leq \dot{\xi}_{C_i}(t) \leq \omega A_i^{max} \end{aligned} \right\} \quad i = 1, \dots, n \quad t = [t_1, t_f]$$

where $\dot{\xi}_{C_i}(t)$ is the time derivative of the controller error and ω is the frequency of the sinusoid:

$$\omega = \frac{2\pi}{T_p} \quad (5.26)$$

Note that since T_p is of order 40 seconds, ω will be approximately 10 rad.min^{-1} . In equation 5.26, ϵ is sufficiently small to force the constraint to be applied near the point of zero error since at this point the derivative of ξ_{C_i} is close to its maximum value (ωA_i). In the cases presented an ϵ of 0.1 was used.

2. Cyclic constraints

The aim of this optimisation problem is to determine the optimal design and cyclic operation of the wastewater plant. A suitable set of endpoint constraints must be applied at the final time to enforce the cyclic operation of the process. Since the initial values of the states in the diversion tank are optimised, a set of endpoint constraints is required to guarantee that these states are cyclic:

$$\begin{aligned} C_d(t_0) &= C_d(t_f) \\ V_{diversion}(t_0) &= V_{diversion}(t_f) \end{aligned} \quad (5.27)$$

Cyclic operation of the states in the wastewater treatment tanks is implicit because:

- the setpoints in each reactor are time invariant,
- the concentration load profile is specified in the problem as being cyclic,
- the controllers are constrained such that at t_f they are near steady-state.

5.4.7 Implementation

5.4.7.1 FORTRAN model generation

Using the unit models described in section 5.3 a complete model of the candidate process was developed within the gPROMS modelling environment. In addition to the underlying equations describing specific process units in the model, a number of ancillary

equations were included. The additional equations provide information to be used by the optimiser in calculating values for the objective function and constraints.

A set of algebraic variables to be used in the endpoint constraints is defined from the time derivatives of the states as follows:

$$\begin{aligned} C_{d_{dot}} &= k_1 \frac{d}{dt}(C_c) \\ V_{diversion_{dot}} &= k_2 \frac{d}{dt}(V_{diversion}) \end{aligned} \quad (5.28)$$

where k_i is a scaling factor and the subscript $_{dot}$ is used to indicate that the algebraic variable represents a time derivative. By explicitly defining these variables, the corresponding sensitivity values (which are required by the optimisation algorithm) will be calculated by DASOLV during the integration process (see chapter 3).

The two path constraints specified in the optimisation problem (see section 5.4.6.2.1) were added in the form given in section 3.5.3. After validating and testing the model within gPROMS, the system equations were generated in DASOLV format using the gPROMS “CODE” facility.

5.4.7.2 Specification of load disturbance profile

Figure 5.4 in section 5.4.2 shows the concentration load profile for the cyclic operation. This section briefly describes the way in which this profile is incorporated in the optimisation.

The time discretisation to be used for this problem (see section 5.4.5) may be summarised as follows. At time t_0 the plant is operating at steady state at which time the inlet concentration load steps from $-0.2N$ to $-0.25N$. This new load remains in effect until time t_1 at which time the concentration load returns to its original value until t_f when the cycle is complete. To achieve this, the concentration load was set up in the optimisation input file as a time-varying control parameter (see section 3.5.3). The breakpoint values for this control were then defined to give the required profile, and were excluded from the set of available optimisation parameters. In doing this, the desired disturbance profile was superimposed on the operating cycle.

5.5 Optimal cyclic operation of base case treatment system

A candidate wastewater treatment system was developed in section 5.4.1. The feasibility of this structure will now be validated using the dynamic optimisation formulation for cyclic operation presented in section 5.4.2. Optimisation of this candidate will

also enable the design of the CSTR units and their respective control systems to be fully optimised. The resulting system will then form a base-case for the remainder of the design procedure in which an integrated diversion system is considered.

5.5.1 Implementation of the base case

The optimisation framework presented previously is general in that it allows consideration of an n CSTR treatment system plus diversion system. The base case system requires cyclic optimisation of a three CSTR treatment system without diversion facilities. This is achieved in the framework presented as follows:

- all the return flows from the diversion system back to the CSTRs are set to zero and removed from the set of optimisation parameters,
- the initial state values in the diversion system are removed from the set of optimisation parameters,
- the maximum volume parameter in the diversion system is set equal to the initial value of the volume state and is removed as an optimisation parameter.

Application of the above measures removes the diversion system as a degree of freedom in the problem. Because the path constraint on the diversion tank maximum volume (equation 5.23) still applies, the bounding used prohibits any material from entering the diversion tank. In other words the design specifications must be satisfied by the three CSTR treatment system alone. It should be noted that since the diversion mechanism is activated outside pH limits of 5.1 and 8.9, the actual solution found will be slightly conservative.

5.5.2 Solution and results

The optimisation problem was initialised using Cohen-Coon tuning parameters in the PI controllers and the CSTR volumes calculated from equation 5.16. Initial guesses for the setpoints in each CSTR were -1.0, -0.1 and 0.0 respectively. Convergence to the solution was rapid, requiring fifteen optimiser cycles and a processing time of 33 CPU minutes on a Sparcstation10. The optimal design produced is summarised in table 5.2. The objective function value for the base case wastewater treatment system is **30805** pounds sterling per year. The total volume of the system was found to be $33.4 m^3$ - just inside the constraint on total volume. The individual CSTR volumes found by optimisation are, on average, approximately 9% greater than those found using Walsh's approximation

<i>Parameter</i>	CSTR 1	CSTR 2	CSTR 3
Volume	11.3 m^3	11.5 m^3	10.6 m^3
Setpoint	-2.66 mN	$-3.39 \times 10^{-1} \text{ mN}$	$3.78 \times 10^{-4} \text{ mN}$
PI gain	$1.65 \times 10^{-2} \frac{\text{m}^3}{(\text{mN} \cdot \text{min})}$	$3.09 \times 10^{-3} \frac{\text{m}^3}{(\text{mN} \cdot \text{min})}$	$1.31 \times 10^{-2} \frac{\text{m}^3}{(\text{mN} \cdot \text{min})}$
PI reset time	1.43 min	0.57 min	0.77 min

Table 5.2: Optimal design for wastewater treatment system base-case

(see equation 5.16). Given that an approximate analysis has given values of the volume to within ten percent of the actual solution, it seems an appropriate mechanism for the initial screening of design alternatives.

Figure 5.6 shows the time profiles of the load disturbance and of the effluent concentrations from each of the three CSTRs in the process. The graphs show the system initially at the nominal steady-state for the process. One minute into the cycle, the concentration load drops. The process recovers to near steady-state some four minutes later and at this point the load steps back up to the nominal level. Recovery of the process back to the nominal steady-state takes a further four minutes. It should be noted that the CSTR 3 exit concentration shown in figure 5.6 remains within the specified pH bounds ($\pm 1 \times 10^{-2} \text{ mN}$) at all times.




Figure 5.6: Cyclic profiles for wastewater process base case

A feasible base case system has been found which gives an upper (pessimistic) bound on the economics of the process. This base case is now used as the starting point for the remainder of the design procedure which considers integration of a diversion system. This will be discussed in the next section.

5.6 Optimal cyclic operation of wastewater systems with diversion

In the previous section an optimal wastewater treatment system was found. This section uses this result as a basis to search for an economically superior solution which may be feasible when a diversion system is integrated into the design.

5.6.1 Economic bounding

At this stage in the design, the potential economic benefits of incorporating a diversion system are not known. Clearly, explicit enumeration of all possible designs (*i.e.* n CSTRs plus diversion) using the dynamic optimisation framework described earlier would be very numerically intensive and is therefore highly undesirable. Simple economic bounding properties do, however, exist which enable the search space to be greatly reduced. These will now be discussed.

Let the cost of the base case system found earlier equal C_{base} . Consider now the cost of adding a diversion system to this base case design. The minimum (annualised) cost of adding such a system is equal to the fixed cost of the diversion tank, *i.e.* £ (20000/3). The cost of the wastewater system can only be reduced compared to the base case by reducing the volumes in each of the CSTRs. Recall that, to ensure good mixing, there is a minimum residence time constraint on each CSTR. Given that the flowrate through the system is at least $3\text{ m}^3.\text{min}^{-1}$ (as defined in the problem specification), the minimum volume of a given CSTR will be 3 m^3 . This volume corresponds to an annualised CSTR cost of £ 8105. Therefore, it may be stated that the *minimum* annualised cost of a process containing n CSTRs and diversion system is:

$$C_{min} = 8150n + 6667 \quad (5.29)$$

Note that the minimum annualised cost of an n CSTR wastewater treatment system is simply $8150n$.

Given that a feasible solution exists with cost C_{base} , then an alternative system

with n CSTRs and diversion need only be considered further if:

$$C_{min} \leq C_{base} \quad (5.30)$$

Consider now the addition of a diversion system to the three CSTR base case treatment system optimised earlier. The minimum annualised cost of such a system using equation 5.29 is £ 31117. This lies above the base case value of £ 30805. Using equation 5.30, this system need not be considered as optimisation cannot produce an economically superior solution. This system will, however, be used later in this case study to reinforce general observations which will be made regarding systems incorporating diversion. Clearly, systems containing more than three CSTRs may also be excluded from the search space as, according to equation 5.29, the cost increases with increasing n . Only two possible systems remain to be considered as viable economic alternatives:

1. two CSTRs and diversion system,
2. one CSTR and diversion system.

Both these system have minimum costs below the base case level. The dynamic optimisation formulation presented earlier will now be applied to each of these systems in turn. Results will be presented for each system and following this a general discussion will be given.

5.6.2 System One: Two CSTRs and diversion

Convergence to the solution of this dynamic optimisation problem required some 663 CP minutes on a Sparcstation10. The design parameters obtained for the wastewater treatment system are shown in table 5.3 below, those for the corresponding diversion system are shown in table 5.4. The objective function value for this system is **25922**

<i>Parameter</i>	CSTR 1	CSTR 2
Volume	$3.79 m^3$	$5.25 m^3$
Setpoint	$-2.49 mN$	$-4.76 \times 10^{-4} mN$
PI gain	$2.48 \times 10^{-3} m^3/(mN.min)$	$4.16 \times 10^{-3} m^3/(mN.min)$
PI reset time	$0.41 min$	$0.29 min$

Table 5.3: Optimal design for system one: treatment system

pounds sterling per year, well below the base case cost of £ 30805. It should be noted that the minimum possible cost for the two CSTR system with diversion, found from

<i>Parameter</i>	Value
Maximum volume	5.15 m^3
Initial volume	0.71 m^3
Initial concentration	0.24 mN

Table 5.4: Optimal design for system one: diversion system

equation 5.29, is £ 22967. The total volume of this system is 14.2 m^3 . Figures 5.7 and 5.8 show profiles resulting from the optimisation for the states and the controls during the operating cycle.




Figure 5.7: System one: profiles during operating cycle - 1




Figure 5.8: System one: profiles during operating cycle - 2

5.6.3 System Two: One CSTR and diversion

Convergence to the solution for this problem required 411 CP minutes on a Sparcstation10. The design parameters obtained for the wastewater treatment system are shown in table 5.5 below, those for the corresponding diversion system are shown in table 5.6. The objective function value for this system is **17572** pounds sterling

<i>Parameter</i>	CSTR 1
Volume	$4.38 m^3$
Setpoint	$9.11 \times 10^{-4} mN$
PI gain	$3.21 \times 10^{-3} m^3/(mN.min)$
PI reset time	$0.45 min$

Table 5.5: Optimal design for system two: treatment system

<i>Parameter</i>	Value
Maximum volume	$6.10 m^3$
Initial volume	$0.16 m^3$
Initial concentration	$1.28 mN$

Table 5.6: Optimal design for system two: diversion system

per year, this is below the previous best cost of £ 25922. It should be noted that the minimum possible cost for the one CSTR system with diversion, found from equation 5.29, is £ 14817. The total volume of this system is $10.5 m^3$. Figures 5.9 and 5.10 show profiles resulting from the optimisation for the states and the controls during the operating cycle of this process.

Recall that in section 5.2 the assumption of monotonicity was made and that the resulting design developed in this work has therefore been based upon the worst-case disturbance. The monotonicity assumption was tested for the single CSTR design by application of a concentration disturbance of $0.225 M$ acid (50% of the worst-case change) and found to be unfalsified.




Figure 5.9: System two: profiles during operating cycle - 1




Figure 5.10: System two: profiles during operating cycle - 2

5.6.4 Discussion of optimisation results

In the previous sections, results for the two wastewater plants with integrated diversion were presented. These two sets of results will now be considered together so that comparisons may be drawn and trends identified.

The graphs presented show profiles for the load change, the effluent concentration of the final CSTR (which is also shown re-scaled), the flow of material actually discharged by the plant, the state trajectories in the diversion tank and the recycle flows back to the treatment system. Cost and volume results are summarised in table 5.7 below.

<i>System</i>	Cost (£)	Total volume (m^3)
Base case	30805	33.4
2 CSTRs + diversion	25922	14.2
1 CSTR + diversion	17572	10.5

Table 5.7: Summary of wastewater design results

With the problem specification in mind, the important features of both sets of results presented are that:

- At no time are the pH limits violated. As shown in figure 5.7 and figure 5.9, as the effluent concentration of the final CSTR approaches $\pm 10^{-2} mN$ the discharge flowrate is forced towards zero by the diversion mechanism.
- Both systems are well within the total volume limit of $35 m^3$.

The integration of a diversion system has enabled the problem specification to be met *both* more cheaply (a 45% cost saving over the base case wastewater treatment system) and with a greatly reduced total volume ($\frac{1}{3}$ of the base case volume). The diversion system essentially enables off-spec effluent to undergo a second pass through the treatment system for re-processing. Because of this, a more crude initial attenuation in concentration is possible, which facilitates size (and hence cost) reductions in the CSTRs.

The similarity in the profiles of the diversion system states for each case allows a fairly general explanation of the role played by the diversion system. The off-spec material leaving the CSTRs is forced into the diversion system. The diversion tank volume rises and its concentration tracks the concentration changes in the final treatment CSTR. The maximum diversion capacity required is kept to a minimum by recycling material to the CSTRs as soon as the load disturbance hits. Once the effect of the load disturbance has passed, the additional material accumulated in the diversion tank continues to be bled

back into the treatment system for re-processing. When the system has settled to a new steady-state the corresponding flows to and from the diversion system are zero (*i.e.* the diversion system is at steady-state). This settling time is 6.2 minutes for the two CSTR system and 12.6 minutes for the one CSTR system. Recall that the settling time for the base case was just four minutes. At this point in the optimisation cycle the load jumps back to its nominal value. Similar behaviour to that described for the first load change results. The settling times to steady-state are also approximately the same as for the first load change.

A further common feature of both sets of results which is not shown in the figures is that in all CSTRs the minimum residence time constraint of one minute becomes active. The residence time is defined as the ratio of vessel volume to flowrate and is added to ensure good mixing in the CSTRs. When this constraint is active, it effectively limits the minimum possible volume of a given CSTR. It can be seen from the figures that most of the recycled material in the two CSTR system is fed only to the second CSTR. This is sensible since (in the presence of an active residence time constraint) recycling material through the first CSTR would require an increase in its volume. The lower recycle rates used in the one CSTR system are probably attributable to the lower disturbance attenuation possible using only one CSTR. An effect of the lower recycle flow used is slower draining of the diversion tank and hence a factor of two increase in the settling time to steady-state over the two CSTR system.

The observations presented above may be summarised as follows.

