



City Research Online

City St George's, University of London

Citation: Knight, R. A. (2008). Optimisation methods for staff scheduling and rostering: an employee-friendly approach. (Unpublished Doctoral thesis, City University London)

This is the submitted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://city-ref-test.eprints-hosting.org/id/eprint/12268/>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).

OPTIMISATION METHODS
FOR STAFF SCHEDULING AND
ROSTERING:
AN EMPLOYEE-FRIENDLY
APPROACH

Roger Alan Knight

A DISSERTATION IN PARTIAL
FULFILLMENT OF THE REQUIREMENTS FOR
THE DEGREE OF DOCTOR OF PHILOSOPHY

SUBMITTED TO CASS BUSINESS SCHOOL, CITY UNIVERSITY,
FACULTY OF ACTUARIAL SCIENCE AND INSURANCE

November 2008

Contents

1	Introduction	1
1.1	The Call Centre industry	2
1.2	The “Pooling Principle”	3
1.2.1	Illustrative Examples	3
1.3	The Call Centre rostering problem	5
1.4	Aims and Objectives	8
1.5	Contributions	10
1.6	Thesis Outline	11
2	Call centre management strategies, methods and models	14
2.1	The cost reduction strategy	14
2.2	The revenue generation strategy	17
2.3	Operations management	19
2.3.1	Employee work rates	20
2.3.2	Work Schedules	22
2.4	Motivation for Study	23
2.5	Chapter Summary	24
3	Employee Friendly Rostering Goals	26
3.1	Legislative Requirements	27

CONTENTS

3.2	Guidelines and Principles	28
3.2.1	The Ergonomic Background	30
3.2.2	Ergonomic Rostering Goals	32
3.3	Generalised Preferences	33
3.4	General Considerations	34
3.5	Analysis of Goals	34
3.6	Chapter Summary	37
4	Literature Review	38
4.1	Decomposition	38
4.2	Forecasting and Staffing	40
4.3	Tour Scheduling	40
4.3.1	Key Model	43
4.4	Tour Assignment	45
4.5	Integrated Models	46
4.6	Chapter Summary	47
5	Tour Scheduling	48
5.1	Tour Scheduling Goals	50
5.1.1	Patterns of Off-Duty	51
5.1.2	Shift Sequences	52
5.1.3	Break Times	54
5.1.4	Individual Preferences	55
5.2	Tour Scheduling Formulation KG1	56
5.2.1	Under staffing	56
5.2.2	Notation	57
5.2.3	Integer Programming Formulation, KG1	60
5.2.4	Goal Seeking: model parameterisation	63
5.3	Quality of the Solution	67
5.3.1	Comparisons with Brusco & Jacobs model: A Telecomms call centre	67

5.3.2	Comparisons with Commercial Package:	
	A Telecomms call centre	69
5.3.3	Comparison with Brusco & Jacobs:	
	A Financial call centre	72
5.3.4	Fitness of the model: practical requirements	73
5.4	Chapter Summary	75
6	Tour Assignment	77
6.1	Pre-Assignment process	78
6.2	The Tour Assignment Model	80
6.2.1	Tour Assignment goals	81
6.2.2	Hard constraints	83
6.2.3	Soft constraints	84
	6.2.3.1 Week-week continuity	85
	6.2.3.2 Fairness	86
6.2.4	The Cost Function	88
6.3	Post-Assignment processing	90
6.4	Empirical Example	92
	6.4.1 Initial weight settings	93
	6.4.2 Alternative weight settings	95
6.5	Chapter Summary	96
7	Application to Nurse Rostering	98
7.1	Practical approaches to Nurse Rostering	100
	7.1.1 Heuristic Approaches	101
	7.1.2 Mathematical Programming Approaches	108
7.2	The Nurse Rostering Model	117
	7.2.1 Stage 1 - Attendance Rostering	120
	7.2.2 Stage 2 - The Shift Allocation Model	127
7.3	Chapter Summary	133

8	The Nurse Rostering Problem	135
8.1	Introduction	136
8.2	Benchmark Problems	140
8.2.1	Problem Descriptions and Rostering Constraints	140
8.2.2	General Approach	143
8.3	Analysis of Problem Structure of “GPost”	145
8.4	Optimal Solutions	148
8.4.1	Optimal Solution for “GPost”	149
8.4.2	Optimal Solution for “ORTEC01”	150
8.5	Month to Month Continuity	152
8.5.1	Advantages	153
8.5.2	Approach	154
8.5.3	Empirical Example	155
8.6	Chapter Summary	159
9	Conclusions and Future Research Direction	161
9.1	Thesis Summary and Conclusions	161
9.2	Future Directions	167
A	The Brusco-Jacobs Tour Scheduling model	170
A.1	Shift start times	171
A.2	Break Scheduling	171
A.3	The Mathematical Model	173
A.3.1	Formulation for P2, for specific shift types and breaks . . .	175
B	Tour Reduction for Tour Scheduling examples	178
B.1	Example 1 - “Telecomms” call centre	179
B.2	Example 2 - “Bank” call centre	182
	References	194

List of Tables

1.1	Numbers of call versus Agent Occupancy	4
3.1	Legislative requirements for employee rostering	28
3.2	Ergonomic principles and practices for employee rostering	32
3.3	Generalised shift and tour preferences in employee rostering	33
5.1	Goals for Tour Scheduling model	50
5.2	Solutions for alternative cost parameter weightings, excluding shifts goal	64
5.3	Solutions for alternative cost parameter weightings, including goal of minimising the number of shifts	66
5.4	Comparison of results from KG1 and P2 for “Telecomms” call centre	68
5.5	Comparison of KG1 and Q-Max for “Telecomms” centre	71
5.6	Comparison of results from KG1 and P2 for “Bank” call centre	72
6.1	Tour combinations for night shift re-distribution	80
6.2	Goals for Tour Assignment model	82
6.3	Hard constraints for Tour Assignment model	84
6.4	Soft constraints for Tour Assignment model	87
6.5	Penalty weights for Tour Assignment example	93
6.6	Violations of high penalty constraints	94
6.7	Distribution of unsociable shifts	95
8.1	Soft constraints for Nurse Rostering benchmarks	143

LIST OF TABLES

8.2 Penalty exemptions for Nurse Rostering benchmarks 144

8.3 Hard constraints for Nurse Rostering benchmarks 144

8.4 Penalties for problem “GPost” 149

8.5 Penalties for problem “ORTEC01” 151

8.6 Historical solutions to ORTEC01 151

8.7 Penalties for “ORTEC01” with continuity 159

A.1 Shift types per band, b=5 171

List of Figures

5.1	Break windows for an 8 hour shift	54
8.1	Optimal solution for problem “GPost”	149
8.2	Optimal solution for problem “ORTEC01”	150
8.3	December-January continuity, “ORTEC01”	157
8.4	Optimal solution for “ORTEC01” with month-to-month continuity	159
B.1	Incomplete tours for “Telecomms” (1)	180
B.2	Incomplete tours for “Telecomms” (2)	180
B.3	Reduced tours for “Telecomms”	181
B.4	Repairing a backward rotation	182
B.5	Incomplete tours for “Bank”	183
B.6	Reduced tours for “Bank”	183

Acknowledgements

I would like to express my gratitude to my supervisor, Professor Celia Glass, for her staunch support during the last three years. Without her guidance, suggestions and ideas, this thesis would not have been possible. I acknowledge the support of the other members of the faculty, academics in other departments of the University, and my fellow PhD research colleagues. I have enjoyed working with you all.

I offer my sincere thanks to my wife Janice, daughter Elizabeth and son Edmund for their patience, understanding, and encouragement during the course of my research. I also thank my wider circle of friends and relations who have provided welcome support.

I dedicate this work to my parents Jean and Alan who taught me the importance of education and encouraged me in my academic studies.

Declaration

I grant powers of discretion to the University Librarian to allow the thesis to be copied in whole or in part without further reference to me. This permission covers only single copies made for study purposes, subject to normal conditions of acknowledgment.

Abstract

The growth in the global call centre industry over the last twenty years has been huge. The main motivating factor for businesses to introduce call centres as their main vehicle for handling customer contacts has been that call centres are inherently efficient. Since the mid-1980's, UK businesses have sought to establish competitive advantage by using call centres to reduce the cost of managing their customer contacts. Over the last decade or so, however, an alternative strategy has emerged based not on cost-reduction and efficiency, but on revenue generation and service quality. This new strategy places high value on customer and staff retention.

This thesis is concerned with the operations management task of employee rostering. We argue that traditional models for producing rosters for call centre employees are designed to support the older efficiency-based culture, and are inappropriate for call centres adopting the more recent quality-based culture. We show how the use of methods and models drawn from conflicting management philosophies contributes to the high level of employee turnover, and inhibits the drive for service quality. Our primary contributions are to identify a set of rostering goals which reflect the interests of the employees, and to quantitatively represent these goals in a system of mathematical rostering models designed to support the revenue generation strategy. Our models are implemented using the robust Mixed Integer Programming methodology. In addition, we adapt our model to address the related problem of nurse rostering, and solve two benchmark problems to optimality. We demonstrate that our model generates rosters of a higher quality than the alternatives, at no additional cost.

Chapter 1

Introduction

This thesis is concerned with the methods and models used to calculate attendance rosters for the telephone answering agents in call centres, and for hospital nurses. We provide a review of the academic research on call centre rostering models and identify a research gap. The nature of our contribution is to develop alternative models which systematically combine to offer a novel approach to the rostering of telephone service agents in call centres, in support of an emerging call centre management culture. We provide an analysis of the models used for rostering hospital nurses, and further develop our own model for application in this field. In this introductory chapter we begin by providing a background to the development of the call centre industry, particularly in the UK, and introduce the call centre employee rostering problem. We set the context for the theoretical part of the thesis, and identify the objectives which we will pursue.

1.1 The Call Centre industry

The earliest form of call centres were the operator-controlled telephone exchanges installed by national telephone operating companies such as AT&T in the USA, and the UK Post Office. These offices provided a service whereby a calling customer could be connected to a second party upon request. The provision, via a team of human agents, of a telephone-based, on-demand service characterises the call centre. A definition of a modern call centre is given by Taylor & Bain (1999):

“A dedicated operation in which computer-utilising employees receive inbound - or make outbound - telephone calls, with those calls processed and controlled either by an Automatic Call Distribution (ACD) or predictive dialing system.”

Over the last 20 years, the growth in the UK call centre industry has been phenomenal and is continuing, despite the fact that call centre services can now be conducted via the internet (for example, railway timetable enquiries), are being “offshored” (that is, dealt with by overseas call centres), or fully automated (such as the payment of household bills). A report on the UK industry by ContactBabel (2006b) estimates that 1 million people are directly employed in UK call centres at the end of 2007. The figure for the US industry is over 5 million (ContactBabel (2006a)). The value of the UK call centre industry is estimated at 13.3 billion.

The rapid growth of the global industry has been enabled by the deregulation of national telecommunications industries together with technological advances in low-cost networking and digital switching equipment. However, the reason that businesses have been motivated to adopt call centres as their primary vehicle for handling customer contacts is they are inherently efficient. We will now explain why this is the case.

1.2 The “Pooling Principle”

Call centres are inherently efficient due to the economies which arise when a number of small groups of servers are pooled together into a single larger group. For a general outline of the mathematics of call centres, see Koole (2005). For a detailed description of this “pooling principle”, see for example Cattani & Schmidt (2005). Here, we offer a brief, intuitive explanation as follows.

Consider an insurance company operating a number of small, high-street branches. To speak to a member of staff, customers must call their local branch. With only a small number of staff present, there is a fair chance that no-one will be immediately available to answer the call, and the customer must wait for answer. But what if the customer could be connected to a member of staff in *any* of the branches, rather than just the local one? In this case, the customer only has to wait if *all* the staff in *all* the branches are unavailable, which is less likely than just the local branch being busy. So, “pooling” the staff resources means that on average, callers will be answered more quickly when all the staff are formed into a single large team serving the whole customer base, than if the same staff are divided into several independent units each dedicated to a small group of customers. Conversely, the same level of service (that is, the same average customer delay) can be provided by a centralised team *using fewer staff resources*.

1.2.1 Illustrative Examples

To illustrate the relationship between workload, staffing and service levels, we now give some examples obtained from the “Erlang-C” staffing model. The model

1.2 The “Pooling Principle”

assumes that the call arrivals form a poisson distribution, with exponentially distributed service times, and that delayed calls queue until answered. Let us assume that these conditions hold for a short period of time, say $\frac{1}{4}$ hr. Further assume that calls take an average of 5 minutes to handle, and that our service target is to answer a minimum of 80 percent of calls within 20 seconds of their arrival. Then we can calculate, using Erlang-C, the number of agents required for different call volumes as shown in Table 1.1, together with the predicted service level (i.e. the percentage of calls to be answered within the 20 second threshold), and the agent *occupancy* level, which is the percentage of time that agents will be working on calls (the remainder being spent waiting for calls to come in).

No. of calls per $\frac{1}{4}$ hr	No of agents Required	Service Level (%)	Average Occupancy (%)
2	2	84.8	33.3
6	4	84.8	50.0
18	9	84.0	66.6
30	14	86.7	71.4
60	25	85.0	80.0
150	57	84.5	87.7
300	108	85.0	92.6
1000	345	84.5	96.7

Table 1.1: Numbers of call versus Agent Occupancy

We can see that in order to maintain a consistent level of service, the productivity of the operation (as expressed by the occupancy figure) dramatically increases as the workload increases. The economy of scale is a result of the reduced effects of the random elements in the call arrival process and duration, as the number of transactions increases.

These scale economies give rise to the possibility of reducing staff costs by

1.3 The Call Centre rostering problem

creating a centralised point for customer contact - a call centre. In chapter 2 we describe how some businesses have utilised these opportunities in the pursuit of low-cost competitive advantage, and how others are pursuing an alternative strategy based on revenue generation. We will then go on to assess the applicability of traditional call centre rostering methods and models within the context of these two alternative strategies. We now describe the call centre employee rostering problem.

1.3 The Call Centre rostering problem

The employee rostering problem generally is to decide the hours of attendance at work, for each of a group of employees. Ernst *et al.* (2004a) define rostering as

“[the process of] allocating suitably qualified staff to meet a time dependent demand for different services while observing industrial workplace agreements and attempting to satisfy individual work preferences”.

In a call centre, customer demand will typically vary very markedly both within a day and from day-to-day. Some call centres operate 24 hours per day, 7 days per week. The vast majority operate outside of the business hours of 9:00-17:00, Monday to Friday. The problem for call centres is therefore to provide the best balance of staff resources across the days of the week and, and within each day, to stagger the shift start and finish times of each employee so as to match the increase and decrease of workload at the beginning and end of the day. Lunch and coffee breaks are used to align the level of staff provision with the staffing requirement throughout the day. Additionally, some employees may be scheduled

1.3 The Call Centre rostering problem

to attend training sessions or team meetings, or to perform non-demand activities. The shift start and end time stagger, the scheduling of short duration activities and the variability in the presented workload require a high degree of resolution in call centre rosters, which are typically calculated on a quarter-hourly basis.

In order to handle this level of complexity, the process of devising employee attendance details is traditionally decomposed into four main sub-problems, as first described by Buffa *et al.* (1976).

1. The *Forecasting* problem is to produce a quarter-hourly workload forecast for a future week. In call centres, the amount of work presented at any time is not generally under management control, rather customers call at a time of their own choosing. Employee rosters can therefore be regarded as being independent of the forecasting process and so we are not concerned with the forecasting problem in this thesis.
2. The *Staffing* problem is to take the quarter-hour workload forecasts and to calculate the number of staff that will be required during each quarter-hour interval, in order to meet some target of service quality. As was mentioned earlier, call centres have existed in some form since the early 20th century, and queuing-theory based staffing models date back to the same period. The “Erlang-C” model, first published in Erlang (1948), remains a cornerstone of many practical software applications. Modern call centre staffing models are still primarily based on queuing theory, with a number of alternative models, supporting different management strategies, formalised by Borst *et al.* (2004).
3. The *Tour Scheduling* problem is to cover the interval staffing requirements

1.3 The Call Centre rostering problem

with a set of weekly “tours”, each of which is a set of work attendance details including days to be worked, shift start and finish times, break times, etc. for one employee. Weekly tour scheduling models have evolved from the set covering formulation for the daily shift scheduling problem presented by Dantzig (1954). Modern approaches such as Brusco & Jacobs (2000) combine days-off scheduling, daily shift scheduling, and break scheduling in a single Integer Programming model.

4. The *Tour Assignment* problem is to determine which of the employees should work which of the scheduled tours. This problem has received little attention in the literature, perhaps because the two-dimensional assignment problem (of which our Tour Assignment problem is an example) has been well understood since Kuhn (1955). Some more recent papers address the prioritisation of assignment, based on seniority (Thompson (1997a)) or productivity (Goodale & Thompson (2004)).

There are alternative interpretations as to which of these components constitute the “employee rostering” problem, with Thompson (1995b), Grossman *et al.* (2001), and Ernst *et al.* (2004a) offering differing viewpoints. In this thesis we will use the term “employee rostering” to refer to the *tour scheduling* and *tour assignment* problems. We recognise, however, that the choice of model used to solve the *staffing* problem will impact upon the roster solution, and we will discuss this issue in chapter 2. Our interpretation of the call centre employee rostering problem in this thesis is in line with that understood by researchers in other employee rostering domains, such as nurse rostering. Burke *et al.* (2004b), for example, define nurse rostering as

“the short-term timetabling of staff, with a typical horizon of a few weeks...we will use the term nurse rostering to refer to this process (the allocation of nurses to periods of work over several weeks)”.

Our view of the call centre rostering problem is therefore to take the interval staffing requirements as input data, and to provide the attendance details for each employee as output. The calculation of an “optimal” solution (that is, the one which matches our rostering objectives most closely), within the framework of a set of “constraints” (which determine the boundaries of acceptability of the solution), is characteristic of the set of problems that lie within the field of discrete mathematics known as *combinatorial optimisation*.

1.4 Aims and Objectives

The goal of this thesis is to explore the applicability of traditional call centre employee rostering models and methodologies in the context of alternative call centre management strategies, in particular those which emphasise revenue generation over low-cost efficiency, and which hold employee retention to be an important goal. Our specific aims are,

1. *To explore the broad literature on call centre management in order to align the call centre operational management task of “work force management” with modern call centre management strategies and policies.* The literature on call centre management generally will enable to us to assess the role of rostering methods and models in the context of the overall goals being pursued by call centre executives. Our objective is to inform future dialogue between HR scholars and those concerned with call centre operations.

2. *To investigate the “human aspects” of employee rostering, in order to specify a set of rostering goals consistent with the culture of modern call centres in Europe and North America.* We will consider the legal and contractual requirements relating to call centre employee rostering, known “best practice” guidelines, and ergonomic rostering principles. The objective is to define a set of goals which can be quantitatively represented in a system of call centre rostering models.
3. *To review in detail the literature on call centre rostering models.* The objective is to obtain a clear understanding of the developments in call centre rostering models to date. We will assess the extent to which current models are able to take account of our rostering goals, in order to identify current shortcomings and to frame our own contribution.
4. *To develop an employee-friendly model for the tour scheduling problem.* Current approaches to tour scheduling aim simply to minimise the number of tours and, hence, employees. They take little account of the preferences of the employees except in the expression of hard rules, which form the framework within which the roster can vary. This is rather black-and-white, and our objective is to explore the possibilities for accommodating the measures by which employees judge the quality of a set of tour schedules.
5. *To apply our rostering goals to the tour assignment problem.* Some commercial workforce management systems take account of individual employee preferences when assigning weekly tours. However, the methods are not formalised and the range of options is limited. Our objective is to improve upon the consistency with which goals are applied in addressing both tour

scheduling and tour assignment.

6. *To integrate the tour scheduling and tour assignment models.* One problem with the decomposition of call centre rostering into these two separate stages is that when employee preferences and requests are taken into account at the tour assignment stage, the choice of assignment is restricted to those tours which have already been scheduled. Our objective is to identify methods which improve the likelihood that preferences and requests can be accommodated.
7. *To draw comparisons between the call centre rostering problem and the nurse rostering problem.* These two problems have a number of similarities and it will be interesting to ascertain the extent to which techniques, models, and general insights relating to either one are applicable to the other. Our objective is to assess the extent to which call centre rostering methods can be applied, or developed, to address instances of nurse rostering problems.

1.5 Contributions

1. The identification of call centre operational management systems and models which support the alternative call centre management strategies. We highlight the incompatibility between the traditional call centre rostering approaches widely used in practice, and the revenue-generation strategy. This provides an inter-disciplinary link between the research areas of call centre HR and operations management.

2. A formalised list of the goals of an employee-friendly, call centre rostering system. The list is categorised by legal requirements, best practice guidelines and ergonomic rostering principles, and generalised preferences. This set of goals more closely resemble those found in nurse rostering than in current call centre rostering models.
3. An Integer Programming model for the Tour Scheduling problem, which takes account of our employee-friendly rostering goals. Results indicate that this model produces higher quality rosters than other models in the literature, using no more tours or employees. This contribution is presented in the form of a working paper, Glass & Knight (2008a).
4. A Tour Assignment model which integrates with the Tour Scheduling model, and which operates to meet the same, consistent set of rostering goals.
5. A Nurse Rostering model which combines aspects of our Tour Scheduling and Assignment models, for application to this alternative domain. Results obtained on two benchmark problems are superior to those obtained previously by other models. Elements of this contribution form our paper, Glass & Knight (2008b), submitted to the European Journal of Operational Research in September 2008.

1.6 Thesis Outline

This section gives an overview of the content of the thesis. In chapter 2 we draw upon the research literature in call centre management to identify the alternative strategies being pursued by call centre executives. We assess the extent to which

the methods and models used in the operational management tasks related to employee rostering are supportive of the business objectives.

In chapter 3 we analyse known best practice in rostering, together with “ergonomic” rostering principles, in order to construct a set of rostering goals for incorporation into a system of rostering models which take account of the employees perspective.

In chapter 4 we present a detailed review of the literature on call centre rostering models. We highlight those model characteristics which align with the cost reduction strategy, and identify areas where there are conflicts between the goals of the current rostering models, and those goals which we have identified as being “employee-friendly”.

Our tour scheduling model forms the basis of a submitted paper, Glass & Knight (2008a), which forms the body of chapter 5. The key conclusion is that our tour scheduling model produces higher quality rosters than alternative models, at no extra cost. This conclusion is supported by results obtained from comparative exercises using data from a real world call centre. The model is based on an Integer Programming formulation which calculates a set of weekly tours, each of which has a single shift start time on each working day, but may contain fewer working days than are strictly required by the contractual rules. The idea is to subsequently re-combine these partial tours in order to reduce the total number.

In chapter 6 we present our tour assignment model. The cost of assigning each of the scheduled tours to each of the available employees is calculated so as to reflect our overall rostering goals. The tours are then assigned using a standard formulation of the two-dimensional assignment model, i.e. the so-called Hungarian method Kuhn (1955). Once the tours have been assigned, we seek to

improve the quality of the roster by exchanging blocks of shifts (referred to as “stints”) between employees. In this way, the original tours are to some extent reconstructed, and so the tour scheduling and assignment models can be regarded as being more tightly integrated than in previous approaches.

In chapters 7 and 8 we turn our attention to the nurse rostering problem. In the first of these two chapters, we point out the similarities between the nurse and call centre employee rostering problems. We present a critical evaluation of a series of papers concerned with practical nurse rostering systems, and give the formulation of a two-stage Integer Programming model for the nurse rostering problem. Chapter 8 is formed by the second of our submitted papers, Glass & Knight (2008b), in which we present an analysis of the structure of two benchmark instances of the nurse rostering problem. We give the optimal solutions to these instances, both of which were previously unknown and were obtained with less than 10 minutes execution time on a standard desktop PC (2.67 Ghz, P4, 512Mb) using CPLEX v10.

In chapter 9 we summarise the content, conclusions, and contributions of the thesis, identify those aspects which are in need of further development, and indicate directions for future research.

Chapter 2

Call centre management strategies, methods and models

In this chapter we draw upon the research literature relating to call centre management, in order to identify the strategic aims being pursued by call centre executives. We discuss the emergence of a call centre culture which places high value on both customer and employee retention, and demonstrate inconsistencies between these values and the traditional operations management models for the call centre employee rostering problem. We briefly discuss the factors which determine customer satisfaction in call centres, and highlight the need for future research in this area.

2.1 The cost reduction strategy

As was described in the introductory chapter, call centres are inherently efficient due to scale economies which accrue as call centres increase in size (known as the “pooling principle”). The strategic use of call centers to acquire low-cost competitive advantage by reducing operating costs through the realisation of these efficiencies began in the UK in the early 1980’s and was identified by Taylor &

2.1 The cost reduction strategy

Bain (1999). The finance companies First Direct and Direct Line were two of the first to follow this strategy. In order to stay competitive, others in the financial sector followed this lead, and the idea was adopted by companies in other sectors of industry in order to gain advantage over their own business rivals. This process led to the rapid expansion of the call centre industry. In order to sustain, or challenge competitive advantage, call centres were pressured to reduce costs still further.

In spite of their high-tech image, call centres are essentially manpower intensive operations, and the main component of the operating cost of a UK call centre is the cost of employment, estimated at 72% for UK call centres by the Department of Trade and Industry, DTI (2004). An obvious way to reduce costs is therefore to reduce the size of the work force. Miozzo & Ramirez (2003) calculated that a one second reduction in the average handling time of Directory Enquiry calls would save British Telecom £2 million annually in staff costs. In call centres in general, and in particular those in highly competitive industries, the commercial pressures have led call centres to adopt a mass-production model, characterised by jobs involving a high degree of standardisation, constant pressure to meet throughput-based quality targets, and invasive levels of monitoring, as described by Batt & Moynihan (2002).

The aim of standardisation is to minimise the size of the workforce by driving down the amount of resource required to complete each transaction. Standardisation in call centres is pursued through the implementation of detailed transaction processes, which are rigidly enforced by scripting systems that guide the agent through each customer transaction, ensuring that the correct activities are carried out in the right sequence, for every call. Measures of quality relate to the

2.1 The cost reduction strategy

length of time that the customers wait for answer; the average time spent queuing (ASA, or Average Speed of Answer), the percentage of calls answered within some time threshold (SL, or Service Level), and the percentage of callers abandoning. Further pressure is brought to bear on workforce costs by the use of management metrics. According to leading call centre texts (for example Cleveland & Mayben (1997), Anton (1997)), amongst the typical call centre Key Performance Indicators (KPI's) are: average talk time, average post-talk time, calls per 8hr shift and Agent Adherence (that is, the percentage of working time when agents are logged into the ACD system). These metrics are monitored on an individual as well as group basis, and are used to reinforce the drive for standardisation.

Perhaps unsurprisingly, the combination of repetitive tasks and a stressful working environment has resulted in high levels of employee turnover in those call centres which operate under the mass-production model. A recent global report by Holman *et al.* (2007) found the average annual staff turnover rates to be as high as 36% for call centers with low-quality, "low discretion - high monitoring" jobs. On the other hand, the high level of standardisation based on rigid transaction processes implies that a low level of skill is required to perform the job, and hence the cost of recruiting and training new employees should in theory be low. Indeed, Wallace & Eagleson (2000) found that some call centre managers regard a high level of turnover as desirable, since this allows the replacement of burnt-out employees with fresh, enthusiastic ones. Bristow *et al.* (2000) argue that businesses following the low-cost strategy via the mass-production model locate their call centres in regions of the UK that have a large pool of low-cost labour, in order to ensure a constant supply of fresh recruits. The mass-production model therefore ought to be, for some call centres at least, "turnover-proof".

2.2 The revenue generation strategy

The mass-production model lends itself most readily to call centres with high volume, low-complexity call types. Information services, catalogue sales, airline ticket sales, insurance quotations and bank account enquiries are examples. This generation of UK call centres dates back to the early 1980's, but they are reducing in size and number as new technology enables increasing automation. Miozzo & Ramirez (2003) point to the role of automatic salutation equipment and internet access in reducing the demand on UK directory enquiry call centres, and the use of speech recognition technology to remove the need for US directory call centres altogether.

2.2 The revenue generation strategy

More recently, an alternative call centre management strategy has emerged, as business executives have come to realise the potential for revenue generation offered by call centres. Call centre services have become more customer focused, cross-selling of products to existing customers is increasing, and customer retention is now a key issue. Mehrotra (1997) points to the increasing adoption of this strategy, and argues that since call centres provide direct personal contact between a business and its customers, the quality of the call center experience will affect the customers' perception of the business as a whole. The success of this strategy, Mehrotra argues, is founded on high standards of customer service and is achieved via the mass-customisation production model, as described by Pine (1999). In this model, a range of specialised services are offered, and customer transactions are to some extent tailored to the needs of the individual caller, in contrast with the standardised approach of the mass-production model. We note

2.2 The revenue generation strategy

that some non-commercial call centres such as those in the public sector adopt a similar set of policies, not in order to generate business, but to provide a flagship for their organisation. In this thesis, we will use the term “revenue generation” to include the strategies of such organisations.

The key element in the delivery of service quality through mass-customisation is the employment of a skilled, knowledgeable, and well motivated workforce, who are empowered to use their discretion and judgment to conduct each call transaction toward a productive outcome. An important aspect of the introduction of mass-customisation in call centres is the replacement of rigid scripting systems by Customer Relationship Management (CRM) systems which present the telephone agent with customer-specific details such as their sales or service history, enabling the agent to decide for themselves whether to engage the customer, for example, in cross-selling. This production model is complemented and supported by Human Resource (HR) management policy, implemented through High Commitment Management (HCM) techniques. These include selective recruitment, incentive schemes relating to call handling quality as well as quantity, on-going training programs, team working, and the identification of career paths. The application of these techniques in call centres has been identified by, for example, Houlihan (2002), Deery & Kinnie (2002), Batt & Moynihan (2002).

Unlike call centres operating the mass-production model, these call centres are hit hard by staff turnover, as the cost of recruiting and training new employees reflects the high levels of skill and knowledge required to do the job effectively. Although Hutchinson *et al.* (2000), Batt (2002) and Guthrie (2001) all found evidence linking HCM practices with reduced levels of employee turnover, the continuing high cost remains a major issue in the call centre industry. Taking

account of the fact that newly recruited employees are less effective than their more experienced colleagues, the total cost of turnover and absenteeism to the UK call centre industry has been estimated at over £1 billion annually (CMInsight (2004)).

2.3 Operations management

It is our contention in this thesis that the paradigm change in call centre strategy is hampered by the application of operations management methods rooted in the earlier low-cost efficiency based strategy, to call centres pursuing the revenue generation, high-quality strategy. Batt & Moynihan (2002) support this position and offer the critique that:

“Operations management has focused on developing algorithms for efficient staffing patterns. The entire calculus focuses on increasing volume and minimising labour costs.”

These goals are, of course, entirely in line with the low cost, efficiency driven strategy described earlier. Accurate forecasting and rostering, allied to the economies of scale achievable in large call centres mean that employees in some call centres can be expected to spend over 90% of their available time dealing directly with customers, while still answering the large majority of callers either immediately or after a very short wait. Gans *et al.* (2003) describe such levels of occupancy as being achievable, on average, in large *best practice* call centres. However, we now present evidence suggesting that very high work rates and unsympathetic work schedules are not good for the employees.

2.3.1 Employee work rates

As was outlined in the introductory chapter, the staffing problem is to calculate the number of employees required during each quarter-hourly staffing interval. Since the number of customer calls arriving in each interval is “given” (the customers call at a time of their own choosing), the average work rate is determined by the number of employees available, which in turn is based on the requirement calculated by the staffing model. So, while the staffing problem itself is outside the scope of this thesis, the choice of staffing model will have a marked effect on job quality, which is our central interest. To illustrate the effects on the employees of very high work rates, we quote Deery *et al.* (2002):

“The speed and pace of work was a particularly significant factor in the depletion of emotional resources. Should organizations fail to address the determinants of emotional exhaustion it is likely that employees may adopt a strategy of withdrawal as a mechanism of coping with the work environment. This could involve work absences and a depersonalized approach to customers. Such an outcome is inconsistent with both rising customer expectations about service quality and the standards now being specified by organizations for service performance.”

Gans *et al.* (2003) described three alternative *staffing regimes*, formalised by Borst *et al.* (2004), to be applied according to the relative importance attached to service quality and staffing costs. We present a brief outline of these regimes, in order to illustrate the importance of adopting the staffing regime appropriate to the overall strategy.

In the *Quality and Efficiency Driven (QED)* staffing regime, the idea is to define some service level target (reflecting the amount of time customers spend waiting for answer), which is acceptable to the customer base. The minimum number

2.3 Operations management

of employees required to deliver that level of service, during each quarter-hour staffing interval is then calculated using some staffing model. Since the planned service level remains constant, the average employee workload increases as the total workload increases, in line with the pooling principle. This is the regime traditionally operated in practice, implemented using staffing models based on queuing theory such as the Erlang-C model (Erlang (1948)) or simulation, and can result in very high work rates in large call centres. At the extreme of this is the *Efficiency Driven (ED)* regime, which aims to maintain work rates at close to 100%, with the majority of callers having to wait for answer. This is neither “employee-friendly” nor customer focused.

The alternative is the *The Quality Driven (QD)* staffing regime. Here, the aim is to ensure that the large majority of customers are answered immediately and do not need to queue, with the assumption that the cost of causing customers to queue is high. According to the authors, this goal is best achieved by building a fixed proportion of slack (or *waiting time*) into the staffing calculation. The proportion of waiting time should be high enough to ensure the delivery of a good service level at the least busy times of day. Since the average employee work rate remains constant at this level, the service level will improve as the workload increases. The QD staffing regime is designed to give high quality of service and avoids the need to drive work rates too high, and therefore fits in with the revenue-generation strategy. Moreover, the additional slack afforded by the QD regime allows greater scope for employees to use discretion in the handling of individual calls. This is a key element of HCM, and thus the QD staffing regime is systematically supportive of HCM HR policies and the revenue-generation strategy. There are, however, two issues that are yet to be fully

addressed.

- One issue is that the authors are primarily concerned with “large” call centres and the behaviour of queuing models under heavy workload conditions. However, the majority of UK call centres are small, with fewer than 50 employees (DTI (2004)). Moreover, some of the largest call centres operate a 24-hr, 7-day service and at certain times, such as during the night and at weekends, may experience low levels of demand.
- A second issue is that while it is generally assumed that measures of the amount of time that call centre customers spend waiting for answer are strong indicators of customer satisfaction, there is little empirical evidence to support this assumption. A study by Feinberg *et al.* (2000) of the factors contributing to call centre customer satisfaction found that of the queuing-related metrics, only call abandonments had a significant (though weak) influence on customer satisfaction, with the main determinant being “first call resolution”.

We therefore perceive that further research is required in order to develop staffing models which support small call centres operating the revenue generation strategy, and which reflect more directly the needs of the customers and the employees.

2.3.2 Work Schedules

Work schedules are the work attendance details produced by the call centre employee rostering models, and therefore the central concern of this thesis. To illustrate the effects on the employees of unsympathetic work schedules, we quote Totterdell (2005):

“Unfortunately, some work schedules can seriously compromise the health and productivity of employees. Problems are most likely to occur when work schedules are unsympathetic to the body clock, do not allow sufficient time for physiological and psychological recovery, and do not take account of employees’ preferences.”