1. The minimum residence time constraint is active on all CSTRs in both systems and this limits the minimum total volume of the system.
2. Given that (1) above applies, recycling most material to the last CSTR in the series keeps the volumes of the preceding CSTRs to a minimum.
3. As the number of CSTRs (and hence disturbance attenuation) increases, the maximum diversion capacity is reduced.
4. As the number of CSTRs increases the settling time to steady-state is reduced.
5. The greatest cost savings arise from being able to remove a unit from the process (and hence the associated fixed cost). The diversion system enables feasibility of operation to be retained.

The generalisation of these observations to other wastewater treatment systems with diversion is now investigated by considering a three CSTR system with diversion.

Recall that previously this system was excluded from further consideration by the economic bound given by equations 5.29 and 5.30. For reasons of brevity, only key findings from the optimisation of this third system are summarised here.

1. CSTR 1 has a minimum residence time of 1.2 minutes, the minimum residence time constraint of 1 minute is active on CSTRs 2 and 3.
2. most of the recycle from the diversion tank is returned to CSTR 3,
3. The maximum diversion capacity required is $3.46 m^3$ (*c.f.* $5.1 m^3$ for the 2 CSTR system and $6.1 m^3$ for the 1 CSTR system).
4. The settling time to steady-state from first load change is 5 minutes (*c.f.* 6 minutes for the 2 CSTR system and 13 minutes for the 1 CSTR system).
5. Total cost of system is £33898 (recall that the minimum cost of this system is £31117).

Apart from point one given above in which the one minute residence time bound is inactive, the 3 CSTR solution is consistent with the trends suggested previously.

5.7 Discussion

In this case study the problem of wastewater treatment plant design has been discussed. Traditional design practices were outlined in which a CSTR based chemical treatment system is designed and a diversion system added as a fail-safe mechanism. It was proposed that an integrated approach, which considered the diversion system design and operation together with that of the continuous wastewater treatment system, could yield economic benefits.

The following design procedure was used to rapidly determine the optimal design and operation of a system meeting the design specifications:

1. Apply Walsh approximation (equation 5.16) to select number of CSTRs for wastewater treatment system.
2. Use dynamic optimisation to determine the optimal design and cyclic operation of the structure found in 1. This now becomes an upper economic bound.
3. Use economic bounding equations to eliminate integrated treatment/diversion systems with cost greater than 2 above.

4. Enumerate remaining treatment/diversion systems using dynamic optimisation formulation for cyclic operation.

Application of this procedure required the solution of three dynamic optimisation problems with a total CP requirement of 1100 CP minutes on a Sparcstation10 to determine the economically optimal design. The optimal design has a total volume which is one third that of the optimal wastewater treatment system without diversion. The cost is reduced by almost a factor of two.

Although the single CSTR design produced has a very small footprint with significant cost reductions having been achieved, a number of other factors must be considered, some of which have been discussed previously.

- The system settling time is large, increasing from 4 minutes for the base case to 13 minutes.
- The optimal open-loop control policy applied to the diversion system is restricted in its application to the disturbance considered. A “traditional” continuous wastewater system operating under closed-loop PI control is better able to deal with a wider range of disturbances.
- The optimal design may be particularly sensitive to the effects of non-ideal reagent valve behaviour such as hysteresis.

In summary, the integrated design and operation of a diversion system with a continuous chemical treatment system has produced a significant intensification of the process in terms of economics and overall size. The points made above acknowledge that penalties must be paid for this intensification (*e.g.* longer settling times and potential sensitivity to non-ideal valve behaviour). Also acknowledged is the fact that the use of dynamic optimisation in the approach gives a rather idealised solution to the problem. A significant difference between a traditional wastewater process and a fully integrated solution has, however, been shown. It does not seem unreasonable to suggest that there are potential economic benefits to be found in the “middle-ground” lying between these two extremes. For example, it may be possible to gain some of the economic benefits found here by integrating the diversion system in a sub-optimal way, possibly by deriving more general control laws from the trajectories found by dynamic optimisation (*c.f.* Sargent and Sullivan, 1979).

Chapter 6

Design of a reactor network for distributed manufacturing

6.1 Introduction

This chapter considers the optimal design and operation of a distributed process plant capable of instantaneous response to a customer's intermittent demand for product. In this sense the case study represents a “worst-case” demand profile, since the process must be capable of automatic startup and shutdown. The form of the case study is intended to encompass many of the key process systems aspects of distributed manufacturing discussed in chapter 2.

The process design in this case study evolves from a problem definition based upon the customer's requirements from the plant, which is presented in section 6.2.1. Following this general statement of the design problem, the types of unit which will form the final process are discussed in section 6.3 together with their respective mathematical models. Questions of numerical robustness of the models in an optimisation framework are addressed in section 6.4.2.

After discussion of the two-stage design procedure to be employed in section 6.5, the form of the dynamic optimisation problem for cyclic operation is considered specifically for this case study. A generalised form of this optimal cyclic operation problem is then developed in terms of generic process building blocks. In section 6.7 a candidate process structure is developed and the optimal cyclic operation of this structure is then considered in 6.8. Section 6.8.2 presents the results of the cyclic optimisation and discusses the concept of hold-states. Following from this discussion, section 6.8.4 considers the cyclic operation of the candidate process in the presence of uncertainty.

6.2 General Problem Definition

6.2.1 Customer requirements

The owner of a medium sized industrial plant believes substantial economic improvements can be made by replacing one of the feedstocks to the plant with material B . Demand of B for the plant will be $12.3 \text{ litres.min}^{-1}$. The demand for B is intermittent, but it must be available on demand and lie within purity limits of 81.69% to 84.90% (w/w) of B . The plant operates for 365 days per year with one twelve hour shift per day.

6.2.2 Other operational constraints

With the economic incentives to switch feedstocks to B come additional operational issues which must be addressed.

1. B is known to degrade if kept in storage for long periods. In addition, the degradation accelerates as temperature increases and is considered significant above temperatures of 48.89°C (120°F).
2. Revised environmental regulations applying at the business site allow a maximum chemical inventory of B of 50 litres in a given vessel. In the event of a vessel failure this is the largest discharge of material to the environment that can be tolerated.
3. Total inventory of B for the entire site is limited to 400 litres .

These issues clearly have a considerable impact on the customer. Even if the maximum inventory of B were kept it would be sufficient to maintain demand for little more than 30 minutes. Difficulties in storing the feedstock B further compound the problems of meeting supply requirements. The possibilities for producing B on demand on the site are therefore to be explored.

Feedstock B may be generated from an environmentally acceptable feed component A via the exothermic, catalytic liquid phase reaction:



The remainder of this chapter will consider the issue of designing a distributed plant capable of meeting the feedstock demand requirements of the customer's plant.

6.3 Process building blocks

Production of B is based around the first order reaction scheme presented above. As previous work using this reaction scheme (de Hennin, 1992; Heath, 1993) has concentrated on continuous stirred tank reactors (CSTRs), a CSTR based design will be developed and evaluated. The general CSTR unit to be used is shown in figure 6.1. The unit incorporates two inlet streams: one for fresh feed of A and the other for connection to the preceding upstream unit. A single exit stream is provided. Each CSTR has facilities for cooling through the use of a cooling jacket with water as the cooling medium. Cooling water is available at a temperature of 20°C and has a maximum allowable discharge temperature of 48.89°C . Fresh feed of A is available at a temperature of 20°C and a concentration of 1kg.litre^{-1} . As mentioned previously, the reaction is catalytically driven. The catalyst may be introduced into or removed from the reacting liquid within each CSTR to achieve effective control over the reaction.

It is assumed in this case study that the instantaneous response demanded by the customer can be achieved by simply storing the required product. Given that product storage buffering may be needed, a product storage tank must be allowed for in the system. This is shown in figure 6.1 and will be discussed in section 6.4.1.2.




Figure 6.1: Process building blocks

6.4 Modelling the Problem

Having introduced the general design problem to be addressed and considered the operational constraints and material specifications which apply, specific issues in modelling the CSTRs and storage tank will now be discussed. These modelling aspects will be dealt with on two levels:

1. modelling the physical behaviour of the system,
2. modelling for numerical robustness.

6.4.1 Physical model of the system

6.4.1.1 Dynamic CSTR model

The dynamic behaviour of a CSTR is represented by a set of differential and algebraic equations based on dynamic mass and energy balances around the unit. In the notation below the superscript 2 refers to the properties (flow, concentration and temperature) of the fresh feed and the superscript 1 to the stream from the upstream unit.

1. Overall material balance.

$$\frac{dV_R}{dt} = F_{in}^1 + F_{in}^2 - F_{out} \quad (6.2)$$

2. Feed component material balance.

$$V_R \frac{dC_{out}}{dt} + C_{out} \frac{dV_R}{dt} = F_{in}^1 C_{in}^1 + F_{in}^2 C_{in}^2 - F_{out} C_{out} - k V_R C_{out} \quad (6.3)$$

3. Energy balance.

$$V_R \frac{dT_{out}}{dt} + T_{out} \frac{dV_R}{dt} = F_{in}^1 T_{in}^1 + F_{in}^2 T_{in}^2 - F_{out} T_{out} - k V_R C_{out} \frac{(\Delta H_R)}{\rho C_p} - \frac{Q}{\rho C_p} \quad (6.4)$$

In equations 6.3 and 6.4, the reaction rate, k , is defined as,

$$k = k_0 e^{\frac{-E_a}{R(460+T_{out})}} \quad (6.5)$$

The rate of heat removal, Q , used in equation 6.4 is given by,

$$Q = U A_H (T_{out} - T_c) \quad (6.6)$$

where the temperature of the cooling water, T_c , is defined as the mean of the inlet and outlet cooling water temperatures,

$$T_c = \frac{1}{2}(T_{c,out} + T_{c,in}) \quad (6.7)$$

An energy balance on the jacket side of the reactor gives,

$$Q = \rho C_p F_c (T_{c,out} - T_{c,in}) \quad (6.8)$$

where F_c is the volumetric flow of cooling water, Finally, the flow leaving the reactor is proportional to the square root of the reactor holdup,

$$F_{out} = C_v k_v V_R^{\frac{1}{2}} \quad (6.9)$$

where C_v is a valve constant and k_v is the valve position. In experiments using exothermic reactors to teach control techniques Brisk (1974) observed that,

under all relevant operating conditions the exit cooling water temperature approached the reactor temperature to $0.5 \pm 0.1^\circ F$.

This assumption is used as a simplification in the modelling of this problem. The expression for the rate of heat removal, Q , then becomes,

$$Q = \rho C_p F_c (T_{out} - T_{c,in} - 0.5) \quad (6.10)$$

which enables equations 6.6, 6.7 and 6.8 to be removed from the system. Note that in this expression the temperature units used are $^\circ F$. In addition, equation 6.5 describing the rate of reaction is extended to incorporate the effect of the catalyst as follows,

$$k = \left[k_0 e^{\frac{-E_a}{R(460+T_{out})}} \right] Catalyst_{control} \quad (6.11)$$

where the variable $Catalyst_{control}$ denotes the catalyst in the reactor and may take a value from zero to one. A $Catalyst_{control}$ value of one gives the maximum reaction rate for a given temperature, for a value of zero there is no reaction. Table 6.1 below summarises the kinetic parameters used in the model.

6.4.1.2 Storage tank

It is assumed that no energy transfer or reaction occur in the storage tank. The storage tank may then be modelled simply by dynamic overall and component mass balances in a similar fashion to the CSTR model above:

<i>Parameter</i>	<i>Value</i>
K_0	$5.4 \times 10^{11} \text{ min}^{-1}$
E_a/R	$16000 \text{ }^\circ R$
ΔH_R	$-6089 \text{ kJ(kmol)}^{-1}$

Table 6.1: Reactor parameters for case study

1. overall material balance,

$$\frac{dV_{store}}{dt} = F_{in} - F_{product} \quad (6.12)$$

2. feed component material balance,

$$V_{store} \frac{dC_{product}}{dt} + C_{product} \frac{dV_{store}}{dt} = F_{in} C_{in} - F_{product} C_{product} \quad (6.13)$$

The product flow rate from the storage tank is determined solely by the customer and is assumed to be a pumped flow. The implications of this will be considered in the following section.

6.4.2 Numerical robustness

In the preceding section the models which describe the physical behaviour of the process units were presented. It is important, however, to also consider the context in which these models are to be applied. As discussed previously, a model which accurately represents the physical system and is robust during the course of a dynamic simulation will not necessarily be numerically robust in the regions that may be entered during solution of the dynamic optimisation problem. It is therefore important to ensure that models to be used for optimisation purposes are as numerically reliable as possible. Two specific instances arise in the models described above, one of which has been discussed previously.

6.4.2.1 Preventing negative storage holdup

In order to ensure that the diversion tank is numerically robust during the optimisation, a proportionally controlled make-up stream is added to the dynamic model to maintain the level in the vessel above a pre-specified lower limit. A detailed discussion on this subject and the form of the equations implemented can be found in section 4.4.2.3.1. A path constraint on the tank holdup ensures that the proportional controller does not operate at the solution point (see 4.5.5.2.1).

6.4.2.2 Undefined derivatives from head-driven exit flow

Recall that previously the exit flow from the CSTR was defined as being proportional to the square root of the holdup in the reactor:

$$F_{out} = C_v k_v V_R^{\frac{1}{2}} \quad (6.14)$$

Consider a situation in which the reactor is drained such that the holdup approaches zero. The value of F_{out} remains well defined throughout the draining operation and as the holdup reaches zero the flow out is defined by 6.14 as zero, thus preventing volume reducing further and becoming negative. However, the derivative of F_{out} with respect to V_R is given by:

$$\frac{dF_{out}}{dV_R} = -\frac{1}{2} C_v k_v V_R^{-\frac{1}{2}} \quad (6.15)$$

As V_R tends to zero the value of this derivative becomes undefined such that any step taken by the integrator will result in a bounds violation of the state, V_R , which is not permitted to become negative. Further discussion on the problem of rapid transients in the integration of dynamic systems may be found in (Jarvis, 1993).

Since this circumstance cannot be precluded during the course of the optimisation, a more robust form of equation 6.4.2.2 is employed:

$$F_{out} = \begin{cases} C_v k_v V_R^{\frac{1}{2}} & \text{if } V_R > V_{R_h} \\ \frac{C_v k_v V_{R_h}^{\frac{1}{2}} (V_R - V_{R_l})}{(V_{R_h} - V_{R_l})} & \text{otherwise} \end{cases} \quad (6.16)$$

Should the holdup fall to V_{R_h} , the value of F_{out} is defined as linear function of V_R . Note that the first derivative with respect to time at the switching point is continuous. Most importantly, the derivative of F_{out} with respect to V_R remains well defined for all V_R . The variable V_{R_l} takes a value between zero and V_{R_h} and fixes a lower limit on V_R . Setting V_{R_l} slightly positive enhances the robustness introduced by the conditional equation above by preventing V_R from becoming exactly zero. The values chosen for V_{R_h} and V_{R_l} for this case study were 1×10^{-3} and 5×10^{-4} respectively.

6.5 Design procedure

At the most general level the problem to be solved is to determine the economically optimal combination of,

1. unit type and topology,
2. unit design,
3. cyclic operation of the complete process.

As discussed earlier, the design procedure used in this work does not attempt to simultaneously solve the combined structural and operational problem described above. Recall that the two stages of the proposed procedure are to,

1. develop a candidate process structure,
2. apply a dynamic optimisation formulation to determine the corresponding optimal operating cycle and design for that candidate.

By decomposing the problem into two levels, potentially expensive dynamic optimisation techniques are applied only when a viable process structure is in place. The resulting second stage optimisation problem is greatly simplified by the absence of combinatorially explosive integer decisions used to determine the process structure.

In the remainder of this chapter the proposed design procedure is used to design a plant capable of fulfilling the customer requirements outlined in 6.2.1. The specified customer requirements from the problem description are considered in the context of an optimisation framework for cyclic operation. The cyclic optimisation problem is then derived for this case study in a general form for a process design which incorporates n CSTRs in series. Following this, the development of potential candidate structures for this problem is considered in detail. Useful bounding properties are given and the candidate structure to be carried forward to stage two is presented. The cyclic optimisation formulation is then applied specifically to this candidate to determine the optimal operating cycle. From this optimisation, desirable design attributes for this class of distributed plant (hold-states) are established. Investigation of the process operation under the effects of non-ideality further demonstrates the relevance of these attributes.

6.6 Optimal cyclic operation

The previous sections in this chapter have dealt with specification of the problem in terms of customer requirements and external constraints. In addition, robust modelling of the unit operations forming the process has been discussed. This section considers the way in which information in the problem description is mapped into the optimisation framework described earlier in chapter 3. Figure 6.2 shows how the problem description




Figure 6.2: Diagram showing original operating cycle (top) and the collapsed cycle which is implemented (bottom).

maps to the cyclic operation framework. At time t_0 the customer demands product. The process then moves from its shutdown-state to the production steady-state, which is reached at time t_1 . At this time the controls are frozen. At time $t_{product}$ no further product is required and demand drops to zero. Simultaneously, the process moves back to the shutdown-state defined at time t_0 . This state is reached at time t_{cycle} . The cycle for the actual process is collapsed as shown in figure 6.2 to avoid unnecessary and time consuming integration through the steady state period during solution. Provided the correct function for the process operating cost is used (as will be discussed later), the two cycles are identical for the purposes of determining the optimal cyclic operation.

As previously discussed, the process to be designed can be described in general by the use of two generic building blocks; a CSTR block and storage tank block. Each unit in the process is modelled by the generic building block for that type of unit. In a similar manner, the cyclic optimisation problem may also be written in a general form with each occurrence of a given type of unit making an associated contribution to the objective function, the number of constraints and number of optimisation parameters in the problem. The general form of the problem containing n CSTRs and storage tank is developed in the following section.

6.6.1 Objective function

As discussed in section 2.3 one of the underlying requirements when posing an integrated design problem is that the problem can be *written entirely in terms of costs and constraints*. The motivation for solving such a problem is to find the cheapest process design satisfying the constraints on the design. The objective function for the optimisation should therefore be an economic one. The specific form of this objective will now be considered.

Recall that the customer requirements given in 6.2.1 specify a production rate of $12.3 \text{ litres.min}^{-1}$. The precise profile for the demand over the course of a year is, however, uncertain. Figure 6.2 shows a nominal operating cycle for the process. The cycle consists of three phases: two transient phases (start up and shut down) and one steady phase. In any given year there will a number of such cycles, N_{cycle} . Each cycle has associated with it a period for which product is demanded, $t_{product}$, as shown in the figure. The final design of the process will be a function of the form of this demand profile. For example, a plant which is optimally designed for a demand profile consisting of a large number of short production runs (large N_{cycle} and small $t_{product}$) will typically have different characteristics to one whose design is based upon much fewer runs of longer

duration (large $t_{product}$ small N_{cycle}). Qualitatively, increasing N_{cycle} while reducing the length of $t_{product}$ would be expected to increase economic impact of the transient phases in the operating cycle. Conversely, as $t_{product}$ increases and N_{cycle} is reduced it is not unreasonable to expect the economics of the optimal production steady-state to dominate. The economic objective function to be used should therefore take account of the form of the demand profile. The cost of a single operating cycle of the process may be written as:

$$\text{operating cost per cycle} = C_{su} + C_{sd} + (t_{product} - t_1)C_{ss}$$

where C_{su} is the cost of the startup phase, C_{sd} is the cost of the shutdown phase, C_{ss} is the cost per unit time of the steady-state operation and t_1 is the time to start up (which must be less than $t_{product}$). The expected annualised operating cost of the process is therefore given by:

$$\begin{aligned} E(\text{annual operating cost}) &= E(N_{cycle})(C_{su} + C_{sd}) + \\ &E(N_{cycle}(t_{product} - t_1)) C_{ss} \end{aligned} \quad (6.17)$$

If $t_{product}$ and N_{cycle} are independent then:

$$E(N_{cycle} \cdot t_{product}) = E(N_{cycle})E(t_{product}) \quad (6.18)$$

The *total* expected annualised cost of the process including capital costs is therefore given by:

$$\begin{aligned} E(\text{annualised cost}) &= E(N_{cycle})(C_{su} + C_{sd}) + C_{capital} + \\ &E(N_{cycle})(E(t_{product}) - t_1) C_{ss} \end{aligned} \quad (6.19)$$

where $C_{capital}$ is the annualised process capital cost.

The optimisation objective function is then:

$$J = \min_{u(t), d} E\{\text{annualised cost}\} \quad (6.20)$$

where $u(t)$ are the time varying controls acting on the process and d is the set of design parameters.

For this case study it is assumed that, on average, the plant will undergo one operating cycle per day with a mean product delivery time of twelve hours. This corresponds to values for $E(N_{cycle})$ and $E(t_{product})$ of 365 and 12 hours respectively.

6.6.1.1 Capital costs

As stated previously, the final design will consist of a number of CSTRs followed (potentially) by a product storage tank. The capital cost of each CSTR building block as shown in figure 6.1 is given by,

$$C_{CSTR} = 1 \times 10^5 + 9.629 \times 10^3 (V_R^{max})^{0.7} \quad \text{cost units} \quad (6.21)$$

where V_R^{max} is the maximum level of holdup in the CSTR (in *litres*) during the operating cycle. The cost of the storage tank building block is defined in a similar way,

$$C_{store} = 5 \times 10^4 + 4.814 \times 10^3 (V_{store}^{max})^{0.7} \quad \text{cost units} \quad (6.22)$$

where V_{store}^{max} is the maximum storage requirement (in *litres*) during the operating cycle. For a system containing n CSTRs and taking the most general case when product storage is required, the total annualised capital cost may be written as,

$$C_{capital} = \frac{1}{3} \left\{ \sum_{i=1}^n C_{CSTR_i} + C_{store} \right\} \quad \text{cost units} \quad (6.23)$$

The $\frac{1}{3}$ pre-multiplier is a capital charge factor (CCF) which approximates “the time value of money” ((Douglas, 1988)), and converts the total capital cost to an annualised equivalent.

6.6.1.2 Operating costs

The operating cost may also be considered for a single CSTR building block. In each CSTR the fresh feed, cooling water flow and catalyst all contribute to the operating cost. Fresh feed of A is costed at $7.063 \times 10^{-2} \text{cost units.litre}^{-1}$. Cooling water is costed at $3.532 \times 10^{-4} \text{cost units.litre}^{-1}$. During use the catalyst will become poisoned and will periodically need to be renewed. The costs associated with this are incorporated as follows: when being “consumed” by the liquid reactant, the catalyst is costed at $1.766 \times 10^{-2} \text{cost units.min}^{-1} \text{.litre}^{-1}$, otherwise it has zero cost. Note that in the case of catalyst cost, the unit litre^{-1} is per unit reactor volume since the catalyst required is assumed to be in proportion to the maximum reactor holdup. Product pumping costs from the storage tank are neglected so that this unit has zero operating cost. The operating cost per unit time for a single CSTR is therefore given by,

$$C_{CSTR_{op}}(t) = 7.063 \times 10^{-2} F_{in}^2(t) + 3.532 \times 10^{-4} F_c(t) + 1.766 \times 10^{-2} V_R^{max} Catalyst_{control}(t) \quad (6.24)$$

where, as before, $Catalyst_{control}$ takes a value from 0 to 1 and indicates the presence of catalyst in the reactor. The cost of an operating cycle per CSTR is then given by,

$$C_{CSTR_{op}} = 7.063 \times 10^{-2} \int_0^{t_{cycle}} F_{in}^2(t) dt + 3.532 \times 10^{-4} \int_0^{t_{cycle}} F_c(t) dt + 1.766 \times 10^{-2} V_R^{max} \int_0^{t_{cycle}} Catalyst_{control}(t) dt \quad (6.25)$$

where t_{cycle} is the length of a *process* operating cycle. As discussed earlier, the process operates at steady-state for a period of time $(t_{product} - t_1)$. The optimisation cycle considered is of length t_f , where $t_f = t_{cycle} - (t_{product} - t_1)$, since (as stated previously) integration over the steady-state period of operation is unnecessary. The cost of the steady-state period must, however, be incorporated into the objective function. Equation 6.25 is therefore modified to reflect this.

$$C_{CSTR_{op}} = \int_0^{t_f} (7.063 \times 10^{-2} F_{in}^2(t) + 3.532 \times 10^{-4} F_c(t) + 1.766 \times 10^{-2} V_R^{max} Catalyst_{control}(t)) dt + (7.063 \times 10^{-2} F_{in}^2(t_1) + 3.532 \times 10^{-4} F_c(t_1) + 1.766 \times 10^{-2} V_R^{max} Catalyst_{control}(t_1))(t_{product} - t_1)$$

The total annual operating cost given that the plant completes N_{cycle} per year is then:

$$C_{op} = N_{cycle} \sum_{i=1}^n C_{CSTR_{op_i}} \quad (6.26)$$

The objective function (6.20) to be used is therefore,

$$\min_{u(t),d} \{C_{op} + C_{capital}\} \quad (6.27)$$

It should be noted that in the objective function described above the shutdown-state of the plant at time t_f is assumed to have an associated zero operating cost. As will be seen later in this chapter this assumption may not always apply. This case will be discussed in depth in section 6.9.3.

6.6.2 Plant parameters for optimisation

Determination of the best combination of plant design and operating policy is achieved by the manipulation of three distinct sets of design parameters in the process:

1. time varying optimisation parameters (*controls*),
2. the plant initial state,
3. equipment sizes.

Similarly to the discussion of the objective function earlier, analysis of the optimisation parameters to be used in this case study may be considered in terms of each process building block. Four controls are available on each CSTR (see figure 6.1) as follows:

1. flow of fresh A to the reactor,
2. flow of cooling water through the cooling jacket,
3. exit flow from reactor (through manipulation of the control valve position),
4. position of catalyst system.

In the case of the product storage tank no control variables are available. The product demand from storage is set by the customer and the feed to the storage is supplied by the exit flow of the preceding reactor (which is a control variable for that reactor).

Each CSTR unit is described by three state variables: volume, concentration and temperature. The storage tank is described by just two state variables; volume and concentration. In the cyclic operation framework proposed the initial values for all of these states are included as optimisation parameters.

Finally, the volume of all CSTR units in the plant together with that of the storage tank are optimised. The equipment size is given by the maximum value of the holdup in a unit during the operating cycle. Thus, the actual optimisation parameter used for each unit is V_R^{max} . Locating the maximum volume in each unit during operation will be discussed in section 6.6.4.1.

6.6.3 Control parameterisation and time discretisation

Recall from chapter 3 that a feasible path control parameterisation method will be used to solve the dynamic optimisation problem. This section considers the choice of both the functional form of the controls and the manner in which the time horizon is discretised.

6.6.3.1 Control parameterisation

As discussed in section 3.5.3.1 piecewise linear controls will be implemented as these give the highest precision while still ensuring the control variables remain correctly bounded using just simple bounds on the break points.

6.6.3.2 Time discretisation

As with the functional form of the controls, the number of discretisation intervals must also be fixed *a priori* in the problem description. Although a complete operating cycle is being optimised, it essentially consists of three phases: startup to production steady-state, product ramped to zero and shutdown. The number of variable time intervals chosen will therefore be split equally between the startup and shutdown phases of the cycle. Thirteen discretisation intervals are used, which allows 6 intervals for each of the startup and shutdown phases. This is anticipated to give sufficient flexibility during solution whilst giving an acceptable number of optimisation parameters (recall from section 3.5.3.1 this should ideally be below 200) for designs incorporating up to 3 CSTRs.

In addition to fixing the number of time discretisation intervals to be used, the lengths of those intervals which are optimised must also be bounded. An upper bound is arbitrarily set at 80 minutes. The lower bound on the interval length was set at 0.1 minutes. These bounds were inactive at the optimum solution. The fixed time interval over which the product flow is ramped to zero is given a value of 1×10^{-3} minutes to represent the product flow being cut-off by the customer.

6.6.4 Constraints

The optimal design and operation of the reactor network is subject to a number of constraints arising from two distinct sources:

1. operational constraints on the plant itself,
2. constraints required by the cyclic optimisation framework.

In this section both sources of constraints will be discussed in relation to the context of the optimisation problem to be solved.

6.6.4.1 Path Constraints

Path constraints are those constraints on the system which must be met for all times in the optimisation. The majority of these result from operational requirements.

Recall from section 6.2.1 that a number of factors have led to on-site manufacture of B being considered, namely:

- limits on inventory in individual vessels and on total site inventory,
- degradation of B in storage, which becomes significant above 48.89°C .

The implication of the inventory limits is clear, the second factor is incorporated by restricting the maximum temperature of any vessel in which B is produced to 48.89°C . These requirements must be met at all times and may be represented as inequality path constraints in the optimisation. The constraints may be generalised for n reactor units as follows.

$$\left. \begin{array}{l} T_{out_i} \leq 48.89 \\ V_{R_i}^{max} \leq 50 \end{array} \right\} i = 1, n \quad (6.28)$$

No heating or cooling occurs in the storage unit and so only the maximum inventory constraint is applied:

$$V_{store}^{max} \leq 50 \quad (6.29)$$

The final operational constraint on the plant is that the customer requires all feedstock supplied to contain between 81.69% and 84.90% (w/w) of B . Again, this requirement must apply at all times and is written as the path constraint:

$$81.69 \leq C_{product} \leq 84.90 \quad (6.30)$$

where $C_{product}$ is the concentration of B in the product. In addition to those path constraints arising naturally as a result of the plant operation, a number of features of the optimisation framework itself are expressed through the use of path constraints. A means of “marking” the maximum holdup in a unit is required both for costing the plant capital (as part of the objective function) and for use in the inventory constraints described above. This is achieved by the following path constraint:

$$V_{R_i}(t) \leq V_{R_{max_i}} \quad i = 1, n \quad (6.31)$$

where $V_{R_{max_i}}$ is a time invariant optimisation parameter. The inventory constraints may then be met simply by placing appropriate bounds on the parameter $V_{R_{max_i}}$.

Recall that numerical robustness in the storage tank model was achieved through a proportionally controlled make-up stream (see section 4.4.2.3.1 for discussion). This was added to prevent negative holdup in this tank. At the solution point the make-up stream must not be active and a further path constraint must be added to ensure this. The form of this constraint was discussed previously in section 4.5.5.2.1.

6.6.4.2 Equality point constraints

The introduction to this section discussed the way in which this case study fits into the general framework proposed in chapter 3. The remaining constraints to be

included in the problem definition are a specific requirement of this framework. Briefly, in the context of this case study the purpose of these constraints is to:

1. ensure that a steady-state is reached following the startup phase,
2. ensure that a steady-state is reached following the shutdown phase,
3. ensure that the complete operation is cyclic.