The quote from Batt and Moynihan at the beginning of this section referred to “*efficient staffing patterns*” and “*minimising labour costs*”, making the point that traditional rostering models are not designed to take account of the employees viewpoint. A central aim of this thesis is to identify a set of goals which are sympathetic to the needs of the employee, and to incorporate these goals into a system of “employee-friendly” rostering models, thus addressing this important issue of unsympathetic shift patterns. In chapters 4 to 6 we will pursue this aim, but we now complete this chapter with a statement of our motivation.

2.4 Motivation for Study

The previous section confirms the importance of applying staffing and rostering models which are appropriate to the management strategy being pursued, and clearly illustrate an inconsistency between the models traditionally used by call centre operations management, and management strategies which place value on employee retention and service quality.

One effect of the application of techniques drawn from conflicting management philosophies is the high cost of employee turnover, as described earlier. Another effect, according to Kinnie *et al.* (2000), is to cause conflicts and tensions, negatively impacting upon employee attitudes and, hence, service quality. We have a third motivating factor. The UK government’s Health and Safety Executive (HSE

(2006)) point out that although there is no specific health and safety legislation on shift working:

“Employers have legal responsibilities to ensure the health and safety at work of their employees, and this includes removing or controlling the risk of fatigue by organising and planning shift working arrangements.”

The adoption of employee-friendly rostering models will enable employers to discharge this duty of care.

2.5 Chapter Summary

We have described the strategy used by early call centres to obtain competitive advantage through low-cost efficiencies. This strategy is supported by the mass-production model, with standardised transaction processes based on scripting systems, high levels of monitoring and constant pressure to meet numerous targets relating to a range of metrics. HR policies are mainly “administrative”, and operations management systems are based on the maximisation of throughput and the minimisation of staff costs. Jobs are highly stressed and repetitive, though the low skill content means that these operations are designed to be turnover proof.

We have identified an alternative strategy based on revenue generation through customer retention and an expanding customer base, and product cross-selling. This is implemented via the provision of a high quality of service based on the mass-customisation production model, with employee-discretion enabled by flexible Customer Relations Management (CRM) systems and High Commitment

Management HR policies. We have highlighted the need for research into staffing models to support this paradigm at the operations management level.

The motivation for our research is to align the operations management task of employee rostering with the revenue generation strategy. We would expect such a contribution to the consistent implementation of call centre management policy to have a positive impact on absenteeism and turnover, act as an enabling factor in the delivery of service quality, and to improve the social responsibility of call centre management.

There is need for inter-disciplinary links between the research areas of call centre HR and operations management, as highlighted, for example, by the HR scholars Batt & Moynihan (2002):

“In order to examine the relationship between quality and efficiency, we need to better understand the use of technology and operations management in call centres. ...HR and industrial relations scholars...have focused on understanding the organisation of work and HR practices, with little attention to understanding the logic of operations management or the technology that undergirds call centre operations”.

The overall contribution of this chapter has been to initiate such an interdisciplinary connection.

Chapter 3

Employee Friendly Rostering Goals

In the previous chapter we identified the emergence of a call centre management strategy which aims to generate revenue by deploying a knowledgeable and committed workforce in order to deliver high quality customer service in a relatively low-stress working environment. We highlighted the role of an “employee-friendly” rostering system in supporting this business strategy at the operations management level.

In this chapter, we will expand on the concept of “employee-friendly” rostering and identify a specific set of goals for such a rostering system. We derive these goals from three sources:

- First, there are the legislative requirements that must be complied with. For European call centres, these regulations are covered by the European Working Time Directive (93/104/EC).
- Second, there are “good practice” guidelines for roster construction in general, which relate to the health and well-being of the employees. These are summarised by the UK Health and Safety Executive (HSE (2006)).

3.1 Legislative Requirements

- Third, there are those additional factors which, in our own experience, are important to employees and by which they judge the quality of a roster. These represent good practice within the call centre industry, and we refer to them as generalised preferences.

The set of rostering goals which we establish in this chapter underpins an entire rostering methodology, and has the potential to form the basis of an innovative commercial Work Force Management System (WFMS) with particular applicability to those call centres adopting the revenue generation strategy outlined in the previous chapter. Later in this chapter we present an analysis of how all three sets of goals relate to the call centre employee rostering problem. First, we will identify or derive our specific goals.

3.1 Legislative Requirements

The main legislative instrument governing employee rostering in the UK is the European Working Time Regulations¹. There are seven basic provisions, five of which impact on call centre rostering. These form our first set of goals, as listed in Table 3.1.

¹The latest information can be found at www.berr.gov.uk/employment/employment-legislation/working-time-regs/index.html

3.2 Guidelines and Principles

WTR1.	A limit of an average of 48hrs per week per worker.
WTR2.	For night workers, an average of 8hrs work in each 24hr period.
WTR3.	A right to 11 hours consecutive rest between shifts.
WTR4.	A right to a day off each week.
WTR5.	A right to a rest break if the working day is longer than 6 hours.

Table 3.1: Legislative requirements for employee rostering

In addition to these five requirements, all workers are entitled to 4 weeks paid leave per year, and night shift workers are entitled to free health assessments. However, since these two requirements do not concern rostering, we do not include them as rostering goals. A recent (April 2003) amendment to the regulations restricts young workers under the age of 18 to work not more than 8 hours per day or 40 hours per week, and not to be rostered to work between the hours of 11pm and 7am.

3.2 Guidelines and Principles

The UK Health and Safety Executive (HSE) has identified a number of good practice guidelines for employee rostering in general. These guidelines are described in detail in HSE (2006), and are listed on the HSE website¹. For ease of reference, we summarise them below.

- Plan an appropriate and varied workload.
- Offer a choice of permanent or rotating shifts and try to avoid permanent night shifts.
- Rotate shifts every every 2-3 days or 3-4 weeks, otherwise adopt forward rotation.

¹www.hse.gov.uk/humanfactors/shiftwork/goodpractice.htm

3.2 Guidelines and Principles

- Avoid early morning starts.
- Limit shifts to 12hrs, or 8hrs for night shifts.
- Schedule regular breaks, and allow some choice of break timings.
- Consider the needs of vulnerable workers.
- The maximum number of consecutive shifts should be 5-7. Night and early shifts should be restricted to 2-3 day stints.
- Allow two full nights sleep when switching from night to day shifts.
- Build regular free weekends into the schedule.

While some of the guidelines are specific, others are somewhat vague and require clarification, particularly if they are to form goals which we can quantitatively represent in a mathematical model. We now develop this set of goals by referring to primary sources in order to seek clarification, and to identify any additional considerations not mentioned by the HSE. The primary sources for the good practice guidelines are to be found in the ergonomic literature. To quote from the International Ergonomics Association¹,

“Ergonomics is the scientific discipline concerned with the understanding of interactions among humans and other elements of a system, and the profession that applies theory, principles, data and methods to design in order to optimize human well-being and overall system performance.”

The reference to both human well-being and system performance is in line with a central tenet of this thesis: that some rosters are better for the employees than others, and that well-designed rosters have a positive outcome for the business, its customers and employees.

¹The International Ergonomics Association, www.iea.cc

3.2.1 The Ergonomic Background

While the literature on the ergonomic effects of shift work is extensive, few papers relate directly to the issue of roster design and evaluation. A significant amount of this research relates to night shifts. Wedderburn (1981) reported that shift-workers placed a higher value on off-duty time during nights, evenings and weekends than during weekday daylight hours. A more recent study by Baker *et al.* (2003) confirmed that this remains the case. In other words, employees in general prefer not to work night shifts or weekends. One reason given for this is that rotation onto and off of night shifts affects the body clock, and another is that these shifts generally interfere with the employees' social lives. Nachreiner *et al.* (1993) produced a rostering model which recognised this goal, and so was an early precursor of our own approach.

A procedure for the evaluation of rosters, based on 14 weighted criteria, was presented by Schoenfelder & Knauth (1993). Knauth (1993, 1996) proposed a more general set of recommendations relating to night working, length of shifts, start time of morning shifts, consecutive working days, and speed and direction of rotation. Folkard & Tucker (2003) plotted the risks associated with different types of shifts. Their study showed that alertness and productivity reduced at night, and concluded that the length of night shifts, number of consecutive night shifts, and night shift breaks need to be considered together during the rostering process. The consensus at this point is that night shifts are not only unpopular, but they can be harmful to health and productivity, particularly if rostered in poorly designed sequences.

Kundi (2003) identified a number of “comparative laws” as the basis for eval-

3.2 Guidelines and Principles

uating the relative merits of rosters, and derived 10 principles which form the basis of good ergonomic roster design. Many of these principles are similar or identical to those stated in previous research and in the HSE guidelines, while offering more detail. Kundi's paper draws on the previous research in the field - he mentions in particular Folkard (1992), Knauth (1993, 1996) and Schoenfelder & Knauth (1993) - and presents a "state of the art" position. Kundi's 10 principles are summarised as follows:

- The total number of night hours should be as low as possible.
- The necessary night work should be as equally divided among the work force as possible.
- There should be no more than three night shifts in a row, occasionally four, or at most five if this gives rise to other beneficial conditions (such as a better arrangement of days off).
- There should be no more than six working days in a row, occasionally seven if this gives rise to other beneficial conditions.
- After a sequence of night shifts there should be at least 32 and preferably 48 hours off duty.
- Before a shift change of 6 hours or more there should be at least 18 and preferably 24 hours off duty.
- Shifts should be no longer than 12 hours, night shifts should be shorter.
- Morning shifts should not begin before 6:00am, night shifts should not begin later than 11:00pm.
- The number of weekend off duty days should be maximised.
- The number of contiguous off duty days should as often as possible be 2-3 days.

We are now in a position to present our ergonomic rostering goals.

3.2.2 Ergonomic Rostering Goals

Taking the HSE guidelines and the ergonomic sources (and Kundi in particular) together, we are now in a position to state our set of “good practice” rostering goals, which we list in Table 3.2 below. Two terms require definition at this point:

- *Night shifts.* Kogi & Thurman (1993) define these as being at least 7 hours in length, covering the period 24:00 to 05:00.
- *Backward rotation.* This means that an employee has a shift starting earlier on one day than their shift on the previous day. (A *forward rotation* means that an earlier shift is followed by a later one).

ERG01.	Offer a choice of permanent or rotating shifts and try to avoid permanent night shifts.
ERG02.	Earliest day shift start time = 07:00hrs.
ERG03.	Latest night shift start time = 23:00hrs.
ERG04.	Maximum length of night shift = 8hrs.
ERG05.	Plan an appropriate and varied workload.
ERG06.	Minimise the total of night and weekend hours.
ERG07.	Avoid backward rotation.
ERG08.	Schedule night shifts in stints of 2-3 days.
ERG09.	Allow 48hrs off duty after night shifts.
ERG10.	The maximum number of consecutive shifts should be 5-7.
ERG11.	The necessary night work should be as equally divided among the workforce as possible.
ERG12.	Build regular free weekends into the schedule.
ERG13.	Schedule regular breaks.
ERG14.	Allow some choice of breaks.

Table 3.2: Ergonomic principles and practices for employee rostering

We will provide an analysis of the impact of these goals on the rostering process in section 3.5.

3.3 Generalised Preferences

The set of ergonomic goals listed in the previous section relate to those aspects of roster design that impact upon employees well-being. Experience of the UK call centre industry tells us that there are additional criteria by which employees judge the quality of a roster, and as such reflect general preferences. These criteria relate to the desirability of patterns of off-duty days, and to the managing of shift changes, and are expressed as the goals listed in Table 3.3. Experience tells us

GP1.	Maximise tours with weekends off-duty.
GP2.	Maximise tours with consecutive off-duties.
GP3.	Minimise isolated work days.
GP4.	Minimise shift changes on consecutive working days.
GP5.	Where possible, adopt forward rotation.

Table 3.3: Generalised shift and tour preferences in employee rostering

that call centre employees generally prefer to have the same shift start time on consecutive working days, with shift changes taking place after an off-duty period. Goal GP4 expresses this preference, and goal GP5 aims for forward rotation as an alternative when the same start time goal cannot be met. There is some overlap between these goals and the ergonomic guidelines outlined in the previous section. This indicates a general agreement between those aspects of roster design that are in some way *good* for the employees (ergonomics) and what they actually *like* (preferences).

3.4 General Considerations

The ergonomic guidelines and literature give little indication of the importance of taking account of *individual* requests for shift start and finish times or off-duty days. However, it is our perception that call centre employees have an expectation that such requests will be considered during the rostering process. In practice, failure to do so results in numerous exchanges of shifts between employees, which can be difficult to keep track of. We therefore aim to allow for individual requests in our rostering system.

A final point is that there is broad agreement in the ergonomic literature that there is no single “best” roster. Aspects of good practice can often be traded off to improve other aspects. For example, it may be possible to increase the number of off-duties in consecutive pairs, at the expense of longer stints of consecutive attendances. The balance between these factors is a matter of local preference, and a number of researchers (for example Folkard (1992), Kogi (1996) and Baker *et al.* (2004)) recommend worker participation in the roster design process. Our rostering model will therefore need to be flexible enough to allow for the relative prioritisation of the various goals.

3.5 Analysis of Goals

Having identified three sets of rostering goals, we now discuss where in the rostering process these goals need to be modelled. Some of the goals relate to the contractual level and therefore properly belong to HR rather than operations management. This is consistent with our view that, in practice, the successful

implementation of an employee-friendly strategy relies on a co-operative approach between HR and operations management.

Most, if not all of the legislative requirements are implicitly complied with by call centre employment contracts. For example, a typical contract for a full-time employee might specify a 40 hour week, with five 8 hour attendances, each with a 1hr lunch break and two days off duty. This guarantees compliance with WTR1, WTR2, WTR4 and WTR5.

Goals ERG01-ERG04 relate to contractual issues. ERG01 suggests the use of fixed contracts, whereby some employees can work the same days and hours each week. These employees need not therefore be “rostered” each week, rather the coverage they provide can be subtracted from the staffing requirements before the tour scheduling stage commences. Goals ERG02-ERG04 relate to the shift structure which traditionally forms an input to the rostering process, rather than a rostering decision. Hence, while these are not, strictly speaking, rostering goals, it is important that our rostering system is able to handle them as model inputs. The example rosters which we generate in this thesis will all comply with these goals.

Goals ERG05-ERG06 are concerned with planned work-rate, and so relate to the staffing model. As was described in the previous chapter, the adoption of the Quality Driven staffing regime is consistent with goal ERG05. It may be possible to compromise on the service level offered at night and at weekends, thereby reducing the number of employees required during those periods in line with goal ERG06. However, the effects on customer satisfaction of such a compromise are not well researched and, intuitively, may vary between call centres according to the nature of the service provided.

Goals WTR3, ERG07-ERG09 and GP1-GP5 are all concerned with shift sequences. WTR3 imposes a minimum rest period of 11 hours between consecutive shifts. ERG07 calls for us to avoid backward rotation, and is related to two of the general preference goals, GP4 and GP5. The overall preference is for the same shift start time on consecutive working days (GP4), forward shift rotation where this is not possible (GP5), and backward rotation as a last resort. Goals GP1-GP3 imply a similar order of preference for weekends off duty (GP1), consecutive paired offs (GP2), off duties on non-consecutive days (if the first are not possible), then isolated work days as a last resort (GP3). Goals ERG08-ERG09 relate to the structure of tours containing night shifts, where a rest period of 48 hours is recommended following a stint of 2 or 3 night shifts. All of these goals will be included in our tour scheduling and assignment models described in chapters 5 and 6.

Goals ERG10-ERG12 are concerned with the assignment of tours to employees. Goal ERG10 concerns the number of consecutive shifts that can be worked across consecutive weeks. ERG11 and ERG12 relate to the “fairness” of distribution of the unpopular night and weekend shifts. In other words, it is not enough just to *minimise* the total number of these shifts and tours (goals ERG06 and GP1), we need to consider how they are *distributed* among the workforce. These tour assignment issues will be addressed in chapter 6.

Finally, ERG13-ERG14 relate to meal and coffee breaks. Chapter 5 describes how breaks are handled by the tour scheduling model, and our method for assigning breaks to individual employees is described in chapter 6.

3.6 Chapter Summary

The contribution of this chapter has been to identify three sets of rostering goals which represent the human aspects of employee rostering. The goals relate to legislative requirements (WT1-WT5), ergonomic principles and guidelines (ERG01-ERG14), and generalised employee preferences (GP1-GP5). Some of the goals relate to contractual issues, which highlights the importance in practice that a consistent approach is shared by both HR and operations management. Two of the goals relate to staffing models, demonstrating the need for consistency within call centre operations management.

Chapters 5 and 6 of this thesis relate respectively to the two stages of the rostering process with which we are concerned in this thesis, namely the tour scheduling and tour assignment problems. In each of these chapters, we will develop models which implement the 24 rostering goals. Before developing our own models, we next present a review of the literature relating to call centre rostering.

Chapter 4

Literature Review

This thesis addresses the problem of employee rostering in call centres. Our aim is to develop models and approaches to this problem, for call centres adopting a strategy of revenue generation based on high quality of service, in preference to the strategy of cost reduction achieved through efficiency gains. In this chapter we present a review of the literature relating to call centre rostering models, placing emphasis on the connections with management strategy, and with the rostering goals identified in the previous chapter. Our objective in this chapter is to offer an overview of the approaches taken by researchers to date, and to identify research gaps.

4.1 Decomposition

The traditional approaches to call centre rostering attempt to model the trade-off between employee costs, and the cost of “under-staffing” - that is, failure to provide sufficient staff to handle the forecast workload. Typically, the aim

is to allocate shift details (days on and off duty, shift start and finish times, break times, etc.) to individual employees, with regard to legal and contractual rostering obligations, in order to ensure that the offered workload is consistently handled, at the target service level. The complexity involved in this task has led to the decomposition of the whole problem into a series of separate steps.