At time t_1 in the operating cycle the plant is required to be at a steady-state capable of meeting the demand requirements of the customer. Steady-state at this point ensures that this demand is sustainable *ad infinitum*. At the final time in the operating cycle, t_f , customer demand is zero and the plant must have reached a new shutdown steady-state. As at time t_1 , this state is sustainable indefinitely until the customer again requires product. By requiring the state of the plant at the start of the cycle, t_0 , to be the same as that at t_f , the cycle itself is guaranteed to be repeatable indefinitely. These requirements are enforced as equality constraints which must be satisfied at specific points in time (t_1 and t_f) and are applied to the states and their respective time derivatives in each unit. Each CSTR unit has three states describing its temperature (T_{out}), concentration (C_{out}) and holdup (V_R). The storage tank has two states describing its concentration ($C_{product}$) and holdup (V_{store}). Since the dynamic behaviour of the process states is determined using optimal open loop control, exact steady-state requirements must be imposed at t_1 and t_f . For a system of n CSTRs and storage tank the constraints are therefore:

1. steady-state constraints

$$\left. \begin{array}{lcl} \dot{T}_{out_i}(t) & = & 0 \\ \dot{C}_{out_i}(t) & = & 0 \\ \dot{V}_{R_i}(t) & = & 0 \end{array} \right\} i = 1, n \quad \left. \vphantom{\begin{array}{lcl} \dot{T}_{out_i}(t) & = & 0 \\ \dot{C}_{out_i}(t) & = & 0 \\ \dot{V}_{R_i}(t) & = & 0 \end{array}} \right\} t = \{t_1, t_f\} \quad (6.32)$$

$$\left. \begin{array}{lcl} \dot{C}_{product}(t) & = & 0 \\ \dot{V}_{store}(t) & = & 0 \end{array} \right\}$$

2. cyclic operation constraints

$$\left. \begin{array}{lcl} T_{out_i}(t_0) & = & T_{out_i}(t_f) \\ C_{out_i}(t_0) & = & C_{out_i}(t_f) \\ V_{R_i}(t_0) & = & V_{R_i}(t_f) \end{array} \right\} i = 1, n \quad (6.33)$$

$$\begin{array}{lcl} C_{product}(t_0) & = & C_{product}(t_f) \\ V_{store}(t_0) & = & V_{store}(t_f) \end{array}$$

where $\dot{C}$ denotes the time derivative of C etc.

Two final equality constraints are then required to synchronise the times t_1 and t_f with the optimised time discretisation intervals. t_1 must coincide with the end of the sixth time interval (end of the startup period). t_f is the final time of the cycle and must coincide with the end of the thirteenth time interval. If the i th time interval length is denoted by Δt_i , the constraints implemented are,

$$\begin{aligned} t_1 &= \sum_{i=1}^6 \Delta t_i \\ t_f &= \sum_{i=1}^{13} \Delta t_i \end{aligned} \quad (6.34)$$

Viewed in terms of the number of constraints, the problem size may range from 23 constraints¹ for a system with one CSTR and storage tank to 45 constraints for a three CSTR system with storage tank.

6.7 Development of candidate structures

In this section the systematic development of a candidate structure for cyclic optimisation will be discussed.

Recall that, essentially, the aim here is to design a plant which can deliver product of a given purity at a given production rate. If the product demand was continuous the process design would be fairly straightforward - it may therefore be said that it is the *intermittent* demand for instantaneous product which complicates the design problem in this case. During its operation the process will exist in one of four operational modes:

1. steady-state production,
2. steady-state with zero production,
3. transition from mode 2 to 1,
4. transition from mode 1 to 2.

Clearly, the final design will have to be feasible for all of these operational modes. A sensible starting point in developing a candidate structure is therefore to assess the feasibility of the design in each operational mode.

¹excluding simple bounds

1. feasible steady-state production

B is produced from one reactant in a single reaction step. Production of the required rate and quality of product will be realisable using one or more CSTRs, providing a suitable combination of reactor conditions (*i.e.* residence time, reactor temperature) exists. As will be seen later, these reactor conditions can be determined from steady-state simulation or, indeed, from steady-state optimisation.

2. feasible steady-state with zero production

Just as the process must be able to produce product continuously for a variable period of time, it must also be capable of remaining in a static idle state for the period when no product is required (as was discussed earlier in section 6.6). Feasible operation in this mode is easily determined by analysing the steady-state form of the system equations given in section 6.4. Since the product demand is zero, the overall material balance dictates that all material flows in the process must similarly be zero (assuming simple connectivity between units and no recycles). With zero flows in place, the component material balances and energy balances in each CSTR reduce to:

$$0 = kV_R C_{out} \quad (6.35)$$

$$0 = -kV_R C_{out} \frac{(\Delta H_R)}{\rho C_p} \quad (6.36)$$

In other words, for the design to have a feasible shutdown steady-state one of either k , C_{out} or V_R must be moved to zero at that steady-state.

3. feasible transitions from mode 1 to 2 and from mode 2 to 1

It is assumed at this stage of the design procedure that the required transitions between operating points are feasible. Recall also that a premise in this work is that the instantaneous response demanded by the customer can be achieved by storing the required product. Since this process has a single product it is assumed here that a single storage vessel gives sufficient buffering between the customer and the process during transient modes of operation. In other words the storage will act as a source for product during operational mode 3 and a sink for product during operational mode 4.

Given that the assumptions made regarding the transitional periods are correct, any process containing a series of one or more CSTRs and a storage tank should be feasible as a design alternative. It would, however, be computationally inefficient to determine the most economic alternative by enumerating all such designs in the second stage dynamic

optimisation. A method for efficient screening of these many alternatives is needed. Such a method will be now be discussed.

6.7.1 Steady-state based bounding

Consider the design of a second process which will now be referred to as the “steady-state process”. This process is required to produce B at the same rate and purity as the distributed plant being designed in this case study (which will be referred to as the “dynamic process” in the following text). However, the steady-state process will produce B over one continuous period of operation. The total annual production of both processes is identical. The initial exploration of designs is carried out using the heuristic that:

the optimal annualised cost of the steady-state process will be less than or equal to the cost of the dynamic process .

The validity of this assumption will be reviewed later in the light of results. Figure 6.3 shows how this property may be effectively used as a economic bound to screen alternatives. Given that an n CSTR system gives the steady-state economic optimum (as




Figure 6.3: Screening dynamic process alternatives through steady-state optimisation

shown in figure 6.3) then provided the cost of the equivalent dynamic system is less than the next cheapest steady-state system (shown as $n + 1$ CSTRs in the figure) no other dynamic systems need be considered. Determination of the steady-state optimum for a

system with n CSTRs is relatively inexpensive when compared to the cost of dynamic optimisation of a complete operating cycle for the same system. Finding the steady-state optimum number of CSTRs in the series should therefore be both cheap and comparatively straightforward. The optimum steady-state system is then carried forward as the candidate process to which the cyclic optimisation is applied. When the bounding property described above does not hold (*i.e.* the cost of the dynamic process exceeds the second-best steady-state solution) further analysis of the dynamic solution is required. The analysis aims to identify reasons for the increased cost and determine the extent to which these increases will propagate through to the alternative systems.

6.7.2 Steady-state optimisation of the reactor network

In this section the bounding property described above will be applied specifically to the case study to establish a candidate structure for stage two of the design procedure. The optimisation problem to be solved is derived directly from the cyclic optimisation formulation presented in section 6.6. Indeed, the optimisation to be performed here may be considered analogous to determination of the operating point at t_1 in the cyclic operation shown in figure 6.2 .

6.7.2.1 Objective function

The objective function to be used is the minimisation of the total annualised cost of the process.

$$\min_{u,d} \{C_{op}^{ss} + C_{capital}\} \quad (6.37)$$

The individual contributions of capital and operating costs are determined as follows. Equation 6.24 defined the operating cost per minute of a single CSTR block as,

$$C_{CSTR_{op}}^{ss} = 7.063 \times 10^{-2} F_{in}^2 + 3.532 \times 10^{-4} F_c + 1.766 \times 10^{-2} V_R^{max} Catalyst_{control} \quad (6.38)$$

Recall that the total annual production of the steady state process must equal that of the dynamic process. Given this, the annualised operating cost of a system with n CSTRs will be,

$$C_{op} = N_{cycle} t_{product} \sum_{i=1}^n C_{CSTR_{op_i}}^{ss} \quad (6.39)$$

The annualised capital cost of the process is recalculated without a term for the cost of storage, which is not required during steady operation,

$$C_{capital} = \frac{1}{3} \sum_{i=1}^n C_{CSTR_i} \quad \text{cost units} \quad (6.40)$$

where,

$$C_{CSTR_i} = 1 \times 10^5 + 9.629 \times 10^3 (V_{R_i}^{max})^{0.7} \quad \text{cost units} \quad (6.41)$$

6.7.2.2 Optimisation parameters

Referring back to section 6.6.2 it can be seen that parameters available for optimisation are the steady-state values of the controls and the equipment sizes. The state of the plant is no longer a degree of freedom as this is determined at steady-state for a given set of control values. To summarise from section 6.6.2, the optimisation parameters in each reactor will be,

1. flow of fresh A to the reactor,
2. flow of cooling water through the cooling jacket,
3. position of catalyst system,
4. reactor holdup.

6.7.2.3 Constraints

In addition to providing simple bounds on the optimisation variables the same operational constraints must be applied as were used in the dynamic optimisation formulation. The path constraints discussed in section 6.6 are therefore applied as inequality constraints on the steady-state optimisation as follows:

$$\left. \begin{array}{l} T_{out_i} \leq 48.89 \\ V_{R_i}^{max} \leq 50 \end{array} \right\} i = 1, n \quad (6.42)$$

$$81.69 \leq C_{product} \leq 84.90 \quad (6.43)$$

Note that the maximum holdup constraint is actually enforced directly by appropriate bounding of the holdup optimisation parameter.

6.7.2.4 Implementation and solution

The simulation environment SpeedUp was used to carry out this phase of the design as it enabled the optimisation problem to be implemented in a straightforward and efficient manner. The SRQPD optimisation algorithm was used and proved adequate for solution of the steady-state optimisation problem with solution times on a Sun Sparcstation 10 of less than 1 CPU second.

6.7.3 Results from steady-state optimisation

Reactor networks consisting of 1, 2 and 3 CSTRs were optimised at steady-state. The results from these optimisations are presented below in tables 6.2, 6.3 and 6.4. The total annualised costs for each system are summarised in table 6.5. From table 6.5 it can be seen that the most economic solution (with a cost of 649640 cost units) is to use two CSTRs. Having found a minimum cost, further optimisation of systems with more than three CSTR units is not pursued. It is interesting to note that the three solutions presented share a number of features in common:

1. all reactors are operated at the maximum allowable temperature,
2. the B produced is of the lowest allowable purity (81.69%),
3. all fresh feed required enters through the first reactor.

The choice of these operating decisions may be explained as follows.

- Operation of the reactors at the maximum allowable temperature gives the maximum conversion of A to B for a given reactor residence time.
- Producing B of increasing purity requires greater residence times and results in increasing costs.
- Feeding all the fresh A to the first reactor gives a residence time distribution for the system which is as close to a plug flow reactor as possible (G.F. Froment and Bischoff, 1990). This arrangement gives the highest conversion for a given CSTR network involving a first order rate of reaction.

Further investigation reveals the economic penalty which must be paid if the feed of fresh A is split between the reactors in the series. A reactor network with two CSTRs was re-optimised with varying values for the fraction of fresh A fed to the first reactor in the sequence. Graph 6.4 shows the resulting change in the economics of the process. Note that the economics plotted for the case of zero feed to the first CSTR correspond to those for a system with one CSTR only. As a concluding remark on the results from the steady-state analysis, it should be noted that to obtain the required product purity in a system containing just one CSTR would require a CSTR volume of 101 *litres* which violates the maximum holdup limit of 50 *litres*.

<i>Design parameter</i>	
Fresh feed of A	12.32 <i>litres.min</i> ⁻¹
Cooling water flow	3.52 <i>litres.min</i> ⁻¹
Position of catalyst	1.00
Holdup	101.07 <i>litres</i>
<i>Reactor states</i>	
Temperature	48.89 °C
Concentration of B	0.82 <i>kg.l</i> ⁻¹
<i>Economics</i>	
Annual operating cost	709373
Annualised capital cost	114554
Total annualised cost	823927

Table 6.2: Optimal steady-state design for 1 CSTR

<i>Design parameter</i>	CSTR 1	CSTR 2
Fresh feed of A	12.32 <i>litres.min</i> ⁻¹	0.00 <i>litres.min</i> ⁻¹
Cooling water flow	0.00 <i>litres.min</i> ⁻¹	3.52 <i>litres.min</i> ⁻¹
Position of catalyst	1.00	1.0
Holdup	38.95 <i>litres</i>	22.30 <i>litres</i>
<i>Reactor states</i>		
Temperature	48.89 °C	48.89 °C
Concentration of B	0.64 <i>kg.l</i> ⁻¹	0.82 <i>kg.l</i> ⁻¹
<i>Economics</i>		
Annual operating cost	409320	103796
Annualised capital cost	74992	61532
Total annualised cost	484312	165328

Table 6.3: Optimal steady-state design for 2 CSTRs

<i>Design parameter</i>	CSTR 1	CSTR 2	CSTR 3
Fresh feed of A	12.32 <i>litres.min</i> ⁻¹	0.00 <i>litres.min</i> ⁻¹	0.00 <i>litres.min</i> ⁻¹
Cooling water flow	0.0 <i>litres.min</i> ⁻¹	2.07 <i>litres.min</i> ⁻¹	1.45 <i>litres.min</i> ⁻¹
Position of catalyst	1.00	1.00	1.00
Holdup	38.94 <i>litres</i>	9.21 <i>litres</i>	9.21 <i>litres</i>
<i>Reactor states</i>			
Temperature	48.89 °C	48.89 °C	48.89 °C
Concentration of B	0.64 <i>kg.l</i> ⁻¹	0.74 <i>kg.l</i> ⁻¹	0.82 <i>kg.l</i> ⁻¹
<i>Economics</i>			
Annual operating cost	409320	42907	42850
Annualised capital cost	74992	48513	48513
Total annualised cost	484312	91420	91363

Table 6.4: Optimal steady-state design for 3 CSTRs

<i>System</i>	<i>Annualised cost</i>
1 CSTR	823927
2 CSTRs	649640
3 CSTRs	667095

Table 6.5: Summary of optimal reactor network costs




Figure 6.4: Variation of process economics with fraction of fresh feed to first reactor

Through the use of steady-state optimisation the system with lowest annualised cost is identified as having two CSTR units. Based on the steady-state bounding procedure described earlier this system is carried forward as the candidate structure to be considered in the second stage of the overall design procedure. In the next section solution of this optimisation problem will be discussed and results presented.

6.8 Optimal cyclic operation of the candidate structure

This section describes how the optimal cyclic operation of the candidate structure (shown in 6.5) is determined using the dynamic optimisation formulation presented earlier in section 6.6. Recall the premise discussed earlier that for the dynamic system *instantaneous response is achievable through the use of product storage*. Based on this premise, a product storage tank is added to the optimal steady-state solution as shown in figure 6.5.




Figure 6.5: Candidate process for cyclic optimisation

6.8.1 Implementation

6.8.1.1 FORTRAN model generation

Using the unit models described in section 6.4 a complete model of the candidate process was developed within the gPROMS modelling environment. In addition to the underlying equations describing specific process units in the model, a number of ancillary equations were included. The additional equations provide information to be used by the optimiser in calculating values for the objective function and constraints. The total operating cost of the cycle (as given in equation 6.25) is determined by integrating the “per unit time” operating cost over the length of the operating cycle. To achieve this, a

new state variable, C_{op}^{total} is defined in the model as follows,

$$\frac{d}{dt}(C_{op}^{total}) = C_{op}(t) \quad (6.44)$$

where $C_{op}(t)$ is the operating cost per unit time. The optimisation formulation also requires constraints to be applied to the time derivatives of the state variables.

A set of algebraic variables must be defined from the time derivatives of the states as follows,

$$\begin{aligned} T_{dot}^1 &= k_1 \frac{d}{dt}(T_{out_1}) \\ T_{dot}^2 &= k_2 \frac{d}{dt}(T_{out_2}) \\ C_{dot}^1 &= k_3 \frac{d}{dt}(C_{out_1}) \\ C_{dot}^2 &= k_4 \frac{d}{dt}(C_{out_2}) \\ C_{product_{dot}} &= k_5 \frac{d}{dt}(C_{product}) \\ V_{dot}^1 &= k_6 \frac{d}{dt}(V_{R_1}) \\ V_{dot}^2 &= k_7 \frac{d}{dt}(V_{R_2}) \\ V_{store_{dot}} &= k_8 \frac{d}{dt}(V_{store}) \end{aligned} \quad (6.45)$$

where k_i is a scaling factor and T , C and V_R refer to temperature, concentration and holdup respectively. The subscript *dot* is used to indicate that the algebraic variable represents a time derivative, the subscripts 1 and 2 refer to reactors 1 and 2 and the subscript *store* refers to a property in the storage tank. By explicitly defining these variables, the corresponding sensitivity values (which are required by the optimisation algorithm) will be calculated by DASOLV during the integration process (see chapter 3).

The eight path constraints specifically required for the candidate structure, which were described in the optimisation problem formulation of section 6.6.4.1, are also added in the form described in section 3.5.3. After validating and testing the model within gPROMS, the system equations were generated in DASOLV format using the gPROMS “CODE” facility.

6.8.1.2 Specification of demand profile

Figure 6.2 in section 6.6 shows the required demand profile for the cyclic operation. This section briefly describes the way in which this profile is incorporated in the optimisation.

The time discretisation to be used for this problem (see section 6.6.3.2 for further details) may be summarised as follows. Product is required for the first six time intervals. it is ramped down to zero over the fixed seventh time interval and remains at zero for the

final six time intervals. To achieve this, the product demand was set up in the optimisation as a time-varying control variable (see section 3.5.3.1) in the optimisation input file. The breakpoint values for this control were then defined to give the required profile, but were excluded from the set of available optimisation parameters. In doing this, the correct product demand profile was superimposed on the operating cycle.

6.8.2 Results for the candidate structure

Having discussed the problem to be solved and its implementation, the results obtained will now be presented and reviewed.

Solution of the dynamic optimisation was initialised using the optimal steady-state point obtained in section 6.7.3. From this starting point, convergence of the dynamic problem to a specified accuracy of 1×10^{-5} was rapid, requiring only six cycles of the optimiser. This corresponds to a solution time of 20.6 CP minutes on a Sun Sparcstation 10. The economic optimum for the annualised cost of the dynamic system was found to be **668009** cost units (compared to a value from the earlier steady state optimisations of 649640 for 2 CSTRs and 667095 for 3 CSTRs). This optimum will be re-assessed in the discussion which follows.

Figures 6.6 and 6.7 illustrate examples of the behaviour of both the states and controls in the process over the course of an operating cycle and will now be discussed in turn. Figure 6.6 shows graphs of product demand, B concentration in the second reactor, the holdup of material in storage and the time derivative of the holdup in reactor one. The graphs demonstrate some key points.

- The trajectories of the two states shown (concentration and holdup) are indeed cyclic.
- Time t_1 , at which the product demand switches to zero occurs at $1.27mins$. Both at this time and at the final time (t_f) the time derivative shown is zero (as are the other time derivatives). Freezing the control variables at either of these points would therefore give a steady-state until such time as the control values are changed.
- The concentration state shown exhibits only very slight variations through the cycle and has an almost identical value to that determined from steady-state optimisation. This is true of all the states and will be considered further in the discussion which follows.
- The optimised initial value for the storage holdup is on its lower bound ($V_{store_{min}}$) and is almost constant during the cycle (the variation in this state is just 0.015

litres). In other words, the storage serves no significant purpose in the operation of this process. The small violation of the lower bound on the storage during the cycle is within tolerances placed upon the relevant path constraint in the problem specification.

The behaviour of the storage holdup shown in figure 6.6 would suggest that the storage tank is not needed to achieve feasible cyclic operation in this example. Further proof that this is the case can be seen in the state trajectory for concentration of B in the second CSTR. The purity of B in this CSTR does not violate product purity specifications at any time in the cycle. In other words, product of the required purity can feasibly be supplied at all times directly from the second CSTR.

In the light of this analysis, the economic optimum will now be reviewed. The fixed costs associated with the storage tank are subtracted from the optimal cost obtained previously. In doing this a revised optimum of **651342** is calculated. This figure is within 0.2% of the steady-state value. The optimum also lies below the steady-state optimum for 3 CSTRs. Note that if infinitesimal switching times were allowed in the control parameterisation this value would be identical to the steady-state economics. Based on the bounding properties discussed in section 6.7.1 no further structures need to be considered.

Figure 6.7 again shows the product demand as a reference, together with the catalyst controls, coolant flows and fresh feed flows. These graphs demonstrate two further important points:

- At the final time all controls are at zero. The earlier assumption of a zero cost endpoint in the cycle is therefore validated.
- When product is required, the controls take almost identical values to those determined in the steady-state optimisation. When no product is required, all controls are zero.

Table 6.6 summarises the values of all the states in the process at the shutdown steady-state value (corresponding to the initial and final condition in the cycle) and at the production steady-state. The values obtained in the steady-state optimisation are presented for comparison. The difference between the states in the plant at shutdown and during steady production are clearly negligible. The shutdown state of the plant that has been found in this case therefore represents an almost *perfect hold-state* which

has zero associated cost and gives the fastest possible transitions between the two steady-states.

Transition between the two states is achieved simply by switching the controls from zero to the value required for optimum steady-state production (as shown in figure 6.7). A shutdown state with zero cost is clearly desirable given that, in theory, the plant can remain idle for an indefinite period of time. Achieving the fastest possible transitions between steady-states is also an important factor in the design and will now be discussed. One of the assumptions in the operating cycle proposed in this work is that the process must complete the current transition between operating points before the reverse transition can be initiated. The implication here is that the process cannot start up during a shut down operation and, conversely, cannot shut down during a start up operation. Clearly, if the transitions are rapid the process offers greater flexibility to the customer.

6.8.3 Review of bounding assumptions

Earlier in section 6.7.1 the heuristic that the “optimal annualised cost of the steady-state process will be less than or equal to the cost of the dynamic process” was introduced as a means of screening design alternatives. The validity of this assumption will now be reviewed and a new bounding property introduced.

6.8.3.1 Bounding heuristic 1

Recall that the operating cycle determined by the optimisation above contains a perfect hold state. The presence of such a state allows the system to be switched near-instantly between non-production and production operating modes. In no case did dynamic optimisation produce a lower value of the objective than steady-state optimisation and the heuristic may be considered “not invalidated”. While bounding heuristic 1 seems appropriate for cyclic operation incorporating a steady-state period, it may not hold for other forms of operation, *e.g.* batch production.

6.8.3.2 Bounding property 2

A second bounding property may also be given as follows:

If a perfect hold state exists then switching instantly to the steady-state optimum operating point is feasible and the associated cost is an upper bound on all potential operating strategies.

It may therefore be inferred that:

where bounding heuristic 1 holds and a perfect hold state exists then the optimal solution is to instantly switch to the optimal steady-state.




Figure 6.6: Typical state profiles from cyclic optimisation of candidate structure




Figure 6.7: Control profiles from cyclic optimisation of candidate structure

<i>State</i>	<i>production state</i>	<i>shutdown state</i>	<i>Steady-state optimum</i>
CSTR 1: conc B	0.637	0.637	0.637
CSTR 2: conc B	0.817	0.817	0.817
Storage: conc B	0.817	0.817	-
CSTR 1: temp	48.892	48.885	48.889
CSTR 2: temp	48.892	48.884	48.889
CSTR 1:holdup	38.951	38.949	38.938
CSTR 2:holdup	22.281	22.279	22.298
Storage:holdup	1.126	1.139	-

Table 6.6: State variables at shutdown and production steady-states

In this section the ideal case of a process designed for cyclic operation has been established. The presence of a perfect hold-state in the operating cycle was clearly the key feature of the design which facilitated this. In the following section the concept of hold-states will be further explored under the effect of various non-idealities which will be introduced into the problem.

6.8.4 Cyclic operation in the presence of non-ideality

In this section the features of the process which give rise to the perfect hold-state above will be considered. Having identified such features, a small amount of non-ideality will be introduced and the resulting effect investigated.

In section 6.7 the feasibility of operating the process in four distinct modes was considered. One mode was to operate at steady-state with zero-production. Feasibility in this mode was determined from analysis of the steady-state form of the process model. It was stated that for this problem,

for the design to have a feasible shutdown steady-state one of either k , C_{out} or V_R must be moved to zero at that steady-state.

In the case of the ideal solution presented in the preceding section it is the rate of reaction k which is moved to zero at the shutdown steady-state. Since the rate of reaction is defined as,

$$k = \left[k_0 e^{\frac{-E_a}{R(T_{out}+460)}} \right] Catalyst_{control} \quad (6.46)$$

this is rapidly achievable by switching the catalyst controls between one and zero (as shown in figure 6.7). In effect, the catalyst can be used to turn the reaction on and off. When this is carried out simultaneously with the other controls in the process, all driving forces acting on the process states are effectively removed and the states become frozen

at their current values. Assuming that such perfect control over the chemical reaction can be achieved is an obvious oversimplification. However, it has enabled a limiting case to be identified together with the mechanisms by which it may occur.

Suppose now that the catalyst does not cleanly disengage from the reacting liquid so that a small level of the catalyst remains in the reactors giving rise to a low level of *background* reaction. The result of this is that perfect control over the reaction in the CSTRs is no longer attainable. The *background* reaction is modelled as follows,

$$k_{bg} = 5.4 \times 10^{11} e^{\frac{-20000}{460+T_{out}}} \quad (6.47)$$

the overall expression for the total reaction is therefore,

$$k_{new} = k_{bg} + k \quad (6.48)$$

where k is the original reaction rate given by equation 6.11. At the optimum reactor conditions found previously the ratio of catalysed to background reaction rate is of order 1000. Clearly, the presence of this background reaction means that k can no longer take a value of exactly zero when the catalyst is withdrawn. For a shutdown steady-state to be feasible with zero product flow either the reactor holdup or concentration must therefore be capable of moving to zero. By looking at the mechanisms driving these states it can be seen that this is not possible.

1. Reactor holdup

For small values of V_R the exit flow from a CSTR is proportional to V_R . V_R will therefore tend to zero as the period for which it drains tends to infinity (given that no fresh material enters the reactor).

2. Reactor concentration

Reaction is the only mechanism by which C_{out} can be reduced. The time variation of C_{out} is therefore proportional to the value C_{out} itself. As with the reactor holdup, C_{out} will tend to zero as the time for which reaction takes place tends to infinity (given that no fresh material enters the reactor).

It is clear that:

in the case when there is no mechanism to move the process to a feasible shutdown steady-state within a finite time the cyclic optimisation formulation is infeasible.

The cyclic optimisation framework presented in this work cannot, therefore, be considered general. In the section which follows the essential requirements of a more

general form of the problem are discussed. In addition, two more specific approaches are presented which directly address the infeasibility in this problem.

6.9 Optimal cyclic optimisation under uncertainty

As stated previously, the problem posed is infeasible when an exact steady-state does not exist. Suppose now that a set of the steady-state endpoint constraints (*i.e.* those imposed at time t_f) are relaxed by a small amount, ε , such that a feasible solution to the relaxed cyclic operation problem exists. The form of the relaxed constraint is,

$$-\varepsilon \leq \dot{x}_i \leq \varepsilon \quad (6.49)$$

where $\dot{x}_i$ is a state time derivative.




Figure 6.8: Drift in process states during idle periods for the relaxed endpoint problem

Typical state trajectories resulting from the solution of both the exact and relaxed endpoint problems are shown in diagrammatic form in Figure 6.8. Suppose that a feasible operating cycle (shown between t_0 and t_f in the figure) is determined by optimisation. At time t_f the process remains in an idle shutdown state for a period t_{idle} . Under the steady-state conditions imposed at t_f by the exact problem, the process states remain unchanged during the idle period and the same operating cycle may be repeated at time $t_{restart}$. In the case of the relaxed endpoint problem, the state time derivatives at t_f may take a value between $-\varepsilon$ and ε . The drift in these states at a given time can strictly² only be calculated by integrating the model equations through t_{idle} . However, because the length of t_{idle} is uncertain the values of the states at $t_{restart}$ will also be uncertain. Due to this uncertainty in the state values at $t_{restart}$ the feasibility of the original operating cycle cannot be guaranteed. A general solution to the cyclic operation problem, which may include relaxed steady-state endpoint constraints, would therefore require cyclic operation under uncertainty to be considered. Although the solution of this class of problems is not within the scope of this thesis a mathematical form of the problem will now be given.

As stated previously, the presence of an uncertain idle time gives rise to a set of possible states at $t_{restart}$ (see figure 6.8). The state of the plant at $t_{restart}$ then becomes the initial condition from which the plant must re-start the operating cycle. The cycle is repeated from this possible set of starting points x_r , again with an uncertain idle time. The form of a dynamic optimisation problem which incorporates uncertainty in the initial state values and idle time may be given as follows:

$$\begin{aligned} & \min_{d,p,x_0} \{J(d, x_0)\} & (6.50) \\ & s.t. : \\ & \begin{bmatrix} x_r \in \chi(p) \\ c(x_0, t_{idle}, d) \leq 0 \end{bmatrix} \\ & \forall x_0 \in \chi(p) \quad \forall t_{idle} \in [t_{idle}^l, t_{idle}^u] & (6.51) \end{aligned}$$

where d are the design variables, $\chi(p)$ is the set of possible initial states, p is the parameterisation of x_0 , t_{idle}^l and t_{idle}^u are bounds on the idle time, c is vector of constraints and J is annualised cost of the process. Equation 6.51 ensures feasible operation of the process for the set of possible initial conditions and idle times. The solution of such a problem has not been considered in this work. It should be noted that the original problem posed

²for sufficiently small values of ε and t_{idle} a linear projection may be accurate enough to predict the state values at $t_{restart}$

in this chapter is a specific form of the problem described above, *i.e.*:

$$\min_{d, x_0} \{J(d, x_0)\} \quad (6.52)$$

s.t. :

$$\begin{bmatrix} x_r = x_0 \\ c(x_0, d) \leq 0 \end{bmatrix} \quad (6.53)$$

In the sections which follow, two solutions to the non-ideal case with background reaction are presented. In both cases the steady-state endpoint is redefined to give a form of the uncertainty which is more easily incorporated into the existing cyclic optimisation framework and results in a more tractable problem.

6.9.1 Bounding the uncertainty using near zero concentrations at t_f

6.9.1.1 Background

Recall that a feasible steady-state shutdown requires either C_{out} , V_R or k to be moved to zero and that no mechanisms exist in the process to achieve this in finite time. The CSTR reactant concentrations may, however, be driven to a value close to zero in a finite time.