A model of the overall process of call centre employee rostering was first proposed by Buffa *et al.* (1976) who identified four stages: Forecasting, Staffing, Scheduling, and Assignment. These four stages were outlined in the introductory chapter, when we introduced the employee rostering problem. More recently, Thompson (1995b) identified a subsequent stage regarding the real-time aspects of service delivery. Another variation is offered by Grossman *et al.* (2001) who suggest a stage whereby the service level target for each staffing interval is specified. Ernst *et al.* (2004a) point out that the scheduling stage may involve days-off scheduling, daily shift scheduling, or a combination of both. Since call centre employees are typically contracted to work a fixed number of days per week, the combination of days-off and daily shift scheduling, in order to generate weekly “tours”, has become the norm.

Over time, various interpretations have been placed on the term “rostering”, and different boundaries have been placed on the rostering problem. Our interpretation is that “rostering” relates specifically to the tour scheduling and tour assignment problems. This is in line with the meaning recognised by researchers in other employee rostering domains, such as nurse rostering.

4.2 Forecasting and Staffing

Although this thesis is focused on the tour scheduling and tour assignment problems, the forecasting and staffing problems impact on roster design in that together they provide a key input, that is the number of employees required during each quarter-hour interval. The number of calls arriving at the call centre at any time is essentially outside management control and the forecasting problem is to predict, rather than plan, the workload. Hence, the forecasting problem can be regarded as being independent of the rostering problem. In chapter 2 we outlined the alternative “staffing regimes” and identified that the *Quality-Driven* regime is consistent with the management strategy with which we are concerned in this thesis, but expressed concerns over the applicability of this model to smaller call centres.

4.3 Tour Scheduling

The first employee scheduling problem to be addressed in the mathematical literature was formulated as a Linear Programming model by Dantzig (1954). The objective was to minimize the total number of employees to be assigned to a set of equal-length shifts, while satisfying a minimum staffing requirement in each (half-hourly) interval throughout a single day. This set-covering formulation still forms the basis of most modern models, including our own. The generalization to a tour scheduling environment involves one variable for each valid tour, as described by Morris & Showalter (1983). The difficulty with this formulation is that the size of the model quickly expands and becomes impractical. Brusco &

Jacobs (2000) point out that with flexibility in terms of start times, meal break placements, shift lengths and other factors, set covering formulations expand to "millions or billions" of integer decision variables.

Implicit modeling addresses the problem of model size. The idea is that each possible tour does not need to be individually represented in the model, rather additional variables and constraints are added to ensure that the scheduling requirements are met. In this way, the information requirements of the model are reduced. The output from an implicit model is not in the form of completed weekly tours. Instead, a set of components is produced (for example, a list of daily shifts) that is guaranteed by the model constraints to produce a feasible solution when the tours are subsequently constructed. In other words, some degree of post-processing is required.

Bailey (1985) was the first to use implicit modeling in tour scheduling, combining Dantzig's original set covering formulation of the daily shift scheduling problem with the Baker (1976) formulation of the days-off scheduling problem. He further introduced deviational variables to allow some under-staffing in the solution. His post-processing algorithm allocated shifts to tours so as to minimize the difference in start times across the week.

Burns & Carter (1985) developed a new approach to the days-off scheduling problem, identifying three lower bounds on the total number of employees that would be required. These bounds were determined by the number of employees required to meet the demand during the weekend (given a restriction that each employee receives at least A out of B weekends off), the week as a whole (with a restriction that each employee works 5 days out of 7), and on each day individually. Their scheduling algorithm always reaches the highest of these bounds, and

so always optimizes the size of the workforce. However, no consideration is given to the sequences of shifts within tours, or to the variations in shift start times.

Meal breaks were incorporated into an implicit set-covering IP model by Bechtold & Jacobs (1990). They use *forward pass* and *backward pass* constraints to schedule breaks within an earliest and latest start time window. A variation on this approach was developed by Thompson (1995a), who used a minimum and maximum pre- and post-break work stretch to govern the location of the breaks. Another break scheduling formulation was presented by Aykin (1996), who used additional variables to allocate breaks within a break window for each shift type.

Jarrah *et al.* (1994) extended Burns and Carter's day-off scheduling model to combine with the set covering formulation for daily shift scheduling, and the break scheduling algorithm of Bechtold & Jacobs (1990). This combination of days-off scheduling, daily shift scheduling, and break scheduling is the hallmark of a modern tour scheduling model. However, the issue of shift start time variation is still not considered up to this point.

The issue of shift start and finish times within a tour was first addressed by Jacobs & Brusco (1996). They used implicit modeling in conjunction with the specification of overlapping start-time bands to control the *bandwidth* range of start times allowable within a tour. This type of restriction is common in practice, as it is required to ensure compliance with legislation limiting the minimum time that must elapse between consecutive attendances. Jacobs and Brusco's approach is to impose the limit on the range of start times within a tour as a hard constraint (i.e. a constraint which, if violated, would render the solution unacceptable). Thus, no account is taken of the preference for a narrower range of start times. Brusco & Jacobs (2000) combined this model with the earlier break scheduling

model of Bechtold & Jacobs (1990), and applied the model to the scheduling of employees in a call center.

Applications of an implicit IP approach to tour scheduling in other industries are presented by Bard *et al.* (2003) and Isken (2004). Bard applied the model of Jarrah *et al.* to a problem in the US Postal Service, while Isken applied a model based on that of Jacobs & Brusco (1996) to a tour scheduling problem in the healthcare industry. Although Isken introduced a new formulation for handling start time variations, the principle of scheduling within a time-range is retained.

4.3.1 Key Model

The current standard call centre tour scheduling model in the literature is that of Brusco & Jacobs (2000). This model has built on previous work in the literature and adopts a similar approach based a single objective, which is to minimize the number of required tours. As such, the Brusco-Jacobs model, which is described in detail in Appendix A, is not designed to take account of the range of goals that we identified in Chapter 3.

Our goals GP1-GP3 relate to patterns of off-duty. In the Brusco-Jacobs model, each feasible off-duty pattern is explicitly defined in the model, and the single objective means that undesirable patterns may appear in large numbers in the solution. Thus, no account is taken of the general preference for certain types of patterns of off-duty. We recognise, however, that it would not be difficult to update the Brusco-Jacobs model to enable such a distinction between patterns, by assigning differing costs to tours of each of the specified off-duty patterns.

Goals WTR3, ERG07, GP4 and GP5 relate to shift start time variations

within tours. The Brusco-Jacobs model limits the range of shift start times which can appear in a single tour through a bandwidth parameter. Thus, the model can be constrained to ensure that *every* tour contains daily shifts with a common start time, by setting the bandwidth parameter to 1. Although this conforms to goal GP4, it may prove expensive in terms of the number of tours required, since there is no flexibility to vary shift start times within a tour. On the other hand, if start time flexibility is allowed (by increasing the bandwidth), two important issues arise.

- The model does not take account of the preference for similar shift start times on consecutive working days (GP4). Even a bandwidth setting of only 5 could potentially result in a tour with a different start time on each working day of the week, even though the overall start time range is limited to 1 hour. This level of variation does not conform to our goals, and experience suggests that such a tour would be highly unpopular.
- The model does not distinguish between forward and backward rotation (goals GP5 and ERG07). Thus, with relatively high bandwidth settings, rest periods between successive shifts can potentially be unnecessarily short. Moreover, the legal requirement that a minimum of 11 hours should elapse between successive attendances (goal WTR3) implies that, for example, a day shift cannot immediately follow a night shift. If this condition is imposed, the same constraint would prevent a night shift from appearing *anywhere* in the same tour as an early day shift. Since a night shift can reasonably follow a day shift, the constraint removes flexibility and could lead to poor solutions.

In developing our own model, our intention is to address these shortcomings. As will be seen in the following chapter, our model bears some resemblance to the Brusco-Jacobs model formulated specifically for a “bandwidth 1” solution, but with additional intelligence added in order to handle the rostering goals.

4.4 Tour Assignment

Tour assignment, in contrast to tour scheduling, attracts relatively little attention in the literature. Thompson (1997a) presented a method of tour assignment based on individual requests for specific shifts or off duty days. Requests were handled in strict seniority order, and other assignment goals not considered. Goodale & Thompson (2004) subsequently developed an alternative approach based on productivity criteria. The dearth of assignment models perhaps illustrates the relative unimportance of the tour assignment problem from a management viewpoint. Whereas the tour scheduling problem directly affects staff costs (in that the goal is to minimise the number of tours and hence, employees), the assignment problem does not. The concern of call centre Operations Management has traditionally been to ensure that the correct *number* of employees are scheduled for work at any time, without particular regard to *who* is on duty. By taking account of employee preferences at both the tour scheduling and assignment stages, our aim is to move toward the integration of call centre HR and Operations Management functions.

4.5 Integrated Models

A number of models seek to integrate the various stages of the decomposed problem. Thompson (1997b) presented an integrated staffing and daily shift scheduling model using Integer Programming. This model aimed to minimise the cost of a shift schedule, while meeting a service level target (i.e. a percentage of callers to be answered within a given time threshold) across a whole day, rather than in each interval. However, integrated models have not successfully been extended from daily shift scheduling to weekly tour scheduling. In practice, a number of commercial products integrate all of the staffing, tour scheduling, and tour assignment stages using simulation and iterative improvement. In this approach, a discrete-event simulation model is used to bombard candidate rosters with simulated telephone traffic in order to identify periods when the expected service levels may fall below target. The rosters are then adjusted accordingly, the simulation is re-run, and so on iteratively. Experience indicates that in practice, there are issues with unacceptably long run times. We have, for example, observed instances of a single week roster taking over 12 hours to execute.

The PhD thesis of Canon (2006) addresses both the staffing and tour scheduling stages without integration. The tour scheduling model is based on tabu search, with the objective of minimising the total amount of under staffing, while observing legal and contractual obligations. The length of daily shifts is allowed to vary so that the total staffing provision each week is in line with an “annualisation” process which is used to allocate workforce resources across each week of a long period (3 months is mentioned as an example) so that the correct number of hours are worked by each employee over a calendar year. This annualised

process is applied in some French call centres (which are the main concern of the author), but is not common in the UK. As with other approaches, generalised and individual preferences are not allowed for.

The size of the typical call center rostering problem continues to make integrated approaches computationally impractical, and the staffing problem remains, by and large, a research problem in its own right. The hierarchical decomposition into tour scheduling and assignment stages remains the norm in call centre rostering. This is not, however, the case in nurse rostering and we present an integrated model for the nurse rostering problem in chapter 7.

4.6 Chapter Summary

The call centre tour scheduling models in the literature seek to minimize the size of the required workforce with regard only to a small number of hard constraints relating to shift patterns. None of these models take account of employee preferences or ergonomic rostering factors. Although Thompson (1997a) takes account of individual requests at the tour assignment stage, the choice of tours is limited to those previously scheduled and assignments are made in strict seniority order. There are no models in the call centre rostering literature that take account of a set of goals such as those we identified in chapter 3. This supports our view that current rostering models are designed to support the low-cost efficiency strategy outlined in chapter 2, and that there is a need for a new approach which allows for “human factors” when rostering within the revenue-generation, high-quality strategy. In chapters 5 and 6 we will present our tour scheduling and tour assignment models which address this research gap.

Chapter 5

Tour Scheduling

In this chapter we present our tour scheduling model. The chapter forms the basis of our working paper Glass & Knight (2008a).

We present a flexible approach to tour scheduling that takes account of general preferences for certain types of tour. We use an implicit Integer Programming model, with weighted schedule costs representing generalised employee preferences for patterns of off-duty within tours having the same shift start time on each working day. We allow the number of scheduled shifts to be less than are strictly required to complete an initial set of tours, offering the opportunity of subsequently reducing the number of tours by explicitly combining shifts of different types into single tours. We incorporate a break scheduling methodology which improves upon the standard approach of identifying valid windows for each break, by adding additional constraints which restrict the combinations of breaks that can appear in the same tour. Our model allows for the presence of under staffing in the solution.

Using data from a real-world call centre, we compare the results obtained by

our approach with those from the classic ILP tour scheduling model of Brusco & Jacobs (2000), and from a commercial rostering package widely used in the UK and Europe. We use data from a second call centre to demonstrate that the model is not over-dependent on a specific problem instance. The results indicate that our model generates schedules which have a markedly greater incidence of desirable tour patterns, using no more tours than alternative models.

Our overall approach is as follows. We specify, as model inputs, the valid set of *shift types*, which we define as a combination of a shift length and start time, and the set of *tour patterns*, which are weekly patterns of working days and off-duties. The overall function of the tour scheduling model is to schedule daily shifts, and implicitly match them to the scheduled tours of each pattern. This approach is broadly the same as earlier implicit tour scheduling models, including that of Brusco and Jacobs, and is inherently flexible as it allows for the specification of a wide range of shift types and working patterns, thus enabling the rostering of a mixture of contract types including full and part-time staff.

The benefits of our model centre upon the fast and reliable generation of high quality rosters, taking account of a range of ergonomic rostering principles and generalised preferences for shift sequences and off-duty patterns, while requiring a workforce no larger than that reported by alternative models, which aim solely to minimise the size of the workforce. The employee-friendly nature of our roster solutions can be expected to reduce the high levels of staff turnover and absenteeism experienced by many US and European call centres.

5.1 Tour Scheduling Goals

We now describe how the rostering goals introduced in chapter 3 are handled in our model, drawing comparisons between our approach and others in the literature. The set of goals impacting upon the tour scheduling stage of the rostering process are listed in Table 5.1.

Goal	Description
WTR1	A limit of an average of 48hrs per week per worker.
WTR2	For night workers, an average of 8hrs work in each 24hr period.
WTR3	A right to 11 hours consecutive rest between shifts.
WTR4	A right to a day off each week.
WTR5	A right to a rest break if the working day is longer than 6 hours.
GP1	Maximise tours with weekends off-duty
GP2	Maximise tours with consecutive off-duties
GP3	Minimise isolated work days
GP4	Minimise shift changes on consecutive working days
GP5	Adopt forward rotation where possible
ERG02	Earliest day shift start time = 07:00hrs.
ERG03	Latest night shift start time = 23:00hrs.
ERG04	Maximum length of night shift = 8hrs.
ERG07	Avoid backward rotation.
ERG08	Schedule night shifts in stints of 2-3 days.
ERG09	Allow 48hrs off duty after night shifts.
ERG13	Schedule regular breaks.

Table 5.1: Goals for Tour Scheduling model

Many of these goals relate to the shift structures and tour patterns which form inputs to the tour scheduling model. Goals WTR1, WTR2, WTR4 and WTR5 relate to the number of hours or days to be worked each week. These are contractual rather than rostering decisions, and are imposed in the tour scheduling model through the shift type and tour pattern inputs. Similarly, goals ERG02,

ERG03 and ERG04 relate to shift lengths or start times and are therefore inherent in the specified shift types. These considerations indicate a significant overlap between call centre HR and Operations Management functions. We assume that only one shift can be worked on any day, in any tour, and that the number of shifts to be worked per tour is specified within a “contract” (for example, full time employees).

Two of the goals, ERG08 and ERG09, relate to night shift working. It is our intention to deal with these goals at the tour assignment stage (described in chapter 6), as we wish to retain the flexibility to take account of individual preferences when assigning night shift stints. Although goal ERG01 recommends the avoidance of permanent night shifts, we are aware that there may be some employees, for example those with responsibilities as carers, who may prefer (or need) to work a disproportionate number of night shifts.

The remaining goals relate to patterns of off-duty (GP1, GP2, GP3), shift sequences (WTR3, GP4, GP5, ERG07) or to lunch and coffee breaks (WTR5, ERG13). We now consider each of these three areas in turn.

5.1.1 Patterns of Off-Duty

Previous models make no distinction in terms of the generalised preferences for certain patterns of off-duty (GP1-GP3). This approach is rather black-and-white. Possible tours are either ruled out, in which case they will not appear in the solution, or allowed in, in which case they may appear in large numbers. In reality, there will be some tours, which although generally acceptable, will be more desirable than others from the employees perspective. In practice, this

means that rostering managers need to experiment with numerous “what if .. ?” scenarios, varying the upper and lower limits on the number of scheduled tours of each pattern, in order to identify acceptable trade-offs between roster costs and quality. This process can be very time consuming, and is rather hit-and-miss.

Our approach is to handle the off-duty pattern preferences by classifying the tour patterns according to the placement of the off-duty days, and attaching a schedule cost to each class. In this way, we avoid treating the off-duty pattern preferences as hard constraints (that is, constraints which, if violated, result in an unacceptable solution). We assume that the highest preference in terms of off-duty patterns is for full weekends off (GP1), and that that next preference is for other patterns with consecutive days off duty (GP2). We further assume that patterns with two off-duty days separated by a single working day are undesirable (GP3). This approach offers a lot of flexibility in terms of both the identification of desirable and undesirable shift patterns, and the relative cost weightings placed on those patterns. The benefit of our approach is that the incidence of undesirable tours is minimised in the solution, without incurring the economic cost of excluding them altogether, and that the number of highly desirable tours is maximised.

5.1.2 Shift Sequences

In terms of shift sequences, we assume that it is generally preferred that shift start times should be consistent on consecutive working days (GP4). As has been noted by Jacobs & Brusco (1996), operations experiencing roughly the same demand pattern on each day may well find that the number of tours can be minimised

using only single start-time tours. In call centres, this is typically the case from Monday to Friday, and although Saturday and Sunday will probably differ, the volumes of work are so much smaller than on weekdays that these days can largely be accommodated within a single-start time solution. We also note that a single-start time solution will guarantee that goal WTR3 will be met. This is an important goal, as it is based on a legal requirement. As we explained in the literature review in the previous chapter, the “bandwidth” method used by Brusco & Jacobs (2000) to implement WTR3 has serious drawbacks.