Consider the following: if the process material balance is at steady-state at shutdown (which is easily achieved by ensuring zero flows through the process), the state time derivatives for the concentration and temperature are defined as,

$$\frac{dC_{out}}{dt} = -kC_{out} \quad (6.54)$$

$$\frac{dT_{out}}{dt} = -kC_{out} \frac{(\Delta H_R)}{\rho C_p} \quad (6.55)$$

Provided C_{out} is sufficiently small,

1. the values of the state time derivatives will also be small,
2. the values of the states themselves will remain tightly bounded.

Equation 6.54 and 6.55 may be integrated to give explicit expressions for C_{out} and T_{out} at a given period after t_f ,

$$C_{out} = C_0 e^{-kt} \quad (6.56)$$

$$T_{out} = T_0 + \left\{ \frac{C_0 \Delta H_R}{\rho C_p} \right\} [e^{-kt} - 1] \quad (6.57)$$

where the subscript $_0$ refers to the value of the state at time t_f . Using the asymptotic values of these expressions allows C_{out} and T_{out} to be bounded as follows,

$$0 \leq C_{out} \leq C_0 \quad (6.58)$$

$$T_0 \leq T_{out} \leq \left[T_0 - \left\{ \frac{C_0 \Delta H_R}{\rho C_p} \right\} \right] \quad (6.59)$$

Providing C_0 is chosen to be sufficiently small, the values of temperature and concentration at $t_{restart}$ can be very tightly bounded, in effect removing any uncertainty in them.

6.9.1.2 Problem Reformulation

Based on the discussion above, the $\dot{x}(t_f) = 0$ endpoint constraints applied to the solution of the ideal case (see section 6.8.2) are reformulated and implemented as follows:

$$\begin{aligned} \dot{C}_{product}(t_f) &= 0 \\ \dot{V}_{R_1}(t_f) &= 0 \\ \dot{V}_{R_2}(t_f) &= 0 \\ \dot{V}_{store}(t_f) &= 0 \end{aligned} \quad (6.60)$$

The steady-state constraints on temperature and concentration are removed and replaced with the following,

$$\begin{aligned} C_{out_1}(t_f) &\leq 1 \times 10^{-3} \\ C_{out_2}(t_f) &\leq 1 \times 10^{-3} \end{aligned} \quad (6.61)$$

where the subscripts 1 and 2 refer to the first and second reactors. Note that $C_{out_i}(t_f)$ is equivalent to C_0 used above. A value for C_0 of 1×10^{-3} was considered close enough to zero to give tight bounding of the states, whilst avoiding excessively long cycle times.

6.9.1.3 Results

Solution times for the reformulated problem were greater than for the ideal case, the problem requiring some 550 CP minutes on a Sparcstation 10. The total annualised cost of the new system has risen to **726065** cost units, which represents an increase of 12% over the ideal base system. The reasons for this increase are twofold:

1. capital cost increases caused by the large storage requirement,
2. operating cost increases caused by the use of additional catalyst and coolant during the prolonged shutdown phase.

The actual production steady-state point determined by the optimisation is similar to that for the ideal case found previously. This would tend to suggest that for the values of $t_{product}$ and N_{cycle} used in the problem description, the economics of the production steady-state dominate the operating cost. Figures 6.9 and 6.10 show the optimised operating cycle resulting from solution of the reformulated problem. The product demand switches to zero at a time of 13.8 minutes. This time corresponds to t_1 in the operating cycle (the end of the startup phase). Notice that during this time the storage tank is drained of product, reaching a minimum value at t_1 (see figure 6.10). The time to move the process to the shutdown state ($t_f - t_1$) is 21.2 minutes. During this time the storage tank is refilled. Crucially, the incorporation of intermediate product storage has enabled product to be delivered on demand during periods when the process is in transition. Recall from section 6.8.2 that in the cyclic operation of the process, startup is not permitted during the shutdown phase and, conversely, shutdown is not permitted during the startup phase. The impact of the extended startup and shutdown periods needed in this (more expensive) solution is therefore a reduction in the flexibility of the process. From an optimisation point of view, however, one distinct advantage of this reformulation is that it remains a feasible approach for increasingly large values of the background reaction rate (provided C_0 is sufficiently small).




Figure 6.9: State profiles from solution of relaxed problem - 1




Figure 6.10: State profiles from solution of relaxed problem - 2

6.9.2 Propagate all uncertainty into one state

In the first approach presented, uncertainty in the states arising from an inexact steady-state at t_f was effectively removed by forcing the reactant concentrations to near-zero at t_f . In the second approach described here, the uncertain state drift arising from the relaxed endpoint problem is propagated into just one state. The drift in this state is shown to be predictable and a simple mechanism for updating the cyclic operation to take account of the drift is found.

Consider the process decoupled into two parts as shown in figure 6.11,

1. a processing stage,
2. a storage stage.

When reaction takes place, steady-state in the processing stage can only be achieved by the continuous feed of fresh reactant to the reactors. Here it is converted to B before being discharged as product. In the storage stage, a zero product flow constraint at t_f can only be achieved if no material enters this stage. Clearly, when these two stages are coupled the resulting problem is infeasible! Suppose now that the steady-state requirement on the storage holdup at t_f is relaxed. This effectively decouples the steady-state constraints on the processing unit from the zero product flow requirement at time t_f . In other words, during the idle period the steady-state of the process is maintained by processing A at a background level and using the storage tank to accommodate the product B . The uncertain drift in the process states described previously is propagated through solely to the storage holdup. Further, because the processing stage is at steady-state during the idle period the drift in the storage holdup is linear with respect to time. Relaxation of this one constraint enables a feasible cyclic operation to be determined. Once this basic strategy is in place, correction of the hold state following an idle period is trivial. Figure 6.12 shows how the basic operating cycle is annotated to correct the storage holdup. The operating cycle runs between t_0 and t_{cycle} as before. At t_{cycle} the process is in an idle shutdown-state with some B produced which is stored. Product holdup in storage increases linearly until time $t_{restart}$ when the customer again demands product. The required product is initially supplied from storage with the next process operating cycle being delayed until the storage level matching that at t_0 is reached. Given that at t_f the holdup in storage is $V_{store}(t_f)$, after an idle time of t_{idle} the holdup will be,

$$V_{store}(t_f + t_{idle}) = V_{store}(t_f) + t_{idle}Q_{product}^{bg} \quad (6.62)$$

where $Q_{product}^{bg}$ is the volumetric production of B at the background rate. The drain time between $t_{restart}$ and t_0 is then easily calculable from,

$$t_{drain} = \frac{t_{idle} Q_{product}^{bg}}{F_{product} - Q_{product}^{bg}} \quad (6.63)$$

Note that during the draining period the process continues to produce B at a rate of $Q_{product}^{bg}$.

6.9.2.1 Application

The procedure is now demonstrated to obtain an estimate of the cost of such an operating cycle for this case study. The results are based upon a mean annualised idle time of 12 hours.

Cyclic optimisation of the non-ideal process using the relaxed storage endpoint produced similar operating profiles and steady-state operating points to the ideal case. The control values at t_f obtained from dynamic optimisation required to maintain a steady idle state are shown in table 6.7.

<i>Control</i>	<i>Value</i>
CSTR 1 - fresh feed	1.25×10^{-2}
CSTR 2 - fresh feed	0.00
CSTR 1 - Coolant flow	0.00
CSTR 2 - Coolant flow	3.57×10^{-3}

Table 6.7: Control values during idle period

During the idle period product accumulates in the storage tank at a rate of $1.25 \times 10^{-2} \text{ litres.min}^{-1}$. Thus over the whole idle period the product accumulated is 9 litres. Given that the product demand is $12.3 \text{ litres.min}^{-1}$ and applying equation 6.63, the drain period is calculated as 0.73 minutes. This time is insignificant as a proportion of the total cycle time ($\sim 720 \text{ minutes}$). The principal cost associated with the revised operating cycle is the need for a larger storage tank, the annualised cost of which is 24008 cost units (representing a 4% increase over the cost of the ideal case).

Although this method can be used to quickly correct the operating strategy and determine associated costs, the economic implications of the added idle and drain phases of operation are not incorporated at the optimisation stage and the solution must therefore be considered sub-optimal. The final section in this chapter considers the incorporation of this more general operating cycle into a formal cyclic optimisation framework.




Figure 6.11: Decomposition of process structure




Figure 6.12: Annotated cyclic operation

6.9.3 Cyclic optimisation formulation incorporating a relaxed storage endpoint

This section describes the way in which the optimisation formulation is revised to incorporate the effects of the idle and drain operations shown in figure 6.12 when determining the optimal cyclic operation. In the original formulation of the cyclic optimisation problem an expected value for the production period was used to obtain an expected annualised cost. In order to consider the effects of the idle state a mean idle period (t_{idle}) is defined here and is denoted by t_{idle} .

6.9.3.1 Constraints

All constraints are applied as described in section 6.6.4, with two revisions:

1. the steady-state endpoint on storage holdup is relaxed (as discussed earlier),

$$0 \leq \dot{V}_{store}(t_f) \leq \epsilon \quad (6.64)$$

where ϵ is a constant.

2. an additional constraint is placed on V_{store}^{max} such that,

$$V_{store}^{max} \geq V_{store}(t_f) + t_{idle}^{max} \dot{V}_{store}(t_f) \quad (6.65)$$

which ensures that the storage tank is large enough to accommodate product generated for the maximum length of idle period (t_{idle}^{max}).

6.9.3.2 Objective function

The objective function also needs to be modified to correctly represent the costs associated with the operating cycle shown in figure 6.12. The capital cost function remains unchanged from equation 6.23, while constraint 6.65 ensures that the correct V_{store}^{max} is used in calculating the capital equipment cost.

A term for the cost of the idle period of operation must be added to the operating costs given in equation 6.25. Each CSTR has an associated operating cost at the idle state of:

$$C_{CSTR_{op}}^{idle} = (7.063 \times 10^{-2} F_{in}^2 t_f + 3.532 \times 10^{-4} F_c(t_f))(t_{idle}) \quad (6.66)$$

Lastly, the cost of the draining period must be accounted for in the operating cost. Let the period for which storage is drained be represented as t_{drain} (which is determined using equation 6.63). Given that there is a fixed total production period of $t_{product}$, the

steady-state production period ($t_{product} - t_1$) must be reduced by an amount t_{drain} . The total operating cost for each CSTR per cycle is then:

$$C_{CSTR_{op}} = \int_0^{t_f} (7.063 \times 10^{-2} F_{in}^2(t) + 3.532 \times 10^{-4} F_c(t) + 1.766 \times 10^{-2} V_R^{max} \cdot Catalyst_{control}(t)) dt \\ + (7.063 \times 10^{-2} F_{in}^2(t_1) + 3.532 \times 10^{-4} F_c(t_1) + 1.766 \times 10^{-2} V_R^{max} \cdot Catalyst_{control}(t_1)) (t_{product} - t_1 - t_{drain}) + C_{CSTR_{op}}^{idle}$$

Cyclic optimisation of the non-ideal case using the above formulation was not carried out. It is expected, however, that the solution of this problem will result in a lower annualised cost for the process than that the sub-optimal solution calculated in section 6.9.2.1.

6.10 Discussion

The work in this chapter has considered the design and optimal cyclic operation of a reactor network. In this section the key features and results of this work will be reviewed and discussed.

6.10.1 Summary of key results

- The problem definition presented was based around the requirements of the customer. Consideration of these requirements resulted in the possibility of producing B on-demand, *in situ* being examined.
- A feasible solution to the problem was found using a two stage design procedure. The two stages were essentially: development of a candidate structure followed by optimisation of the operating cycle and design of this structure.
- A steady-state bounding heuristic was applied to efficiently screen design alternatives for cyclic operation through a steady-state production phase (6.8.3.1):

The optimal annualised cost of the steady-state process will be less than or equal to the cost of the dynamic process .

- The base-case solution exploited a perfect hold-state in the system which essentially enables the process to be moved instantly to the optimal steady-state operating point.
- A bounding property (6.8.3.1) was then introduced:

If a perfect hold state exists then switching instantly to the steady-state optimum operating point is feasible and the associated cost is an upper bound on all potential operating strategies.

- Reasons for the existence of a perfect hold-state were identified and the cyclic operation of the candidate structure in the presence of non-ideality was considered. This was achieved by inhibiting the perfect control over the reaction which was previously the case. This rendered the original optimisation problem infeasible.
- Recovering feasibility in the problem required a relaxation of the steady-state endpoints. This led to the problem being re-posed as a cyclic optimisation under uncertainty. Two simplified approaches to the problem with uncertainty were presented.

6.10.2 Generality of results

In posing the problem to be solved a number of degrees of freedom in the problem were fixed *a priori*:

- the reaction scheme,
- the type of reactor unit,
- the connectivity of the reactor units.

Specifically, a reactor network consisting of a series of continuous stirred tank reactors is used to carry out a simple first-order exothermic reaction producing B from A . The impact of these assumptions on the generality of the results presented above in section 6.10.1 will now be considered.

The general design problem to be solved was stated in section 6.5 as the determination of the economically optimal combination of:

1. unit type and topology,
2. unit design,
3. cyclic operation of the complete process.

A priori specification of the degrees of freedom given above means that the problem solved in the chapter was to find the optimal:

1. number of units,
2. unit designs,

3. cyclic operation of the complete process.

The key question here is whether the generality of the results presented in this chapter is in any way compromised by the simplifications made to the problem.

6.10.2.1 Development of candidate structures

In section 6.7 four modes of operation for the final process design were identified:

1. steady-state production,
2. steady-state with zero production,
3. transition from mode 2 to 1,
4. transition from mode 1 to 2.

The condition was given that the final design must be feasible for each of these operational modes (a requirement which was shown to be relatively simple to assess). This condition must clearly be satisfied for all processes which are required to operate within the cyclic framework defined for this problem. The existence of feasible steady-states is straightforward to establish given that a model of the system being considered is available. An underlying assumption in this work is that given sufficient storage of each of the required products, the transitional modes of operation will be feasible. The generality of this assumption cannot be guaranteed for all cases, but has not been proved invalid by the work presented.

Bounding heuristic 6.8.3.1 has been shown to hold for the cases considered. In the context of the cyclic operation framework used in this case study the heuristic seems appropriate and is potentially applicable to any process which operates within such a cycle.

6.10.2.2 Cyclic optimisation

The cyclic optimisation framework used in this case study is specific only in the sense that it defines an operating strategy which addresses the needs of the customer. The form of the cyclic optimisation itself is independent of the structure of the process and is therefore not restricted by the assumptions made regarding the type and topology of the units to be used.

Bounding property 6.8.3.2 is a general bounding property which applies to the *type* of operating strategy and is independent of the structure of the process.

6.10.3 Alternative operating strategies

As was seen above, a number of the results obtained may be considered general only in the context of the form of cyclic operation used for this work. A specific feature of these operating cycles is that they pass through a period of steady production. As a conclusion to this chapter, potential benefits of using a batch operating strategy rather than continuous³ production are briefly considered.

In order to sensibly compare the two operational strategies, the batch process will be designed under the same conditions as the “continuous” process, namely:

- production of B will be achieved through the same reaction used previously,
- the same operating constraints on the maximum operating temperature and maximum equipment sizes apply,
- utilities available to the batch reactor are the same as those used in the CSTRs (cooling water and catalyst).

The nominal form of the batch operating cycle is assumed to be as follows. At time t_0 the customer demands product. At this time the batch reactor is empty and the storage tank is full. Product is drained from storage while the reactor is charged with fresh A and reaction commences. At some time in the future, the storage tank will be exhausted and the batch will be completed such that the storage tank may be refilled to continue the supply of product (if more product is required).

Clearly, for such a strategy to be feasible the time required to produce a batch of product must be less than the time it takes to drain the storage tank of stored product. Recall that the maximum inventory allowable in a given vessel is 50 *litres*. At the required product demand rate of 12.3 *litres.min*⁻¹ this corresponds to approximately 4 minutes worth of product supply. So, a batch time of less than 4 minutes is needed for this operating strategy to be feasible.

The optimal cooling profile to give the minimum batch time may be determined from a simple dynamic optimisation. The profile was found for a 50 *litre* reactor and is essentially a bang-bang control (which is consistent with a minimum time problem). Zero cooling is used for 7.26 minutes to give the most rapid change in concentration possible, followed by 1.38 minute period when a coolant flow of 0.40 *litres.min*⁻¹ is applied to prevent violation of the maximum temperature bound until the batch is complete. The minimum batch time is therefore 8.64 minutes. It is important to note here that the

³in the sense that steady-state production occurs in the cycle

concentration and temperature profiles over the period with zero cooling are independent of the reactor volume. Since this period accounts for some 85% of the batch time, using a different size reactor will not significantly alter the minimum batch time.

The nominal batch cycle described above is therefore infeasible. In fact three storage vessels (each of volume 35.9 *litres*) will be needed to supply the product until the batch reactor cycle is complete. A number of batch reactors will be required to operate in parallel to recharge each storage tank at the end of the cycle. A feasible operating strategy using a parallel arrangement incorporating three sets of reactors and storage vessels in series is shown in figure 6.13. The three batch reactors are charged simultaneously at t_0 and reaction proceeds. Product is drawn from each storage vessel in turn until the batch is complete at time t_{dump} , when the contents of each reactor is dumped into its corresponding storage vessel. The cycle may then be repeated. The capital cost implications of such a system are very significant as it requires double the number of equipment items as the continuous system. No factors are apparent which would significantly change the operating cost.

In summary, a feasible design and cyclic operating strategy based around a steady-state period of operation has been presented. This result was then compared to an alternative batch operating strategy. Results using a simple analysis of the batch problem showed that, for this case, a batch operation would not give economic improvements over a cycle using a steady-state period of production.




Figure 6.13: Parallel arrangement for batch operating strategy

Chapter 7

Conclusions and future work

This thesis has considered the development and application of an integrated design procedure for the design and operation of distributed chemical manufacturing facilities. In this chapter the key results of this work will be summarised and conclusions drawn. Finally, areas for future work are identified.

7.1 Results and conclusions

7.1.1 Development of the design procedure

The development of the design procedure used in this work is outlined below.

- The concept of a distributed chemical manufacturing facility was introduced and shown to be an inherently dynamic process. This provided the motivation for an integrated approach to the design of such plants.
- A review of the literature identified two integrated design approaches; controllability about steady-state and explicit consideration of the dynamic operation in the process design. It was concluded that the second approach was appropriate for this work.
- A two stage decomposition of the problem was proposed as the basis for the design procedure. Candidate structures are first developed using heuristics and economic bounding. Dynamic optimisation is then applied at the second stage to determine optimal design and operation of the candidate.
- Cyclic operation enabled infinite time horizon profiles to be effectively captured in a finite time optimisation problem. Single disturbances were considered with the disturbance being representative of the perceived worst case.

- The design procedure developed in chapter 3 has been successfully applied to three case studies. In all the cases considered, the dynamics and operation of the process formed an integral part of the design problem.

7.1.2 Case studies

The design procedure developed was applied successfully to three case studies. The main results for each of these are now summarised.

7.1.2.1 Continuous multiproduct plant

The dynamic optimisation formulation for optimal design and cyclic operation was effectively demonstrated for a process which is typically operated cyclically in industrial practice. Substantial economic improvements over alternative published solutions were obtained. This was achieved by considering the dynamics of the complete system (*i.e.* process and product storage). The optimal reactor response was highly oscillatory and closed-loop unstable, but remained bounded in the finite cycle time determined in the optimisation. A damped controller response was achieved by minimisation of its integral square error subject to a specified maximum increase in the best-case economics.

7.1.2.2 Integrated wastewater treatment plant

The integrated design and operation of the diversion and treatment systems typically found in chemical wastewater was considered. Application of an integrated approach enabled substantial reductions in cost and overall size to be achieved as compared to a conventional system. The proposed design procedure was effectively applied to rapidly find the optimal solution. The approach was flexible enough to enable existing screening techniques, new economic bounding properties and dynamic optimisation to be applied in a systematic way. Analysis of the results allowed a number of qualitative conclusions to be drawn about the general form of the integrated wastewater systems developed.

7.1.2.3 Distributed reactor network

A plant capable of instantaneously supplying an intermittent demand for product was developed. Again the proposed design procedure was efficiently applied. A candidate structure was quickly found using steady-state optimisation and a heuristic on the achievable economics of the dynamic system. The ultimate design produced incorporated a perfect hold state which enabled near-instant switching from the shutdown

state to the steady-state optimum operating point. An upper economic bound followed from this. A more general cyclic operation problem with uncertainty was introduced for the non-ideal case. Two simplifications to the problem enabled a limited solution to this problem. Batch operation was considered as an alternative operating strategy. Using simple qualitative arguments, this system was found to be more costly.

7.1.3 General conclusions

A number of general conclusions may be made based upon the case studies considered in this thesis.

- The design procedure has enabled viable candidates to be quickly developed and solutions to be obtained with minimal computational effort.
- Heuristics and economic bounding rules were straightforward to derive and were shown to be effective in reducing the design search space.
- The procedure is sufficiently flexible to enable existing screening tools to be applied where appropriate.
- Decomposition of the problem to separate structural decisions from unit design and operation gives a tractable second stage optimisation problem for which robust and proven solution techniques exist.

7.2 Areas for future work

The final section in this thesis considers areas where opportunities for further work exist.

- The wastewater treatment case study has demonstrated that substantial reductions in cost and size are possible if diversion and treatment systems are integrated. The use of optimal open loop controls in the diversion system and a limited uncertainty description (a worst case concentration disturbance) makes the results obtained somewhat idealised. It is considered that sufficient opportunities exist in this area to justify further work. Useful extensions would be the incorporation of more general feedback control laws in the diversion system, and a more comprehensive treatment of likely uncertainty in process parameters and disturbances.
- The optimisation framework used in this work has proved to be applicable on a range of problems. It is, however, fairly idealised in requiring point to point cycling of the

process states. The formulation should be extended to a more general operating cycle which considers feasible *regions* for the initial and final process states.

- The multiproduct plant case study brings out the problem of conflicting objectives. In this case, some sacrifice in the economics had to be made to achieve a more damped system response. A systematic approach should be developed to accommodate circumstances where conflicting objectives occur. One possible systematisation of design and operation objectives might be:
 1. ensure feasibility,
 2. find best economics,
 3. given best economic solution, can worthwhile improvements be obtained for a given penalty in the best economics?
- The optimal multiproduct plant design developed was inherently unstable in its operation. The solution was feasible because the operation was cyclic with frequent changes in operating point. Other processes may exist which can be operated in a similar way to reap economic benefits and is a potential area for further investigation.

It is hoped that this initial foray into the design and operation of distributed chemical manufacturing facilities will stimulate further research and enable the potential benefits such processes to be more widely realised in the future.

References

- A. Abbas, P.E. Sawyer, and P.L. Yue. Integrated design and control of a continuous stirred tank reactor. *Trans IChemE*, 71(A):453–456, 1993.
- P.I. Barton. *The Modelling and Simulation of Combined Discrete/Continuous Processes*. PhD thesis, Imperial College, 1992.
- R.S. Benson and J.W. Ponton. Process miniaturisation — A route to total environmental acceptability? *Trans IChemE*, 71(Part A):160, March 1993.
- L. T. Biegler. On the simultaneous solution and optimization of large scale engineering systems. *Computers and Chemical Engineering*, 12(5):357, 1988.
- M.L. Brisk. Teaching reactor design and control. *Chemical Engineering - II*, page 13, 1974.
- A.E. Bryson and Y.C. Ho. *Applied Optimal Control*. Halsted Press, 1975.
- J.G. Cantrell and W.L. Luyben. Design of rapid transition multiproduct processes. *Ind. Eng. Chem. Res.*, 35:2624–2631, 1996.
- C. L. Chen and S. Macchietto. SRQPD Version 1.1 – User’s Guide. Technical report, University of London, September 1991.
- C.A. Crooks and S. Macchietto. Synthesis and validation of operating procedures for chemical plants. Technical report, Imperial College, 1990. Process Systems Engineering Industrial Research Consortium.
- C.A. Crooks, N. Shah, C.C. Pantelides, and S. Macchietto. A combined milp and logic-based approach to the problem of operating procedures for batch plants. Technical report, Imperial College, 1990. Process Systems Engineering Industrial Research Consortium.

- G.C.C. Cushnie. *Removal of Metals from Wastewater. Neutralization and Precipitation*. Noyes, 1984.
- S. de Hennin. Structural decisions in on-line optimisation - annual report. Technical report, Imperial College, 1992.
- I.T. Dedopoulos and N. Shah. Optimal short-term scheduling of maintenance and production for multipurpose plants. *Ind. Eng. Chem. Res.*, 34(1):192–201, 1995.
- V.D. Dimitriadis and E.N. Pistikopoulos. Flexibility analysis of dynamic systems. *Ind. Eng. Chem. Res.*, 34(12):4451–4462, 1995.
- J.M. Douglas. *Conceptual Design of Chemical Processes*. McGraw-Hill, 1988.
- T.R. Elliot and W.L. Luyben. Capacity-based economic approach for the quantitative assessment of process controllability during the conceptual design stage. *Ind. Eng. Chem. Res.*, 34(11):3907–3915, 1995.
- R.W. Fusillo and G.J. Powers. A synthesis method for chemical plant operating procedures. *Computers Chem. Engng*, 11(4), 1987.
- R.W. Fusillo and G.J. Powers. Operating procedure synthesis using local models and distributed goals. *Comput. chem. Engng*, 12(9/10), 1988.
- A. Georgiou and C. A. Floudas. Simultaneous process synthesis and control: Minimisation of disturbance propagation in heat recovery systems. In *FOCAPD*, 1990.
- G.F. Froment and K.B. Bischoff. *Chemical Reactor Analysis and Design*. Wiley, 1990.
- P.E. Gill, W. Murray, and M.H. Wright. *Practical Optimization*. Academic Press, 1981.
- I.E. Grossmann and D.A. Straub. *Recent Developments in the Evaluation and Optimization of Flexible Chemical Processes*. Elsevier, Amsterdam, 1991. p. 49.
- F. Hanssmann. *Operations Research in Production and Control*. Wiley, 1962.
- J. Heath. Interactions between design and control - annual report. Imperial College, 1993.
- J.A. Heath. *Process Control Structure Selection Based on Optimisation*. PhD thesis, Imperial College, 1997.
- J. M. Howell. *The Effect of Energy Integration on the Switchability of Chemical Processes*. PhD thesis, University of London, 1984.

- J.F. Van Impe, B.M. Nicolai, P.A. Vanrolleghem, J.A. Spiet, B. de Moor, and J. Vandewalle. Optimal control of the penicillin g fed-batch fermentation: An analysis of the model of Heijnen et al. *Optimal Control Applications and Methods*, 15:13–34, 1994.
- V.A. Ivanov, V.V. Kafarov, V.L. Perov, and A.A. Reznichenko. Design principles for chemical production startup algorithms. *Automn. Remote Control.*, 41, 1981.
- V.A. Ivanov, V.V. Kafarov, V.L. Perov, and A.A. Reznichenko. On algorithmization of the start-up of chemical productions. *Engng. Cybernet.*, 18, 1981.
- R.B. Jarvis and C.C. Pantelides. *DASOLV, User Manual, Version 1.0*. Imperial College, 1992.
- R.B. Jarvis. *Robust Dynamic Simulation of Chemical Engineering Processes*. PhD thesis, Imperial College, 1993.
- L.M. Király, F. Friedler, and Szoboszlai. Optimal design of multipurpose batch chemical plants. *Computers Chem. Engng*, 13(4/5):527–534, 1989.
- D.E. Kirk. *Optimal Control Theory - An Introduction*. Prentice-Hall, 1970.
- Ch. Kruse, G. Fieg, and G. Wozny. A new time-optimal strategy for column startup and product changeover. *J. Process Cont.*, 6(2/3):187–193, 1996.
- J. Lee and W.F. Ramirez. Optimal fed-batch control of induced foreign protein production by recombinant bacteria. *AIChE Journal*, 40(5), May 1994.
- B. Lee and G.V. Reklaitis. Optimal scheduling of cyclic batch processes for heat integration - I. Basic formulation. *Computers Chem. Engng*, 19(8):883–905, 1995.
- B. Lee and G.V. Reklaitis. Optimal scheduling of cyclic batch processes for heat integration - ii. extended problems. *Computers Chem. Engng*, 19(8):907–931, 1995.
- A.M. Lenhoff and M. Morari. Design of resilient processing plants—I: Process design under consideration of dynamic aspects. *Chemical Engineering Science*, 37(2):245, 1982.
- E.S. Leversund, S. Macchietto, G. Stuart, and S. Skogestad. Optimal control and on-line operation of reactive batch distillation. *Computers chem. Engng.*, 18(Suppl.), 1994.
- J. S. Logsdon and L. T. Biegler. A relaxed reduced space SQP strategy for dynamic optimization problems. *Computers and Chemical Engineering*, 17(4):367, 1993.

- M.L. Luyben and C.A. Floudas. Analyzing the interaction of design and control-1. A multiobjective framework and application to binary distillation synthesis. *Comput. Chem. Engng*, 18(10):933–969, 1994.
- M.L. Luyben and C.L. Floudas. Analyzing the interaction of design and control-2. Reactor-separator-recycle system. *Comput. Chem. Engng*, 18(10):971–994, 1994.
- D.M. Marselle, M. Morari, and D.F. Rudd. Design of resilient processing plants–II. Design and control of energy management systems. *Chem. Eng. Sci.*, 37(2):259–270, 1982.
- K.B. McAuley and J.F. MacGregor. Optimal grade transitions in a gas phase polyethylene reactor. *AIChE Journal*, 38(10):1564–1576, October 1992.
- P. Mizsey and Z. Fonyo. A predictor-based bounding strategy for synthesizing energy integrated total flowsheets. *Comput. Chem. Engng*, 14(11):1303–1310, 1990.
- P. Mizsey and Z. Fonyo. Towards a more realistic overall process synthesis - The combined approach. *Comput. Chem. Engng*, 14(11):1213–1236, 1990.
- M.J. Mohideen, J. Perkins, and E.N. Pistikopoulos. Optimal design of dynamic systems under uncertainty. *AIChE Journal*, 42(8):2251–2272, 1996.
- M. Morari and J.D. Perkins. Design for operations. In *Proc. FOCAPD 94*, Snowmass, CO, 1994.
- M. Morari. Design of resilient processing plants–III: A general framework for the assessment of dynamic resilience. *Chem. Engng Sci.*, 38(11):1881, 1891 1983.
- M. Morari. Flexibility and resiliency of process systems. *Comput. Chem. Engng*, 7(4):423–437, 1983.
- M. Morari. Effect of design on the controllability of chemical plants. In *Interactions Between Process Design and Process Control, IFAC Workshop, London, UK*. IFAC, 1992.
- L. T. Narraway. *Selection of Process Control Structures Based on Economics*. PhD thesis, University of London, 1992.
- N. Nishida and A. Ichikawa. Synthesis of optimal dynamic process systems by a gradient method. *Ind. Eng. Chem., Process Des. Dev.*, 14(3):236, 1975.