Our model allows the scheduling of incomplete tours, that is, tours containing fewer than the required number of shifts. The idea is that the shifts belonging to partially completed tours may potentially be recombined into a smaller number of tours containing shifts of more than one type. Goals GP5 and ERG07 together indicate that forward shift rotation is generally preferred to backward rotation, and these goals can be taken into account when recombining partial tours. In our empirical experiments, which we illustrate later in section 5.3, we have found that only a handful of tours remain incomplete and that only two or three tours can thus be removed. On the other hand, the presence in the roster of a few “spare” shifts allows additional flexibility when making week-to-week tour assignments or considering individual requests for shifts or absences such as emergency leave, and it may be considered preferable to retain this flexibility rather than reduce the number of tours to a strict minimum. With these factors in mind, we have not developed an algorithm for tour reduction, but rather have carried out the process manually.

5.1.3 Break Times

Goal ERG13 calls for the scheduling of “regular” breaks. The standard approach to break scheduling is to define “windows”, relative to the shift start time, within which breaks can be scheduled to take place. The definition of break windows is central to the formulations of both Aykin (1996) and Bechtold & Jacobs (1990). Windows are typically defined by first taking the most regular spread of breaks across the shift as the “ideal” set of break timings. Flexibility is then introduced by allowing the scheduling of a break within a window extended either side of the ideal timing of that break. A standard example of break windows around ideal times is illustrated in Figure 5.1.




Figure 5.1: Break windows for an 8 hour shift

The drawback of this method used in isolation is that although the defined windows may appear reasonable when considering individual breaks, the combinations of break timings that can be scheduled within the same shift is not restricted. If breaks at extreme ends of adjacent break windows are allocated to the same shift, then very long or short continuous work periods can result, e.g. $2\frac{1}{2}$ hrs and $\frac{3}{4}$ hr respectively in the above example. This is inconsistent with our goal of assigning *regular* breaks (ERG13).

In order to overcome the problem of irregular breaks, we utilise the concept of minimum and maximum work stretch. The placement of the break windows

5.1 Tour Scheduling Goals

in Figure 5.1 correspond to a minimum work-stretch of $1\frac{1}{4}$ hrs and a maximum of 2hrs, together with a lunch break window extended a half-hour either side of the of the ideal. However, the work stretch constraints additionally need to be imposed between breaks 1 and 2, and 2 and 3, in order to avoid extreme combinations.

Note that the allocation of three breaks per eight hour shift (as per Figure 5.1) is common practice in many European and US call centres. There are, however, variations on this practice. For example, some employees may prefer to take a shorter lunch break and finish their shift early. Part-time workers may only receive one short break, or even (subject to goal WTR5) none at all. However, the basic method of identifying break windows and imposing work-stretch constraints is flexible enough to handle such variations.

5.1.4 Individual Preferences

In most call centres, there will be some employees who each work the same hours every week. We handle the scheduling of such restricted-availability employees by subtracting their total provision, per interval, from the overall staffing requirements, before the rostering process is executed for the remaining flexible employees. It is our hypothesis that additional personal preferences, such as specific days off, are easily handled at the tour assignment stage. Moreover, if this is proved not to be the case, such idiosyncratic preferences could be included as subsidiary goals within the objective function of the tour scheduling model, or as additional constraints. The UK Health and Safety Executive guidance, HSE (2006) calls for employees to be given some choice as to their break timings. It

is our intention to comply with this guidance by allocating the breaks to specific shifts only after the shifts have been assigned to individual employees. This approach is thematically consistent with our overall desire to integrate the tour scheduling and assignment processes.

5.2 Tour Scheduling Formulation KG1

The various rostering goals described in section 5.1 above may be encapsulated in a Mathematical Program which we now present. For illustrative purposes, the formulation below is based on full time staff working 8 hours per day, 5 days per week. Each shift incorporates three breaks, of duration 1,4,1 intervals respectively. The break windows are as shown in Figure 5.1, and the minimum and maximum work-stretch restrictions of 5 and 8 intervals, i.e. $1\frac{1}{4}$ and 2hrs, respectively, are imposed. In addition, shift start times are constrained to be consistent within a tour. We have incorporated all other generalised preferences without imposing them as hard constraints. They are considered as goals, which are represented as relatively weighted costs within the objective function. The sensitivity of the model to the values of these costs is discussed later in section 5.2.4. Before presenting the mathematical detail of the model, which we refer to as KG1, we next describe how we handle under staffing.

5.2.1 Under staffing

In call centre rostering, the size of the workforce is often restricted by employee availability (when rostering on a week to week basis) or by budgetary restrictions when planning further ahead. Moreover, it is often the case that as the size of the

workforce increases, a diminishing return is obtained in terms of the additional workload which becomes covered. Also, managers often have the flexibility to cover short periods of under staffing either themselves or with overtime workers. For these reasons, we allow the model to include a certain amount of under staffing. However, the nature of the relationship between workload, staffing provision and service level means that a longer period of a small amount of under staffing is generally less damaging to the service level than a shorter period of high under staffing. Hence, we wish to minimise not only the total amount of under staffing across the whole roster, but also the maximum amount of under staffing in any single staffing interval.

5.2.2 Notation

The following notation is used in the model. Shift types and tour patterns are explicitly represented in the model, and each tour contains shifts of a single type. For illustrative purposes, we are assuming that shifts are of a standard length, and are thus distinguished only by their start time, and may therefore be indexed according to the starting interval of duty. We further assume that only one shift can be worked on any day, in any tour. Note that the break windows are relative to a shift start interval of zero, so that the actual break windows for any particular shift type are calculated by adding the shift index to the base break intervals.

Indices

i	day of week,	$i = 1, 2, \dots, w$
j	shift type,	$j \in \mathcal{J}$
k	tour pattern,	$k \in \mathcal{K}$
π	relative staffing interval,	$\pi \in \Pi$

Sets

$\mathcal{J}$	the set of all shift types
$\mathcal{K}$	the set of all tour patterns
$\mathcal{K}^w$	the set of tour patterns with full weekend off-duty.
$\mathcal{K}^c$	the set of tour patterns with other consecutive off-duty days.
$\mathcal{K}^n$	the set of tour patterns with non-consecutive off-duty days.
$\mathcal{K}^o$	the set of tour patterns containing a single, isolated work day.
$\mathcal{T}$	the set of staffing intervals in a day.
Π_1	the set of starting intervals for break 1, relative to the shift start interval = $\{5, 6, 7, 8\}$.
Π_2	the set of starting intervals for break 2, relative to the shift start interval = $\{12, 13, 14, 15, 16\}$.
Π_3	the set of starting intervals for break 3, relative to the shift start interval = $\{23, 24, 25, 26\}$.
Π	the set of starting intervals for all breaks, for shift type 1 = $\Pi_1 \cup \Pi_2 \cup \Pi_3$

Note that the subsets $\mathcal{K}^w, \mathcal{K}^c, \mathcal{K}^n, \mathcal{K}^o$ partition $\mathcal{K}$

Data

a_{jt} = 1 if shift type j covers interval t , 0 otherwise.

c_{ik} = 1 if day i is a working day in pattern k , 0 otherwise.

r_{it} the number of employees required in interval t of day i .

Decision Variables

$b_{ij\pi}$ the number of employees working shift type j on day i ,
starting a break in interval $j + \pi$.

u_{it} the amount of under staffing (employee intervals) in interval t on day i .

u_{max} the maximum amount of under staffing in any one interval.

x_{ij} the number of employees working shift type j on day i .

z_{kj} the number of employees working shift type j in pattern k .

Schedule Cost Parameters

λ_1 tour patterns with a full weekend off duty.

λ_2 tour patterns with other consecutive off-duty days.

λ_3 tour patterns with non-consecutive off-duty days.

λ_4 tour patterns containing a single, isolated work day.

λ_5 the total number of scheduled shifts.

λ_6 the total amount of under staffing.

λ_7 the maximum under staffing in any interval.

5.2.3 Integer Programming Formulation, KG1

Minimise

$$\begin{aligned} & \lambda_1 \sum_{k \in \mathcal{K}^w} \sum_{j \in \mathcal{J}} z_{kj} + \lambda_2 \sum_{k \in \mathcal{K}^c} \sum_{j \in \mathcal{J}} z_{kj} + \lambda_3 \sum_{k \in \mathcal{K}^n} \sum_{j \in \mathcal{J}} z_{kj} + \\ & \lambda_4 \sum_{k \in \mathcal{K}^o} \sum_{j \in \mathcal{J}} z_{kj} + \lambda_5 \sum_{i=1}^w \sum_{j \in \mathcal{J}} x_{ij} + \lambda_6 \sum_{i=1}^w \sum_{t \in \mathcal{T}} u_{it} + \lambda_7 \cdot u_{max} \end{aligned} \quad (5.1)$$

Subject to

$$\begin{aligned} & \sum_{j \in \mathcal{J}} x_{ij} \cdot a_{jt} - \sum_{\pi \in \Pi} b_{i(t-\pi)\pi} \\ & - \sum_{\pi \in \Pi_2} (b_{i(t-\pi-1)\pi} + b_{i(t-\pi-2)\pi} + b_{i(t-\pi-3)\pi}) \geq r_{it} - u_{it} \quad \text{for } i = 1, \dots, w, \forall t \in \mathcal{T} \end{aligned} \quad (5.2)$$

$$u_{max} - u_{it} \geq 0 \quad \text{for } i = 1, \dots, w, \forall t \in \mathcal{T} \quad (5.3)$$

$$\sum_{k \in \mathcal{K}} c_{ik} \cdot z_{kj} - x_{ij} \geq 0 \quad \text{for } i = 1, \dots, w, \forall j \in \mathcal{J} \quad (5.4)$$

$$\sum_{\pi \in \Pi_1} b_{ij\pi} - x_{ij} \geq 0 \quad \text{for } i = 1, \dots, w, \forall j \in \mathcal{J} \quad (5.5)$$

$$\sum_{\pi \in \Pi_2} b_{ij\pi} - x_{ij} \geq 0 \quad \text{for } i = 1, \dots, w, \forall j \in \mathcal{J} \quad (5.6)$$

$$\sum_{\pi \in \Pi_3} b_{ij\pi} - x_{ij} \geq 0 \quad \text{for } i = 1, \dots, w, \forall j \in \mathcal{J} \quad (5.7)$$

5.2 Tour Scheduling Formulation KG1

$$x_{ij} - (b_{ij7} + b_{ij8} + b_{ij12}) \geq 0 \quad \text{for } i = 1, \dots, w, \forall j \in \mathcal{J} \quad (5.8)$$

$$x_{ij} - (b_{ij8} + b_{ij12} + b_{ij13}) \geq 0 \quad \text{for } i = 1, \dots, w, \forall j \in \mathcal{J} \quad (5.9)$$

$$x_{ij} - (b_{ij15} + b_{ij16} + b_{ij23}) \geq 0 \quad \text{for } i = 1, \dots, w, \forall j \in \mathcal{J} \quad (5.10)$$

$$x_{ij} - (b_{ij16} + b_{ij23} + b_{ij24}) \geq 0 \quad \text{for } i = 1, \dots, w, \forall j \in \mathcal{J} \quad (5.11)$$

$$x_{ij} - (b_{ij5} + b_{ij15} + b_{ij16}) \geq 0 \quad \text{for } i = 1, \dots, w, \forall j \in \mathcal{J} \quad (5.12)$$

$$x_{ij} - (b_{ij5} + b_{ij6} + b_{ij16}) \geq 0 \quad \text{for } i = 1, \dots, w, \forall j \in \mathcal{J} \quad (5.13)$$

$$x_{ij} - (b_{ij12} + b_{ij25} + b_{ij26}) \geq 0 \quad \text{for } i = 1, \dots, w, \forall j \in \mathcal{J} \quad (5.14)$$

$$x_{ij} - (b_{ij12} + b_{ij13} + b_{ij26}) \geq 0 \quad \text{for } i = 1, \dots, w, \forall j \in \mathcal{J} \quad (5.15)$$

$b_{ij\pi}, u_{it}, u_{max}, x_{ij}, z_{kj}$ integer and non-negative

The objective function is made up of a number of terms, each being the product of a cost parameter and a quantity. The weights of the cost parameters $\lambda_1, \dots, \lambda_4$ reflect the relative desirability of each of our classes of tour patterns; weekends off-duty, consecutive off-duties, non-consecutive off-duties and isolated work days. The fifth term places a cost on the total number of scheduled shifts, $\sum_{i=1}^w \sum_{j \in \mathcal{J}} x_{ij}$. The idea is that, within a minimised set of tours, any shifts which are surplus to requirements are excluded from the solution. This offers the opportunity to re-combine any partially completed tours with the aim of reducing the number of tours in total. The final two terms of the objective represent the total amount of under staffing during the week, $\sum_{i=1}^w \sum_{t \in \mathcal{T}} u_{it}$, and the maximum under staffing in any particular interval, $u_{max} = \max_t \max_i u_{it}$, respectively. The weights attached to the cost factors $\lambda_1, \dots, \lambda_6$ can be configured to reflect local

5.2 Tour Scheduling Formulation KG1

preferences.

The first set of constraints, (5.2), ensure that the number of employees provided in each interval covers the number required less any under staffing. These coverage constraints take account of the total number of employees at break, by subtracting the the total number of them in each interval from the total provision, ensuring that the remaining staffing total is sufficient to meet the interval staffing requirement. The $b_{ij\pi}$ variables record the intervals when an employee starts each break, and thus the second, third and fourth interval of the lunch break are handled separately.

Constraints (5.3) determine the maximum level of under staffing in any interval of any day, u_{max} , by requiring that each of the individual interval under staffing values u_{it} do not exceed this maximum. Constraints (5.4) ensure that for each day i , the number of scheduled shifts of type j , x_{ij} , can be accommodated within the tours of type j , with those shift patterns k which include day i as a working day. Note that the number of shifts in a tour is allowed to be less than are strictly required, thus allowing tours to potentially remain incomplete. Constraints (5.5-5.7) ensure that a first, second, and third break are scheduled for each scheduled shift.

The remaining constraints represent the work-stretch requirements. Constraints (5.8-5.11) are the $1\frac{1}{4}$ hr minimum work-stretch constraints, and (5.12-5.15) the 2hr maximum work stretch. Both sets of constraints operate on the basis that if we take, for any shift type on any day, a selection of scheduled breaks from two successive break windows, then if the total number of these breaks exceeds the number of scheduled shifts, two of these selected breaks must appear in the same shift. The constraints therefore preclude those break combinations which

are forbidden to appear in the same shift by limiting their total number to the number of scheduled shifts. For example, constraint (5.8) ensures that for each day and shift type, a first break taken either 7 or 8 intervals after the shift start interval cannot appear in the same shift as a second break starting 12 intervals after the shift start. So the total number of these breaks must not exceed the number of scheduled shifts.

5.2.4 Goal Seeking: model parameterisation

We now discuss the setting of the weights for each of the seven cost parameters in the objective function of KG1. Four of the parameters ($\lambda_1, \dots, \lambda_4$) are attached to the number of scheduled tours belonging to the different classes of off-duty patterns, namely those with the weekend off, other consecutive paired days off, non-consecutive days off, and those with an isolated work day. The weights attached to these parameters reflect the relative desirability of each of these off-duty patterns. Between them they ensure that the total number of tours is minimised. The fifth cost parameter, λ_5 , is attached to the total number of scheduled shifts, and the sixth and seventh are attached to the two measures of under staffing; i.e. the total amount (λ_6), and the maximum in any interval (λ_7).

We begin by considering the relative weightings for parameters λ_1 to λ_4 . In order to investigate these weights in isolation, we set the parameter λ_5 weight to zero and constrained the model to exclude any under staffing, by setting variable u_{max} to zero. We conducted three runs with various setting for parameters λ_1 to λ_4 and the results are shown in Table 5.2, where “W” denotes the weight attached to each of the cost parameters λ_1 to λ_5 . We restricted the execution time for each

5.2 Tour Scheduling Formulation KG1

run to 20 minutes, and all results were within 2.5% of optimum within this period. Our experiments are performed using a standard desktop PC with a 2.67GHz P4 processor, and 512Mb RAM. The interval staffing requirements are taken from a call centre in the UK telecommunications industry, employing a workforce of approximately 200 full-time equivalent (FTE).

	Run 1		Run 2		Run 3	
(number of)	W	result	W	result	W	result
Total Tours		170		192		172
λ_1 Weekends off-duty	1	56	1	101	10	80
λ_2 Other consecutive days off	1	24	10	91	11	52
λ_3 Non-consecutive days off	1	61	100	0	13	40
λ_4 Isolated work days	1	29	1000	0	15	0
λ_5 Shifts	0	850	0	893	0	854

Table 5.2: Solutions for alternative cost parameter weightings, excluding shifts goal

The first run (Run 1) treats all tours as being of equal cost (1), making no distinction between tours with different off-duty patterns, and so is similar to the Brusco-Jacobs model for bandwidth 1. The result was a total of 170 tours, with less than 50% of tours having the desirable consecutive off-duty property. Run 2 addresses this by placing an exponentially increasing weight on the lower quality tours. Such exponential weightings are typical of nurse rostering, and are applied in the benchmark instances addressed in Chapters 7 and 8.

The effects of the weightings for Run 2 were as follows. In terms of the quality of tours, not only the undesirable isolated work days, but even the non-consecutive off-duty patterns were entirely eradicated from the solution, thereby increasing the number of tours with the desirable consecutive-off property from

5.2 Tour Scheduling Formulation KG1

below 50% to 100%. This is feasible since all days and intervals can be covered with tours containing consecutive off-duties. The total number of tours has, however, increased from 170 to 192, a 13% increase which would probably be cost-prohibitive in most practical settings. The trade-off is between the schedule quality and the number of scheduled tours. The Run 1 weightings favour tour minimisation over roster quality, while the Run 2 weightings have the opposite effect.

We therefore wish to find a reasonable compromise between roster quality and tour minimisation. Intuitively, we can specify the weights using a combination of a base cost for each tour, representing the cost of employment, with an additional premium to reflect the desirability of each tour type. In Run 3, we apply a cost of 10 for each tour, with added costs of 0,1,3,5 for each of the patterns of off-duty.