- N. Nishida, Y.A. Liu, and A. Ichikawa. Studies in chemical process design and synthesis - ii. optimal synthesis of dynamic systems with uncertainty. *AIChE Journal*, 22(3):539–548, 1976.
- E. O'Shima. Safety supervision of valve operations. *J. Chem. Engng Japan*, 11, 1978.
- S. Palanki, C. Kravaris, and H.Y. Wang. Optimal feedback control of batch reactors with a state inequality constraint and free terminal time. *Chemical Engineering Science*, 49(1), 1994.
- A. Palazoglu and Y. Arkun. A multiobjective approach to design chemical plants with robust dynamic operability characteristics. *Comput. Chem. Engng*, 10(6):567–575, 1986.
- K.P. Papalexandri and E.N. Pistikopoulos. Synthesis and retrofit design of operable heat exchanger networks. 1. Flexibility and structural controllability aspects. *Ind. Eng. Chem. Res.*, 33(7):1718–1737, 1994.
- K.P. Papalexandri and E.N. Pistikopoulos. Synthesis and retrofit design of operable heat exchanger networks. 2. Dynamics and control structure considerations. *Ind. Eng. Chem. Res.*, 33(7):1738–1755, 1994.
- J.M. Pinto and I.E. Grossmann. Optimal scheduling of multistage continuous multiproduct plants. In *Proceedings*, Mount Crested Butte, CO, 18-23 June 1993. 397-402.
- G.P. Pollard and R.W.H. Sargent. Off line computation of optimum controls for a plate distillation column. *Automatica*, 6, 1970.
- Prosys. *Speedup User Manual, volume 1*, July 1991.
- Prosys. *Speedup User Manual, volume 2*, July 1991.
- A. Reeve. Batch control, a recipe for success ? *Process Engng*, 73(1):33–34, 1992.
- G.V. Reklaitis. Progress and issues in computer-aided batch process design. In *FOCAPD Proceedings on Foundations of Computer Aided Process Design*, New York, 1990. Elsevier. 241-275.
- D.W.T. Rippin. Simulation of single and multiproduct batch chemical plants for optimal design and operation. *Computers Chem. Engng*, 7(3):137–156, 1983.
- D.W.T. Rippin. Batch process systems engineering: A retrospective and prospective view. *Computers Chem. Engng*, 17 (suppl.):S1–S13, 1993.

- J.R. Rivas and D.F. Rudd. Synthesis of failure-safe operations. *AIChE Journal*, 20(2), 1974.
- J.R. Rivas, D.F. Rudd, and L.R. Kelly. Computer aided safety interlock systems. *AIChE Journal*, 20(2), 1974.
- N.V. Sahinidis and I.E. Grossmann. Minlp model for cyclic multiproduct scheduling on continuous parallel lines. *Computers Chem. Engng*, 15(2):85–103, 1991.
- R. W. H. Sargent and G. R. Sullivan. Development of feed changeover policies for refinery distillation units. *Ind. Eng. Chem. Process Des. Dev.*, 18(1):113, 1979.
- D.E. Seborg, T.F. Edgar, and D.A. Mellichamp. *Process Dynamics and Control*. Wiley, 1989.
- N. Shah and C.C. Pantelides. Design of multipurpose batch plants with uncertain production requirements. *Ind. Eng. Chem. Res.*, 31:1325–1337, 1992.
- E. Sørensen and S. Skogestad. Optimal operating policies of batch distillation with the emphasis on the cyclic operating policy. In E.S. Yoon, editor, *Proceedings of the 5th International Symposium on Process Systems Engineering*, volume 1, 1994. 449-456.
- R.E. Sparrow, G.J. Forder, and D.W.T Rippin. The choice of equipment sizes for multiproduct batch plants. heuristics vs. branch and bound. *Ind. Eng. Chem., Process Des. Dev.*, 14(3):197–203, 1975.
- G. Stephanopoulos. Synthesis of process flowsheets: an adventure in heuristic design or utopia of mathematical programming? In *FOCAPD*, page 439, 1981.
- G. Stephanopoulos. *Chemical Process Control An Introduction to Theory and Practice*. Prentice Hall, 1984.
- R.E. Swaney and I.E. Grossmann. An index for operational flexibility in chemical process design. *AIChE Journal*, 31:621, 1985.
- S. Vasantharajan and L.T. Biegler. Simultaneous strategies for optimization of differential-algebraic systems with enforcement of error criteria. *Comput. chem. Engng.*, 14(10), 1990.
- V.S. Vassiliadis, R.W.H. Sargent, and C.C. Pantelides. Solution of a class of multistage dynamic optimization problems part 1 - problems without path constraints. Technical report, Imperial College, 1993.

- V.S. Vassiliadis, C.C. Pantelides, and R.W.H. Sargent. Optimization of discrete charge batch reactors. *Computers Chem. Engng.*, 18(Suppl.), 1994.
- V.S. Vassiliadis. *DAEOPT - A Differential-Algebraic Optimal Control Problem Solver, version 2.0*. Imperial College, February 1992.
- V.S. Vassiliadis. *Computational Solution of Dynamic Optimization Problems with General Differential - Algebraic Constraints*. PhD thesis, University of London, 1993.
- V.T. Voudouris and I.E. Grossmann. Optimal synthesis of multiproduct batch plants with cyclic scheduling and inventory constraints. *Ind. Eng. Chem. Res.*, 32:1962–1980, 1993.
- S.P. Walsh and J. D. Perkins. Integrated design of effluent treatment systems. In J. D. Perkins, editor, *Interactions Between Process Design and Process Control, IFAC Workshop, London, UK*. Pergamon Press, 1992.
- S. P. Walsh. *Integrated Design of Effluent Treatment Systems*. PhD thesis, University of London, 1993.
- A.W. Westerberg. A review of process synthesis. In R.G. Squires and G.V. Reklaitis, editors, *Computer applications to chemical engineering*, number 124 in ACS Symp. Ser., pages 53–87. American Chemical Society, 1980.
- A.W. Westerberg. Synthesis in engineering design. *Comput. Chem. Engng*, 13:365, 1989.
- V. White, J. Perkins, and D. Espie. Switchability analysis. In *IPDC'94*, IFAC, pages 254–259, Baltimore, June 1994.
- M.J.D. White, editor. *Advances in Water and Effluent Treatment*, London, December 1993. Mechanical Engineering Publications Limited.
- V. White. *Switchability Analysis of Process Plants*. PhD thesis, Imperial College, 1994.
- S.J. Wilkinson, A. Cortier, N. Shah, and C.C. Pantelides. Integrated production and distribution scheduling on a europe-wide basis. *Computers Chem. Engng*, 20(SB):S1275–S1280, 1996.

Appendix A

Review of operating procedure synthesis literature

This appendix contains a review of operating procedure synthesis methods. These methods were considered as possible alternatives to dynamic optimisation as tools for determining the design and optimal cyclic operation of the process structures developed in this work.

A.1 Valve sequencing operations

Pioneering work in the field of operating procedure synthesis was carried out by Rivas et al (Rivas and Rudd, 1974; Rivas *et al.*, 1974). In their work on “Synthesis of Failure-Safe Operation” the chemical processing system is modelled as a network of valves and connectors. The procedure synthesis problem is broken down into a hierarchy of goals as shown in figure A.1. At the highest and most complex level is the operation objective; this is level 0. Level I goals comprise an ordered set of verbal instructions which will cause the level 0 goal to be met. The level I goals must be specified and ordered by the user. Level II goals are now formed by expressing the level I goals as logical propositions using Boolean statements. Three propositions are used which state that a species is either flowing through, blocked at, or trapped within a given connector. The procedure synthesiser may now use the level II goals to generate a sequence of valve operations. This is achieved by proposing a single valve change and analysing changes in the system. The synthesiser looks for changes which positively relate to one or more of the goals. A safety interlock system, also developed by Rivas and Rudd, is then used to ensure that it is safe to apply the proposed change. The synthesis algorithm is demonstrated on an

industrially significant catalyst reactivation system.




Figure A.1: Goal hierarchy proposed by Rivas and Rudd

Other authors (O'Shima, 1978) have also represented the chemical plant as a network of valves and connectors in order to tackle the procedure synthesis problem. In O'Shima's terminology a connector may represent any portion of the plant blocked by valves, which could mean either a piece of equipment or a pipe. In addition to the flowing, blocked and trapped states used by Rivas and Rudd, a branch flow state is introduced to indicate that the connector is possibly an upstream branch of a flowing route. A path search is employed to determine routes for flow from one point on the network to another. Valves not located on the established route are closed. By assessing the flow states around the connectors and using information about forbidden sequences it is then possible to find a safe valve operation sequence. O'Shima points out that this sequence is not necessarily unique. An example of the use of the algorithm on a flowsheet containing two distillation columns and two alternative packed bed reactors was presented.

Although much of the work on valve sequencing operations was highly innovative, a number of drawbacks in this approach are apparent:

1. The lower level goals in the tree hierarchy (A.1), including the order in which they must be achieved, must be specified by the user. This is an intensive exercise.
2. The chemical plant is modelled in a highly simplistic manner by using valves and

connectors.

3. Process dynamics are not considered.
4. A feasible and not necessarily optimal solution results.
5. The search algorithms used are fairly inefficient.
6. Allowable operating constraints are restricted to forbidding the simultaneous presence of certain species in the same location of the plant.
7. There is no scope for incorporating design elements into the operating procedure generation framework.

Although this technique has been proven to work on industrial sized examples it is very likely that the overly simplistic modelling approach used will not be useful in the context of the integrated design problem.

A.2 State transition sequences

The next class of operating procedure synthesis techniques to be considered are those based on transition networks. The key work in this area was carried out by Ivanov et al (1981a; 1981b). A method for the algorithmisation of chemical plant startup was developed which was capable of being realised on a plant control computer. In this work the concept of the state transition graph was first introduced. The chemical process is described by a system of nonlinear algebraic and differential equations which relate the vector of controlling variables, $\underline{V} = \{u_1, u_2, \dots, u_m\}$, to the vector of process outputs, $\underline{Z} = \{z_1, z_2, \dots, z_n\}$. In addition Ivanov defines a set, $\underline{S}$,

$$\underline{S} = \{s_0, s_1, \dots, s_k\} \quad (\text{A.1})$$

whose elements are states that the system can be in during the startup operation. A control action, u_{ij} , is now defined which causes the system to move from state s_i to state s_j . From this, a *graph of state transitions* may be built up as shown in figure A.2.

Each of the arcs joining the state nodes is given a weighting value. The weighting value used depends on the problem being solved. In optimal time problems, for example, it is the time for transition between nodes i and j . It could equally correspond to material or energy usage during the transition. The optimal startup sequence is found by searching all the paths on the state transition graph leading from the initial to the goal state to find




Figure A.2: State transition graph proposed by Ivanov

the one which gives the smallest sum of weighting coefficients. Ivanov et al demonstrate the proposed algorithm on the startup of a methane conversion process.

The state transition approach to procedure synthesis overcomes some of the inherent problems with the valve sequencing techniques discussed earlier.

1. The chemical process is represented by a system of DAEs.
2. More complex constraints may be used.
3. Optimal solutions will be generated.

The approach does suffer from several weaknesses.

1. The state transition graph is generated manually.
2. A path searching algorithm is still required and is likely to be inefficient, especially as problem sizes become large.
3. There is no obvious extension to incorporate design aspects into the problem.
4. The weighting factors used must be assigned *a priori* which is both difficult and highly subjective.
5. A rigorous model of the plant is required.

A.3 Artificial intelligence methods

This section describes the first attempts to develop a formal method for operating procedure synthesis. The pioneering work in this area was largely due to Powers et al (Fusillo and Powers, 1987; Fusillo and Powers, 1988) and it is this work which will form the bulk of what follows.

Fusillo and Powers (1987) describe four main areas of their formal procedure, namely:

1. system decomposition,
2. system modelling,
3. handling system constraints, and
4. the planning method used.

It is useful to cover each of these topics in turn in this review.

A.3.1 System decomposition

The system is subdivided into processing subsystems in order to simplify the procedure generation problem. Fusillo and Powers define a subsystem as consisting of one or more unit operations which share a common function, which can be isolated from its neighbours. In order to effectively decompose the chemical system Fusillo and Powers introduce the important new concept of a *stationary state*, which a subsystem should have where possible. A stationary state can be characterised by certain criteria:

1. the presence of simultaneous inverse operations, or
2. a subsystem with large thermal energy, pressure or material capacitance.

These terms require some clarification. *Simultaneous inverse operations* occur when a transformation is performed on one part of the process and the inverse transformation is carried out in a connected part of the process. A good example of this is a distillation column where evaporation occurs at the bottom of the column and condensation occurs at the top. The column may be run as a subsystem under total reflux until it can be integrated with the rest of the process. In the case of process capacitance, Fusillo and Powers cite the example of a charged liquid phase reactor whose contents are heated until other reactants are ready to be fed. The idea of stationary states is an important one since they provide convenient stopping-off points during process operations.

A.3.2 Modelling

Fusillo and Powers make use of functional modelling techniques. In this approach a unit operation is modelled as either a sink or source of energy, material or momentum. For each unit a list of effects is given which describes the effect an operation on the unit

will have on the subsystem. A vector of states is used to globally represent the subsystem variables which are, in turn, described by multiple, discrete values.

A.3.3 Constraint representation

Constraints can be both local and global and are established *a priori* based on five factors given by Fusillo and Powers as:

1. preconditions for unit operations e.g. to avoid damage to a unit,
2. requirements for a reaction,
3. production requirements,
4. hazard avoidance,
5. materials of construction.

As an example, a global constraint may be, “*do not mix Cl_2 and CH_4* ”, in order to avoid explosion.

A.3.4 Operating procedure planning

The planning operation is posed as a mean-ends analysis. The initial and goal states of the system are described by vectors of temperature, pressure, composition etc. and the operating procedure aims to find operators which reduce the difference between these states. In the case of process systems the operators are manipulations to process units. By using stationary states as points which the process passes through during an operational change the search space is reduced in size. In addition, the operator search is hierarchical with “macroscopic” operations such as *start compressor* being broken down into tasks such as, *start lube oil and lube flush* etc. The use of hierarchical planning allows tasks to be ordered such that time consuming tasks and those which do not interfere are performed simultaneously.

Implementation of the proposed method is achieved through the “Prototype Operating Procedure Synthesis Program” (POPS). The user is required to give information about the flowsheet structure, constraints and the initial and goal states of the system. A hierarchical model library is used for the unit operations. Fusillo and Powers demonstrate the POPS package on the startup of a chlorination process. The work carried out in developing POPS has been extended (Fusillo and Powers, 1988) to use causal models (in the form of digraphs) and algebraic equations to model the system in addition to the

functional models previously used. Startup of a heat exchanger network is considered as a case study.

The work presented by Fusillo and Powers represents a significant contribution to the field of operating procedure synthesis, but it does have some weaknesses:

1. The dynamics of the system are not considered, instead the system is represented in a quasi steady-state manner.
2. Global, discrete variables are used to describe the subsystem state.
3. The optimality of the solution is not considered.