The result of Run 3, in comparison with Run 1, is as follows. The 29 isolated work days pattern scheduled by Run 1 are no longer required, while the number of weekends off has increased from 56 to 80 (i.e. an increase of 43%), and the other consecutive off-duty tours have more than doubled from 24 to 52. This improvement in roster quality has been achieved at the cost of only two more tours. However, we have yet to place any cost on the scheduled shifts.

With the second set of schedule runs (4-6), we introduce a cost per scheduled shift. The cost weightings and results for these exercises are shown in Table 5.3. Essentially, the idea is to remove any shifts from the schedule, which are needed in order to complete a full tour, but which are otherwise surplus to requirements. The aim is to create opportunities to subsequently reduce the number of tours by re-combining partial tours, allowing tours to contain multiple shift types.

5.2 Tour Scheduling Formulation KG1

	Run 4		Run 5		Run 6	
(number of)	W	result	W	result	W	result
Tours		173		194		171
λ_1 Weekends off-duty	1	58	1	102	10	78
λ_2 Other consecutive days off	1	28	10	92	11	52
λ_3 Non-consecutive days off	1	64	100	0	13	41
λ_4 Isolated work days	1	23	1000	0	15	0
λ_5 Shifts	1	842	1	838	1	841

Table 5.3: Solutions for alternative cost parameter weightings, including goal of minimising the number of shifts

In Run 4, 8 fewer shifts were used than in Run 1, at a cost of 3 additional tours. A total of 842 shifts may possibly allow a reduction to 169 tours, although the overall roster quality is poor, with a high incidence of undesirable tour patterns. Run 5 has high roster quality, and the presence of only 838 shifts implies a lower bound of only 168 tours. However, the reduction from 194 to 168 tours will require many of the attractive tours to be re-assembled in less desirable patterns or with unfavourable shift changes, by a subsequent tour reduction process. Run 6 seeks to improve upon the best compromise solution from the first set of runs, i.e. Run 3. The solution of 171 tours, with a high proportion (over 75%) of desirable tours, and none of the unpopular tours with isolated work days, represents a good compromise between roster quality and tour costs. Since the 841 scheduled shifts can be accommodated within 169 tours, there is also an opportunity to reduce the number of scheduled tours by two.

5.3 Quality of the Solution

We now compare the tours scheduled by our model with those from the Brusco and Jacobs model, and with those produced by a commercial rostering package. In all examples, we use a dataset of interval staffing requirements obtained from the telecommunications call centre introduced earlier. The centre operates 24 hours per day, 7 days per week, in a discontinuous operation where shifts are scheduled to begin and end within a 24hr period running from 07:00am on each day, to 07:00am the following day. We have assumed that the service is to be staffed by full time employees, each working five 8 hour shifts per week, with each shift having 3 breaks, of 15 minutes, 1 hour, and 15 minutes respectively, to be scheduled within the fixed time windows of Figure 5.1. We have further assumed that shifts can start in any interval from 1 to 65, with interval 1 corresponding to 07:00hrs, and 65 to 23:00hrs.

5.3.1 Comparisons with Brusco & Jacobs model: A Telecomms call centre

The first set of exercises compare results obtained from our model, KG1, with those of the tour scheduling model of Brusco & Jacobs (2000). The Brusco Jacobs model is described in detail in Appendix 1. Their formulation, referred to as P2, does not allow for under staffing, or for minimum and maximum work-stretch. In order, therefore, to generate a “like-with-like” comparison, we disabled these aspects of our own model by setting the maximum understaffing variable, u_{max} , to zero. We set the weights for cost parameters λ_1 to λ_5 to 10,11,13,15,1 respectively, as in “Run 6” above. Since consistency in shift start times within tours is one of our rostering goals, we initially set the P2 “bandwidth” to 1, meaning that no

5.3 Quality of the Solution

shift start time variation is allowed within any tour.

A summary of the solutions is given in Table 5.4. The initial solution found by KG1 contains more of the desirable weekends and consecutive-paired off duties than the bandwidth 1 solution for P2, and has eradicated the unpopular isolated work days, using no more tours than P2. In addition, the KG1 solution offers additional flexibility as it contains 14 “spare” shifts, in other words there are 14 opportunities to allow for requests for absences for leave or training, etc, without breaking any coverage constraints.

	KG1 (initial)	P2 (b=1)	KG1 (reduced)	P2 (b=5)
(number of)				
Tours	171	171	169	not
Shifts	841	855	841	found
Weekends off-duty	78	60	76	
Other consecutive off-duties	52	19	42	
Non-consecutive off-duties	41	60	48	
Isolated work days	0	32	3	
Backward rotations	0	0	0	
Forward rotations	0	0	3	

Table 5.4: Comparison of results from KG1 and P2 for “Telecomms” call centre

If, on the other hand, it is considered preferable to reduce the number of required tours by allowing some start time variation within tours, the KG1 solution can be adjusted by combining some of the scheduled tours together to remove the spare shifts from the solution. The manually obtained result from this example was to reduce the number of tours from 171 to 169, by condensing 11 partially completed tours to 9 tours. The relevant shift and tour details are given in Appendix B. Of the 9 “reduced” tours, 3 contain a forward rotation (i.e. a shift change takes place across consecutive working days, but the change is to a later

shift start time), and 3 have a shift change following an off-duty day. None of the tours contain more than 1 shift change. Note that shift changes following a period of one or two off-duty days are considered to attract no additional cost, since this is the preferred time for shift changes to take place. Four tours remain incomplete, offering some flexibility in handling requests for emergency leave. The nature of the shift changes means that the minimum time between ending one shift and starting the next is unchanged from the “bandwidth 1” assumption, highlighting the weakness of the Brusco-Jacobs approach, which makes no distinction between forward and backward rotations, or shift changes following off duty days.

To make a comparison of this new solution for 169 tours, we increased the P2 bandwidth to 5 (intervals), allowing a start time variation of up to 1 hour in a tour. It should be noted that even this relatively small bandwidth could result in a tour containing a different shift start time on each working day, which experience suggests would be highly undesirable. However, the model was unable to improve upon the 171 tours even after 24hrs run time. This supports the findings of Aykin (2000), that the break scheduling formulation of Bechtold & Jacobs (1990) which forms part of P2, is not as effective as his formulation which is similar to our own.

5.3.2 Comparisons with Commercial Package: A Telecomms call centre

In order to assess the results of our model, fully extended to allow for understaffing and to include work-stretch constraints, we now draw comparisons with the results obtained using a commercially available rostering package: the Q-

5.3 Quality of the Solution

Max WorkForce Management System¹. Q-Max is the most widely installed call centre work-force management system in the UK², and is used in over 40 countries around the world. The Q-Max scheduling system is based on local-search heuristics, using a construction and iterative improvement algorithm.

For the purposes of the comparative exercise, we used the same staffing requirement dataset as in the previous exercises. We restricted the number of tours to 165, thus forcing under staffing into the solution and representing a typical weekly rostering problem, where the number of available employees is known. The aim of this exercise is to compare the results not only in terms of the quality of the rostered tours, but also in terms of the under staffing. The Q-Max scheduling system has a single objective which is to minimise the sum-of-squared deviations between staffing requirements and provision in each interval. The direct penalisation of under staffing in both Q-Max and KG1 allows a comparison in these terms.

As with the Brusco-Jacobs model, shift and tour details in Q-Max are treated as hard constraints, and no qualitative distinction is made between the different valid patterns of off-duty. We configured the Q-Max system to allow tours to contain any two days off from seven, and to allow start time variations within a bandwidth of 5. The system allows for minimum and maximum work-stretch restrictions, but does not allow a user-defined lunch window for each shift type. This gives rise to the possibility of scheduling lunch breaks which, by our own definition, are too early or too late. Hence, the Q-Max method has a slight advantage in terms of flexibility, but we report the number of lunches scheduled

¹Q-Max Systems Ltd, www.q-max.co.uk

²according to the Professional Planning Forum, 2008

5.3 Quality of the Solution

“out of window”.

The results of the exercise are shown in Table 5.5. The KG1 solution is better in all respects. KG1 was able to cover the staffing requirements more accurately than the commercial package, in terms of both the total under staffing (employee-intervals) during the week, and the maximum level of under staffing during any interval. If we assume that understaffed intervals would be covered with overtime working, then the KG1 solution represents a saving of over 25 hours of overtime compared to the Q-Max solution.

At the same time, KG1 produced the superior results in terms of the quality of the scheduled tours. A large majority of the package generated tours contained more than 1 shift type, with 46 tours having 3 shift types, and 7 tours having 4 shift types. Many lunch breaks were scheduled at extreme ends of the minimum or maximum work-stretch limits, placing them outside of our lunch window and hence not meeting our goal of scheduling “regular” breaks (ERG13).

(number of)	KG1	Q-Max (bw=5)
Tours	165	165
Shifts	825	825
Weekends off-duty	74	60
Other consecutive off-duties	35	19
Non-consecutive off-duties	53	60
Isolated work days	3	32
Lunches out of window	0	>300
Tours with > 1 start time	0	124
Total under staffing (intervals)	61	163
Maximum under staffing	2	4

Table 5.5: Comparison of KG1 and Q-Max for “Telecomms” centre

5.3.3 Comparison with Brusco & Jacobs: A Financial call centre

In order to demonstrate that the high quality solutions outlined above are not overly dependent on the specific call centre data used, we now repeat the comparative exercise for a second call centre. This call centre is also a 24hr, 7 day operation, this time belonging to the finance sector. We make the same assumptions regarding shifts and tour patterns as previously, and run the same exercise as in 5.3.1, keeping the weights at the default values. The results of this exercise are given in Table 5.6, alongside those obtained from the Brusco-Jacobs model. As before, all runs were restricted to 20 minutes.

	KG1 (initial)	P2 (b=1)	KG1 (reduced)	P2 (b=5)
(number of)				
Tours	115	115	112	not
Shifts	556	575	556	found
Weekends off-duty	61	45	58	
Other consecutive off-duties	37	14	27	
Non-consecutive off-duties	17	38	26	
Isolated work days	0	18	1	
Backward rotations	0	0	0	
Forward rotations	0	0	3	

Table 5.6: Comparison of results from KG1 and P2 for “Bank” call centre

The results are similar to those from the previous exercise in that the initial solution to KG1 (all tours with a single shift start time) is of a higher quality than the Brusco-Jacobs “bandwidth 1” solution in all respects. The number of required tours in the KG1 solution can be reduced from 115 to 112 by allowing 10 tours to contain two shift start times. Of these 10 tours, 3 have a forward rotation (i.e. a change from one shift start time to a later one, on consecutive

working days), while the other 7 contain a shift change following one or two days off-duty which is the preferred time for changes to take place. The partial tours from the initial solution, together with the “reduced” tours, are given in Appendix B. Following the reduction to 112 tours, the solution quality in terms of off-duty patterns remains superior to the Brusco-Jacobs “bandwidth 1” solution. No feasible solution for the Brusco-Jacobs “bandwidth 5” model was found after 20 minutes.

5.3.4 Fitness of the model: practical requirements

We should point out that none of the above results from KG1, P2, or Q-Max represent optimal solutions. The KG1 results were, however, all within a tolerance gap of 2.5% (i.e. they were at worst within 2.5% of optimal, but may have been closer). The question arises as to whether this level of error is acceptable in practice. In order to address this, it is necessary to offer some explanation of the set of call centre planning tasks known collectively as “work-force management” (WFM), of which the rostering task forms part. In the context of our interest in rostering accuracy, we need to focus our attention on the *timescales* in which the various WFM activities are performed.

The rostering process is typically carried out around four to six weeks ahead of time. In other words, each week, an employee is notified of their tour assignment for a single week which is 4-6 weeks ahead. This represents a trade-off; the employees would generally prefer as much notice as possible in order to plan their social lives. On the other hand, a closer rostering horizon allows shorter-notice requests for leave and absences to be taken into account during the rostering process.

5.3 Quality of the Solution

The planner would prefer, by and large, to delay until the last minute in order to reduce uncertainty.

The production of the rosters is, however, not the end of the WFM process. Between the time that the roster is produced and the time it is to be worked, either the workload forecast or the staff availability (or both) may change. Call centres therefore have the choice of allowing for uncertainty by building additional slack resources into the workforce, or handling the uncertainty by retaining the flexibility to adjust the roster in the light of changing circumstances. Most call centres adopt the latter approach. One source of flexibility is to treat a percentage of employees as “reserves”, that is, staff who will be scheduled only at short notice. Other sources of flexibility are applied on the day, or at one or two days notice. For example, breaks and other scheduled activities such as meetings and training sessions can be re-scheduled, overtime can be assigned, and managers and team leaders can lend assistance. In addition, calls can be routed in real-time to back-office support staff, or the workload re-balanced between groups of employees. The degree of flexibility afforded by these planning and control methods is generally assumed to allow planners sufficient scope to adjust the rosters in the light of variations from the original plan. In order to retain control, it is necessary for the degree of roster *flexibility* to outweigh the *variability* in the workload forecast and staff availability. This concept is well understood in the field of cybernetics and is referred to as the *Law of Requisite Variety*, Ashby (1956)¹.

In short, we can conclude that it is not possible to produce a fully optimal

¹available in electronic form as a PDF from Principia Cybernetica, <http://pespmc1.vub.ac.be/ASHBBOOK.html>

roster 4-6 weeks ahead of time, since we will not be in possession of the full details of the problem at that point in time. The aim, therefore, is to produce a roster which is accurate *to the best of our knowledge* at the time of production, such that the roster can be kept “on track” as additional information comes to light. The quantification of the degree of variety within call centre rostering has not, to our knowledge, been studied in detail and we suggest that this would be an interesting and fruitful area for further study.

5.4 Chapter Summary

The research literature on the ergonomics of rostering identifies numerous good practices, and highlights the problems that employees may encounter when these practices are not followed (see for example Totterdell (2005)). To date, however, little research has been conducted as to how these rostering principles can be represented in call-centre rostering models.

We have presented an implicit Integer Programming tour scheduling model which takes account of generalised employee preferences and ergonomic principles. Based on the evidence of two real-world call centre examples, our model produces solutions containing a higher proportion of desirable tour patterns than the model of Jacobs and Brusco, using no more tours. The model takes advantage of the problem structure in terms of the relationship between numbers of shifts and tours, enabling us to use an implicit ILP representation to schedule tours with consistent shift start times. It is our intention to develop an extension of our tour scheduling model to allow for multiple employee skills. The goals represented in our model relate to the patterns of off-duty days and shift start times within

5.4 Chapter Summary

tours, and breaks within shifts. In addition, any “spare” shifts which remain in the tour schedules will add flexibility at the assignment stage, and increase the scope to accommodate individual requests and personal preferences.

The contribution of this chapter is to quantify and realise important employee-friendly criteria in terms of patterns of shift start and finish times, days on and off-duty, and the timing of meal and coffee breaks. We have demonstrated that the quality of employee rosters can be improved with no detriment to customer service levels and at no additional staff cost.

In the next chapter, we will apply a similar approach to the tour assignment problem, thus facilitating the integration of the tour scheduling and assignment stages of the rostering process.

Chapter 6

Tour Assignment

Following the construction of a set of weekly tours, the tours must be assigned to individual employees. In this chapter we present a comprehensive tour assignment model based on the goals identified in chapter 3. We consider week-to-week continuity in terms of the number of consecutive days worked across weeks, and take into account any end of week shift changes. We aim to distribute the tours and shifts with unsocialable hours, in particular weekend and night shifts, in accordance with the ergonomic goals. We intend that each employee can express requests for shift start times, lunch times and off-duty days.

We assume that the number of available employees and the number of scheduled tours are equal. If this is not the case and the number of employees is greater, then additional tours can be added as required. The number of scheduled tours would normally be constrained to not exceed the number of employees. Our assignment process has three stages:

- In a *pre-assignment stage*, we *re-distribute the night shifts* by exchanging stints of night shifts between tours. The aim is to provide a more even

spread of night shifts among the workforce, and to avoid assigning more than one stint of night shifts to the same employee in a single week.

- At the *main assignment stage*, we *assign the tours to the employees* using the Hungarian method for the two-dimensional assignment problem, Kuhn (1955), which is a polynomial time algorithm of $\mathcal{O}(n^3)$ in the number of employees. Our cost function takes account of week-to-week continuity, and a number of “fairness” issues.
- In a *post-assignment stage*, we first attempt to *meet any individual requests* by exchanging stints or shifts between employees. Finally, we *assign breaks to employees*, generally on the basis of “first on duty, first to break”, but also taking account of any requests for lunch breaks at specific times.

These stages are described in the following three sections.

6.1 Pre-Assignment process

The purpose of carrying out this step is to re-distribute the night shift stints among the tours. The aim is twofold: to keep the number of consecutive nights in a stint small, and to avoid the need to assign two stints of night shifts to the same employee in the same week. This is in line with goals ERG08 which states that night shifts should be scheduled in stints of 2-3 days, and ERG10 which states that the necessary night work should be as evenly distributed among the work force as possible. In addition, we need to take account of goal ERG09 which recommends a minimum period of 48hrs off duty following night shifts.

Recall that the output from the tour scheduling model, described in the previous chapter (5), is a set of tours each of which consists of stints with identical start times on each working day of the week. Therefore, night shifts appear in tours consisting exclusively of night shifts and two (or more) off-duty days. The mechanism for achieving the above goals is to combine an all-night tour from the tour scheduling stage, with an all-day tour by exchanging stints between tours.

In order to be able to exchange stints between tours in such a way as to meet our goals for night working, night shifts must only be scheduled in those tour patterns which will allow the successful exchange of stints. At the tour scheduling stage, we therefore prevent night shifts from appearing in tours with an isolated work day, thus ensuring a minimum length of 2 for night shift stints, in accordance with goal ERG08. For simplicity, we also rule out the possibility of scheduling a single night shift at the beginning or end of the week, as this will raise issues of week-to-week continuity. For each tour of five night shifts produced by the tour scheduling module, we wish to identify a day tour pattern with which to exchange shifts.

We now present an analysis of the tour patterns which we are able to recombine in order to meet our goals relating to night shift work. The night shift tour patterns are listed in table 6.1, where “N” denotes a night shift, “D” a day shift (which refers to any shift other than nights), and “O” an off-duty day. The list is not exhaustive and some of the night shift patterns can be matched with alternative day shift patterns. However, since the number of night shift tours will invariably be small compared to the number of days shift tours, it is not difficult to identify suitable pairings. Once the night shift stints have been redistributed,

6.2 The Tour Assignment Model

Night Tour		Day Tour		New Tour 1		New Tour 2
OONNNNN	+	DDDDDOO	$\mapsto$	OODDDNN	+	DDNNNOO
NNOONNN	+	DDOODDD	$\mapsto$	NNOODDD	+	DDOONNN
NNNOONN	+	DDDOODD	$\mapsto$	NNNOODD	+	DDDOONN
NNNNNOO	+	DDOODDD	$\mapsto$	NNOODDD	+	DDNNNOO
ONNONNN	+	DDDOODD	$\mapsto$	DNNNOODD	+	ODDONNN
ONNNONN	+	DDDDOOD	$\mapsto$	DNNNOOD	+	ODDDONN
ONNNNNO	+	DDDOODD	$\mapsto$	ODDNNNO	+	DNNOODD
NNONNNO	+	DDOODDD	$\mapsto$	NNOODDD	+	DDONNNO
NNNONNO	+	DDDOODD	$\mapsto$	NNNOODD	+	DDDONNO

Table 6.1: Tour combinations for night shift re-distribution

we can proceed with the tour assignments.

6.2 The Tour Assignment Model

The standard mathematical model for assigning n tours to n employees is as follows:

Minimise

$$z = \sum_{i=1}^n \sum_{j=1}^n c_{ij} \cdot x_{ij} \quad (6.1)$$

subject to

$$\sum_{i=1}^n x_{ij} = 1 \quad \text{for } j=1, \dots, n \quad (6.2)$$

$$\sum_{j=1}^n x_{ij} = 1 \quad \text{for } i=1, \dots, n \quad (6.3)$$

$$x_{ij} = 0, 1 \quad \text{for } i, j=1, \dots, n \quad (6.4)$$

where c_{ij} is the cost of assigning the j^{th} tour to the i^{th} employee, and

$$x_{ij} = \begin{cases} 1 & \text{if employee } i \text{ works tour } j \\ 0 & \text{otherwise.} \end{cases}$$

The objective function minimises the total assignment costs. Constraint (6.2) ensures that each tour j is assigned to one and only one employee, and constraint (6.3) ensures that each employee i is assigned one and only one tour. In order to apply the mathematical model, we need a function to calculate the cost c_{ij} of assigning each of the tours j to each of the employees i . The cost is designed to represent the rostering goals that relate to week-to-week continuity, and also to represent the fairness of distribution of unsociable night shifts and weekend working. We consider these goals in the following subsection.

6.2.1 Tour Assignment goals

In chapter 3 we identified the complete set of goals for our rostering model. In Table 6.2 we list those goals which apply to the tour assignment stage of the rostering process, either in relation to week-to-week continuity, or to “fairness”.

When assigning a tour to an employee, we need to take account of the continuity from the previous week’s assignment to the current assignment. For example, we need to ensure that a minimum period of 11 hours elapses between a shift finishing on the last day of the previous week, and a shift starting on the first day of the current week, in line with goal WTR3. Similarly, it may be possible to create additional off-duty pairs (goal GP2) by assigning a tour which begins with an off-duty day, to an employee who ended the previous week with an off-duty

6.2 The Tour Assignment Model

Goal	Description
a) week-to-week continuity	
WTR3	Employees have a right to 11 hours rest between attendances
GP2	Maximise tours with consecutive off-duties
GP3	Minimise isolated work days
GP4	Minimise shift changes on consecutive working days
GP5	Adopt forward rotation where possible
ERG08	Schedule night shifts in stints of 2-3 days
ERG09	Allow 48hrs off duty after night shifts
ERG10	The maximum number of consecutive shifts should be 5-7
b) fairness	
ERG11	The necessary night work should be as equally divided among the workforce as possible
ERG12	Build regular free weekends into the schedule

Table 6.2: Goals for Tour Assignment model

day. The first eight of our assignment goals relate to week-to-week continuity, and the last two to the “fairness” of distribution of unsociable night and weekend shifts among the workforce.

First we must decide how to penalise deviations from our rostering goals. In the tour scheduling model, we represented the goals as either hard or soft constraints. Hard constraints are those rules which must be complied with, and any violations result in an unacceptable solution. Soft constraints are those rules which should be met if possible, with a penalty attached to violations, and where the relative weights of penalties reflect the relative importance of the goal. For consistency, we wish to impose the same penalty structure in this tour assignment model as we did previously in the tour scheduling model.

6.2.2 Hard constraints

We consider goals WTR3, ERG08 and ERG09 as having the highest priority of the tour assignment constraints. Goal WTR3 was imposed as a hard constraint in the tour scheduling model, and the goals relating to night shifts (ERG08 and ERG09) were handled by the pre-assignment process described in section 6.1. However, goal ERG10 was not considered during tour scheduling since we were assuming a 5-day working week, and in any case it is not possible to schedule more than 7 days in a week. It is, however, an important issue in tour assignment since, even with a 5-day week, it would be possible to assign consecutive tours which result in a stint of 10 working days. Our experience is that UK call centre agents do not generally expect to work more than 6 consecutive days. For illustrative purposes, we will therefore treat 7 days as an “absolute” limit (“hard” constraint), and penalise a 7 day stint by means of a soft constraint in section 6.2.3 below.

We impose the “hard” constraints by assigning a penalty cost to the term in the objective function, which is high enough (relative to the soft constraint penalties) to ensure that a hard constraint will not be violated unless there is no alternative. We denote this penalty weight as w_0 . Our set of hard constraints, together with the associated goal, measurement of deviation and the penalty weight is shown in Table 6.3.

The question arising at this point is how do we know that there will be a solution that satisfies all of the hard constraints? Our idea is that once the assignments have been finalised for any week, we can analyse the completed roster in order to identify any additional constraints which need to be imposed on the tour scheduling model when generating the tours for the following week, in order

6.2 The Tour Assignment Model

	Hard Constraint	Goal	Violation	Weight
H_1	A minimum of 11 hours between attendances	WTR3	less than 11 hrs	w_0
H_2	Night shifts should be assigned in stints of 2-3 days	ERG08	1,4,5,6 or 7 shifts	w_0
H_3	A minimum of 48hrs off duty after a stint of night shifts	ERG09	less than 48 hrs	w_0
H_4	The maximum number of consecutive shifts = 7	ERG10	more than 7 shifts	w_0

Table 6.3: Hard constraints for Tour Assignment model

to guarantee a feasible continuation. For example, any employees ending a night shift stint on Sunday will need the following Monday and Tuesday off, so the set of tours scheduled for the following week must include a number of such tours at least equal to the number of Sunday night shifts. We observe, however, that in call centres, the number of shifts required to meet the customer demand on Sunday is very low in comparison with Monday, so that there is a great deal of choice in the assignment of tours for those few employees who worked on the previous Sunday. In our experiments, this amount of flexibility has always been sufficient to ensure that a feasible set of assignments can be found, without the need to impose additional constraints during tour scheduling. This is not, however, our experience when addressing the nurse problem, and in chapter 7 we describe how we handle the continuity between rostering periods.

6.2.3 Soft constraints

The remaining goals for tour assignment (listed in Table 6.2) are the “generalised preference” goals GP2-GP5 and ERG10-ERG12. At this assignment stage, the achievement of goals GP2-GP5 and ERG10 is affected only by week-to-week con-

tinuity considerations. The remaining two goals, ERG11 and ERG12, relate to the fairness of distribution of undesirable night and weekend shifts among the workforce. We now explain how we handle the week-to-week, and fairness goals as soft constraints. The calculations required to implement these constraints are given in section 6.2.4.

6.2.3.1 Week-week continuity

Goal GP2 reflects the desirability of consecutive off-duty days. If a tour with a single off-duty on Monday is assigned to an employee who had an off-duty day on Sunday of the previous week, then we will have created an additional off-duty pair in line with goal GP2. Since this is desirable, we associate a penalty weight w_1 with the assignment of a tour with a single off-duty on Monday, to an employee who was working on the previous Sunday. Likewise, we penalise the assignment of a tour with a Monday shift to an employee who had a single off-duty on the previous Sunday.

We also need to avoid creating isolated work days, in accordance with goal GP3. If an employee ended the previous week with Saturday off-duty and a Sunday shift, then assigning a tour with Monday off would create an isolated work day on Sunday. Similarly, it is possible to create an isolated work day on the current Monday. Any such assignments will incur penalty w_2 .

Goals GP4 and GP5 relate to shift changes on consecutive working days, and affect assignments of tours with Monday shifts, to employees who worked the previous Sunday. Our aim is to exactly match the two shift types so that no shift change takes place, in accordance with goal GP4. If this is not possible, we prefer a forward rotation (goal GP5), with a small penalty weight w_3 . The last

resort is to allow a backward rotation, which attracts a larger penalty weight w_4 . The hard constraint relating to the minimum rest period is a limiting factor. We attach weight w_5 to an assignment which results in 7 consecutive working days across the previous and current weeks, in accordance with goal ERG10.

6.2.3.2 Fairness

Goals ERG11 and ERG12 relate to the “fairness” of distribution of the unsociable night and weekend shifts. We approach the modeling of the fairness issues by maintaining a number of items of data for each employee. These “counters” reflect the number of times that night shift stints or weekend shifts have been assigned to an employee. In looking for fairness, we are not particularly interested in the *absolute* numbers of shifts worked (over any period), rather the number of shifts worked by each employee *relative* to the others. For each employee, we keep account of four specific items:

1. The relative number of stints of night shifts worked
2. The relative number of Saturdays worked, where Sunday was off-duty
3. The relative number of Sundays worked, with Saturday off duty, and
4. The relative number of full weekends (Saturday and Sunday).

We associate penalty weights w_6 - w_9 with the assignment of night stints, Saturday, Sunday shifts, and full weekends. Each week, after the tour assignments have been made, we update the values of the four counters for each employee. For each counter in turn, we first ascertain the lowest value for any of the employees. We then subtract this minimum value from the appropriate counter for each

6.2 The Tour Assignment Model

employee, so that the minimum value for any employee is now zero, and the values stored for the rest of employees are relative to zero. This approach has two important benefits:

1. There is no need to identify a period over which the measures are taken.
2. New recruits will start with a zero counter and will quickly blend in.

The full set of soft constraints, together with the measurements of deviation and the penalty weights (W) is shown in table 6.4.

	Soft Constraint	Goal	Violation	W
a) week-to-week continuity				
S_1	Consecutive off-duty days (preferred)	GP2	Single off Sun, working Mon, working Sun, single off Mon	w_1
S_2	Isolated work days (to be avoided)	GP3	Isolated Sun or Mon shift	w_2
S_3	Forward rotation	GP4	Mon shift starts <i>later</i> than Sun shift	w_3
S_4	Backward rotation	GP5	Mon shift starts <i>earlier</i> than Sun shift	w_4
S_5	Number of consecutive shifts worked = 2 to 6	ERG10	Seven consecutive shifts worked	w_5
b) fairness				
S_6	Equal number of stints of night shifts to be worked	ERG11	Additional stints relative to lowest	w_6
S_7	Equal number of Saturday shifts to be worked	ERG12	Additional shifts relative to lowest	w_7
S_8	Equal number of Sunday shifts to be worked	ERG12	Additional shifts relative to lowest	w_8
S_9	Equal number of full weekends to be worked	ERG12	Additional shifts relative to lowest	w_9

Table 6.4: Soft constraints for Tour Assignment model

Some of the soft constraints represent what we have termed “generalised preferences”, which relate to the criteria by which call centre employees judge roster quality *as a group*. One aspect of the assignment process which we have yet to address is the need to take account of *individual* employee preferences. These can be in the form of requests for specific shifts or off-duty days. Our approach to handling such requests is through the exchange of shifts and stints between individual employees once the initial assignments have been made. In other words, the focus of the main tour assignment process is on the ergonomic principles and generalised preference goals outlined in Table 6.2, while any individual requests are handled subsequently by identifying mutually beneficial exchanges. One reason for leaving these exchanges until an initial tour assignment has been made is that this will offer more flexibility, since it may be possible to comply with some requests by exchanging a single shift or a working stint, rather than re-assign a whole tour. We will offer other reasons when we later discuss the post-assignment process.

6.2.4 The Cost Function

Having defined the penalties associated with violations of the hard and soft constraints, we now explain how we apply them in order to calculate the total cost c_{ij} of assigning tour j to employee i . We approach this by considering each employee i in turn, and calculating the cost of assigning to the employee each tour j in turn. In order to evaluate any violations relating to week-to-week continuity, we need to know what the employee was doing at the end of the previous week, and the shift start times and off-duty days for each of the tours under consideration.

Specifically, we need to know:

A) For each employee:

1. The shift type worked on Sunday of the previous week (if on duty).
2. The length of the final stint of the previous week, i.e the number of consecutive days either worked or off-duty.

B) For each tour:

1. The shift type scheduled on Monday (if any) in each tour.
2. The length of the stint of duties (or off-duty days) at the beginning of the week.

Any violations of the hard constraints (H1-H4), or the soft constraints relating to week-to-week continuity (S1-S5), can be evaluated using only this information. Penalties w_0 to w_5 are imposed where violations are detected. The function for applying these penalties is a simple binary decision (each of the relevant constraints is either satisfied or violated).

In terms of the soft constraints relating to fairness (S6-S9), we consider each of the four criteria, namely night shift stints, Saturday shifts, Sunday shifts and full weekends. For each employee i we define data elements e_{in} to hold the relative numbers of these four types of shift (as described earlier), with $n = 6, \dots, 9$ respectively. For each tour j , we define data attributes t_{jn} taking values 0 or 1 according to the presence or absence of the appropriate types of shift. The cost of the fairness aspect of any assignment, i.e. the penalties associated with soft constraints S6-S9, is $\sum_{n=6}^9 e_{in} \cdot t_{jn} \cdot w_n$.

The total cost of assigning the tour is then calculated by summing the penalties relating to week-to-week continuity and fairness. There is a lot of flexibility in the way the weights are configured. The relative values of the weights can be balanced to reflect local preferences. This flexibility is, in itself, one of rostering goals identified in chapter 3 (described under “general considerations”).

6.3 Post-Assignment processing

At this point in the rostering process, a number of individual requirements will have been taken into account. Employees working fixed, as opposed to flexible, contracts will have been allocated their fixed duties from the outset. Those employees on leave, or otherwise unavailable would not have been rostered on duty. However, once the initial tour assignments have been made, any individual requests for shifts and off-duties may be considered, by exchanging stints between employees in the same way as we do with the night shift stints in the pre-assignment stage.

We should note at this point that there are potential dangers in attempting to accommodate each and every individual request and preference. First, if each employee is required to specify a set of standing preferences, then some mechanism for quantifying these preferences must be put in place, and the relevant data collected and maintained. This can be a time consuming process. Second, the whole exercise may prove counter-productive. For example, Bard & Purnomo (2005) describe how those employees clever enough to work the system will indulge in gaming, while those who do not understand the process will fail to express their preferences accurately. The system can then be perceived as being unfair, which

6.3 Post-Assignment processing

is contrary to our stated goals. Silvestro & Silvestro (2008) describe how those employees who are most flexible may come to feel that their goodwill is being taken advantage of, leading to a decline in morale. It is for these reasons that both our tour scheduling and assignment models aim to satisfy *generalised* rather than *individual* preferences, and follow general good practice as identified by the ergonomic literature.

We do, however, recognise that in practice, it is common for employees to exchange shifts and tours between themselves, and it is our intention that a post-assignment process should identify occasions where shift exchanges between employees would be mutually beneficial. This development will be carried out with the co-operation of a call centre partner, an initiative which is underway at the time of writing. For the purposes of this thesis, however, we focus on the goals relating to ergonomic principles, generalised preferences, and the equitable distribution of generally undesirable shifts.

Once the assignment of daily shifts to employees is finalised, we assign the lunch and coffee breaks. This approach aligns with the practice whereby employees are given notice of their forthcoming shift assignments at an earlier date than they are notified of breaks. Typically, shift and off-duty assignments are made perhaps 4-6 weeks in advance, whereas breaks are only assigned with perhaps one week's notice. This gives planners the opportunity to reschedule all breaks should there be significant changes to the call profile or to staff availability in the interim.

The rostering goals relating to break assignments are ERG13, which calls for breaks to be scheduled regularly, and ERG14, which states that employees should be given some choice of break. The constraints imposed on the scheduling

of breaks by the tour scheduling model ensure that breaks are initially scheduled at reasonable times. However, we now re-allocate breaks to individuals taking account of the two goals. We first aim to fulfill individual requests for breaks at specific times, in accordance with goal ERG14. In practice, such requests are submitted, for example, by employees wishing to make a lunchtime appointment. Our approach is to prioritise these requests and, where possible, deal them before the remainder of the breaks are assigned. We are assuming that such requests are relatively few, and that it will still be possible to assign the rest of the breaks without violating any of the constraints relating to break timings which were imposed at the tour scheduling stage. Once these individual requests have been fulfilled, we allocate the remaining breaks on the principle of “first on duty, first to break”, as this will have the effect of allocating breaks regularly throughout each shift, in line with rostering goal ERG13.

6.4 Empirical Example

We took the 171 tours which were scheduled in the comparative exercise described in the previous chapter. We re-distributed the night shift stints as described in 6.1. Initially, we set a penalty weight w_0 of 1000 for the hard constraints, and penalty weights for the soft constraints in the model as shown in table 6.5. The values attached to each of the penalties reflect one view of the relative importance of the constraints. The weights can be (and, indeed should be) changed to reflect local preferences. Given that the solution method used for our assignment model (the Hungarian method) is exact, the solution itself is optimal with regard to the assignment goals and their relative weightings. Thus, any assessment of the

quality of the solution is subjective. A process of establishing an appropriate set of weightings for a particular operation in practice is described by Dowsland & Thompson (2000). Here, following an analysis of the results obtained with this initial set of weights, we will illustrate how changing the weights can impact upon the solution.

6.4.1 Initial weight settings

	Soft Constraint	Weight	Value
S_1	Consecutive off-duty days (preferred)	w_1	10
S_2	Isolated work days (to be avoided)	w_2	10
S_3	Forward rotation	w_3	10
S_4	Backward rotation	w_4	100
S_5	Number of consecutive shifts = 7	w_5	100
S_6	Equal number of stints of night shifts to be worked	w_6	10
S_7	Equal number of Saturday shifts to be worked	w_7	10
S_8	Equal number of Sunday shifts to be worked	w_8	10
S_9	Equal number of full weekends to be worked	w_9	100

Table 6.5: Penalty weights for Tour Assignment example

We wish to assess the extent to which our rostering goals are met when using this model. In order to assess the fairness aspects, we need to analyse the results over a period which is long enough for the some regularity to become established. Some types of undesirable tour, such as those with Saturday off and Sunday worked, are relatively scarce, and so we need to assess the assignments over a fairly long period. Using the same set of tours each time, we generate a roster for 14 weeks. Starting with a random assignment in week 1, there are $171 \times 13 = 2223$ weekly assignments to be made. We recorded the number of violations

6.4 Empirical Example

of the hard constraints (with penalty 1000) and soft constraints with value 100. First, we consider the results in terms of week-to-week continuity. These results are listed in table 6.6.

	Constraint	Value	Deviations
H1	A minimum of 11 hours between attendances	1000	0
H2	Night shifts assigned in stints of 2-3 days	1000	0
H3	A minimum of 48hrs off duty after night shifts	1000	0
H4	Number of consecutive shifts > 7	1000	0
S4	Backward rotation	100	56
S5	Number of consecutive shifts $= 7$	100	32

Table 6.6: Violations of high penalty constraints

We can see that no hard constraints were violated, that only 2.5% percent of assignments resulted in a backward rotation, and that only 1.4% resulted in a 7-day stint. Our goal was to distribute the unsociable night stints and weekend shifts as evenly as possible among the workforce. As discussed above, it may well be possible to re-distribute these stints if there are employees who actually prefer to work them. However, our overall aim at the tour assignment stage is to meet our rostering goals as described in Table 6.2, and the treatment of individual requests is held back until the post-processing stage. In Table 6.7 we report the minimum and maximum numbers of the night and weekend stints and shifts worked by any employee, over the 14 week period.

In terms of fairness of distribution, the difference between the lowest and highest numbers of unpopular assignments was only 1 for three of the criteria, which was the fairest possible distribution. Since at least one employee worked only two Saturday only shifts, while others worked four, the evenness of distribution of this

Shift or stint type		Minimum Number Assigned	Maximum Number Assigned	Mean
S6	Nights (2 or 3 day stints)	1	2	1.1
S7	Saturday (with Sunday off)	2	4	3.6
S8	Sunday (with Saturday off)	0	1	0.2
S9	Full weekend worked	4	5	4.2

Table 6.7: Distribution of unsociable shifts

type of tour has room for improvement. This aim can be achieved by increasing the weight attached to this criteria as we now demonstrate.

6.4.2 Alternative weight settings

The aim of this exercise is to improve upon the fairness of distribution of the tours containing a Saturday shift and Sunday off-duty, and so we increase the weight attached to the assignment cost of these tours, w_7 . We began by increasing the value of this weight from 10 to 50, leaving all other weights unchanged. This had no effect on the distribution of the “Saturday” tours, with the range of assignments remaining at 2-4. We then increased w_7 again, this time to 100. This change had the desired effect of increasing the minimum number of assignments, thereby resulting in a more even distribution and a range of 3 to 4 Saturday tours assigned. In this solution, the ranges of all four undesirable tour types is now only 1, which is the best possible. However, the number of backward rotations increased from 56 to 74, and the number of occurrences of 7 consecutive shifts increased from 32 to 41. This illustrates how the various goals trade-off against each other.

We would emphasise that the effect of the weight changes has been neither to objectively improve the quality of the solution, nor to worsen it. Indeed, both the original solution and this alternative are optimal with respect to the weightings of our stated goals. The preference for one solution over the other is subject to the general opinion of the rostered employees, and some trial-and-error may be required to train the weights for each installation.

6.5 Chapter Summary

In this chapter we have presented our model for assigning the previously scheduled tours, shifts, and breaks to the employees. Our methodology has three stages. The pre-assignment stage aims to distribute the night shifts more evenly, so that each tour contains at most one stint of three night shifts. We then assign the weekly tours to the employees in the second (main) stage, using the Hungarian method. The cost function takes account of continuity from the previous rostering period to the current, and the equitable distribution of unsociable weekend and night shifts among the work force. At the third, post-assignment stage, which remains under development, we intend to consider individual requests for shift start and finish times or off-duties, and assign the scheduled lunch and coffee breaks for each day. We will allow for individual requests for specific break times and meet these first where possible. The rostering process is completed by assigning the remaining breaks on the principle of “first on duty, first to break”.

The contribution of this chapter has been to present a comprehensive tour assignment model which takes account of twelve employee-friendly rostering goals. The model differs from others in the literature as it takes account of generalised

6.5 Chapter Summary

preferences for patterns of off-duty, week-to-week continuity, fairness of distribution of unsociable shifts, and individual requests for shifts, days off duty, and lunch breaks. The restructuring of scheduled tours during the pre- and post-assignment stages introduces an element of integration into the tour scheduling and assignment processes.

Chapter 7

Application to Nurse Rostering

In chapter 3 we identified a total of 24 goals for a call centre employee rostering model, and in chapters 5 and 6 we applied them systematically to the tour scheduling and tour assignment stages of the call centre rostering process. We note, however, that the 5 goals relating to government legislation, and the 14 drawn from ergonomic principles and guidelines are applicable not only to call centre rostering, but also to any employee rostering domain. Many of our “employee-friendly” goals form standard constraints in benchmark instances of the nurse rostering problem, and even the generalised preferences are comparable. There are further similarities between the call centre, and nurse rostering problems:

- Both address the need to match the correct number of appropriately skilled employees to a demand for service, at all times.
- The majority of hospital wards, and many call centres, operate 24-hour days, 7 days-per week.

-
- The level of the workload, in both cases, varies throughout the day, from day-to-day, from week-to-week and often seasonally.
 - Like nurses, call centre agents are employed on a mixture of full time and part time contracts, and may be skilled to handle different types of work.
 - Shift restrictions enforced by employment law are the same in both cases, and staff contracts do not differ radically between UK call centre agents and nurses.
 - Staff resources are limited in both hospital wards and call centres.

These similarities in the overall nature of the two problems, together with the similarity in the goals of nurse rostering models and our own call centre rostering model, invite the interchange of methodologies and techniques between the two domains. Our investigations in this area have led to our development of a nurse rostering model, and some aspects of this model have been submitted for publication. In order to make it easy for the reader to locate our published material in the wider context of this thesis, we have included our paper in it's entirety as a separate chapter to follow this one. Taken together, this current chapter and the paper constitute a single contribution. The two chapters are organised as follows. In this chapter, we begin with an analysis of the techniques used in those models which are able to address benchmark problem instances of real-world complexity. In section 7.2 we then present the formulation of our own Integer Programming model, tailored to one specific problem instance referred to as "ORTEC01". In the first part of our paper (included here as chapter 8) we present some insights into two benchmark problem instances, which strongly influenced the design of

our model, and we present our solutions to these two instances. In the second part of the paper, we offer a critique of the methodology adopted by approaches which address an isolated, discontinuous rostering period, and demonstrate how our model can be adapted to give better results in a continuous rostering environment.

7.1 Practical approaches to Nurse Rostering

There is a large body of literature on nurse rostering (see Burke *et al.* (2004b) for an overview). However, as was pointed out by Burke *et al.* (2008), there are relatively few papers which directly address problems of real world complexity. Since our concern is with real world rostering problems, we have focused our attention in two areas. First, we consider heuristic approaches, focussing on a series of papers associated with the commercial rostering package known as Impakt PLANE. The methodology implemented in this system is itself referred to as ANROM. Second, we investigate those papers which take a similar approach to our own call centre models, namely those based on Mathematical Programming. Early research in this area led to the development of the ANSOS nurse rostering system implemented in a number of US hospitals. Our overall aims are to gain insights into the nature of the problem at a practical level, to analyse the techniques currently used by academics in the mainstream of nurse rostering research, and to assess the opportunity for our own contribution.

7.1.1 Heuristic Approaches

A series of papers Burke *et al.* (1999, 2001a,b, 2002, 2003, 2004a), Causemaecker & Vanden Berghe (2003) and Causemaecker *et al.* (2005) produced jointly between Nottingham University and Katholieke Universiteit Lueven (Gent, Belgium) relate to a commercial rostering package known as *Plane*. The Plane product was commercially developed by Impakt ¹ and GET ² primarily for the Belgian market to meet demand for a product which would schedule staff individually, taking account of individual preferences and availability, rather than constructing a cyclical rotation. In doing this, the product takes account of a number of constraints relating to the continuation from one week to the next within the planning period, which is typically one month. Nurses are categorised according to grade (e.g. senior nurse, junior nurse) and a separate staffing requirement is established for each grade in terms of a minimum and a preferred level per shift. There is a degree of flexibility in that some grades can provide cover for others if necessary, but there is an associated cost reflected in the evaluation function. Coverage constraints must always be satisfied, and all attendance considerations are treated as soft constraints with cost parameters under user control. These attendance constraints are divided into a number of categories relating to hours, days, weekends, shifts. Personal requests have higher priority than other constraints when resolving conflicts.

The Plane product is introduced in Burke *et al.* (1999). The large number of modifiable constraints, the number of possible daily duties (6 to 15) and the need for fast execution times led the authors to develop a heuristic method based

¹Impakt N.V., Ham 64, B-9000 Gent

²GET, General Engineering and Technologie, Antwerpse Steenweg 107, B-2390 Oostmalle

7.1 Practical approaches to Nurse Rostering

on tabu search and manual adjustment. An initial feasible solution is generated using one of three methods; using either the previous or current planning period as a start point, or by generating a new solution randomly. The tabu search algorithm seeks to exchange duties and off-duty days between staff members on a daily basis. Following a successful exchange the areas involved are placed on the tabu list. This method outperformed a steepest-descent algorithm, but further improvements in schedule quality could be identified by inspection of the results. Two additional heuristics, referred to as *diversifications* were added to improve weekend working, and to improve the schedule for the individual with the worst result. A final step was implemented to shuffle duties between individuals, but as this step resulted in unacceptably long execution times the intention was to implement it as a system option.

Robustness issues arising from the use of tabu search heuristics in practice are described in Burke *et al.* (2001b), which introduces a memetic approach based on the implementation of a steepest-descent heuristic within the framework of a genetic algorithm. Previously, a problem had been encountered which related to the scheduling of each grade separately; the approach introduced here is to schedule all grades concurrently, with resources borrowed from adjacent grades to cover for absences and unavailability. The memetic algorithm begins by randomly generating a number of feasible schedules. Each is then improved using the tabu search algorithm. The resulting schedules are used to create offspring, by pairwise comparison of schedule “rows” (i.e. the individual schedules created for each staff member) from parent solutions, with tournament selection. The new solutions are then repaired to make them feasible in terms of the staffing constraints, and the process repeats. Variations on this basic approach included

7.1 Practical approaches to Nurse Rostering

randomly varying the sequence in which grades are scheduled, randomly selecting the child schedule from the two parents, and constructing child schedules from “best-placed” events within the two parents (i.e. the two parents are combined, rather than one selected in its entirety). In the latter case, feasibility is restored by the random addition of additional shifts, which also serves to increase diversity. Experiments were conducted with different numbers of event-swaps and one algorithm (4-swaps) found to perform best. The greedy-shuffling step from the earlier tabu search hybrid was re-introduced, both as a local improvement meta-heuristic and as a post-processing step on the optimum solution, and resulted in further improvements. This new hybrid algorithm gave better or similar results on a number of real-world problems, compared to the earlier hybrid tabu search algorithm. In addition to the improved result, the ability of the new algorithm to schedule all staff grades concurrently was seen as a major benefit; previously, user judgment had been needed to identify the best sequence, and attempts to automate these decisions within the tabu search algorithm had been unsuccessful. The downside of the new algorithm was its execution time; the solution to the most complex problem took over two hours to compute. Although solutions would probably be found more quickly on a modern computer, the slow execution times observed in 2001 motivated the authors to move on to the next step.

In Burke *et al.* (2001a), the authors point out that the application of evolutionary algorithms requires a large number of evaluations of the fitness function. They introduce an efficient evaluation method based on the concepts of *Time Units* and *Numberings*. A time unit is a unique combination of a shift and day within the planning period so that any attendance can be identified in binary terms. A numbering is a numeric representation of each constraint which can

7.1 Practical approaches to Nurse Rostering

be compared to the values of the time units. This representation allows all constraints to be checked in a single pass, rather than requiring each to be checked separately. This allows for fast evaluation, and is memory-efficient in that multiple constraints can be handled by the same numberings.

The aim of Burke *et al.* (2002) was to refine the optimisation algorithm in order to overcome practical problems. In earlier versions of Plane, schedule quality was measured by the number of constraints violated. The new approach is to use *compromise programming* to measure the distance from an ideal solution. Weights are attached to the criteria to reflect their relative importance in dimensionless units, and a best and worst case for each criteria given values of 0 and 1 (times the weight). The idea is that this improves on the success-failure approach by taking account of the range of possible violations. The search is driven by the tabu search algorithm with meta-heuristics to address weekend working and to improve the schedule for the individual with the worst result. Experimentation with various weightings gave encouraging results, and further work will offer guidance on the setting of weights.

The number and variety of scheduling constraints, together with the range of criteria weightings, created difficulties in that the search space has deep, narrow valleys which were hard to find, with solutions often becoming trapped in local optima. This problem is discussed in Burke *et al.* (2003), which proposes to expose hidden parts of the solution space by introducing problem-specific neighbourhoods. The basic tabu-search heuristic looks to move a scheduled shift from one staff member to another staff member who was previously off duty on that day. The neighbourhood for this heuristic is the set of schedules, for the selected day, for all staff members of appropriate grade. Additional meta-heuristics were

7.1 Practical approaches to Nurse Rostering

previously introduced to improve weekend working, to improve the worst-case schedule and to “shuffle” shifts between staff members. Further meta-heuristics are introduced in this paper, such as the balancing of over- and under-time. Neighbourhoods are defined for each of the meta-heuristics, and each neighbourhood is paired with either the steepest-descent or tabu search algorithm. The method generates an initial feasible solution randomly, and then moves from one algorithm-neighbourhood to another and searches for improvements in each. The selection of the algorithm-neighbourhood to move to is determined by a variable which keeps track of the success rate of previous visits. Results indicate that changing neighbourhoods enables better solutions to be found, which remained hidden when searching a single neighbourhood.

In Causemaecker & Vanden Berghe (2003) the writers accept that in practice, it is sometimes preferable to relax the hard staffing constraints if this leads to a marked improvement in the evaluation of the soft schedule constraints. A pre-processing check begins by checking the number of staff available, at each grade, for each day to be scheduled. First, the maximum number of staff available must exceed or equal the number specified, and a check is made to ensure that a feasible solution exists. If this validation is passed, the algorithm then proceeds to look at the number of staff who are expecting to be either on, or off duty according to recent work history, or special requests such as holidays. The algorithm reports on any inconsistencies between requested staffing levels and this new calculation of staff availability, and recommends increases or decreases in the requirement as appropriate. The user then has the choice of relaxing the coverage constraints or accepting that some violations of important schedule constraints will be unavoidable. Once the schedule has been generated, a post-processing algorithm

7.1 Practical approaches to Nurse Rostering

can suggest either additional shifts which will improve the overall coverage at the expense of soft constraints, or additional work-hours which will improve the schedule for staff members with less than the expected weekly total hours worked at the expense of overstaffing.

The whole series of papers is condensed into a single summary paper Burke *et al.* (2004a), where the methodology is now referred to as ANROM (Advanced Nurse Rostering Model). The authors claim that the multi-criteria approach incorporates an extensive set of realistic constraints into a flexible model, offering a major improvement for hospital planners in comparison with methods that employ a single evaluation function. ANROM is commercially available and widely used. The main drawbacks pointed out by the writers are that the model requires insight into the characteristics of data (presumably in the setting of the cost weightings and the “distance measures” for each constraint) and that pre-scheduling feasibility checks need to be performed on the hard staffing requirements, and adjustments made if necessary. The writers accept that run times can be lengthy, but point out that these can be tailored to the time available by selecting the heuristics to be used.

Comments

The series of papers described here were produced over a period spanning 4-5 years. It is clear that by the end of this period the initial idea of treating staffing levels as hard constraints, and schedule considerations as soft constraints, has been somewhat eroded. From the outset there was an initial check to ensure that at least one feasible solution existed, with users required to change the staffing

7.1 Practical approaches to Nurse Rostering

levels if necessary. Additional checks added in Causemaecker *et al.* (2005) enabled users to manage trade-offs between violations of staffing constraints and some of the more important schedule constraints such as holiday requests, with the result that the staffing constraints are less hard than was originally intended. On the other hand, some soft constraints are clearly much more important than others, and a wide range of cost weightings created a solution space contoured such that good solutions were often hidden (Burke *et al.* (2003)). It would be interesting to consider whether some of the more important schedule criteria could be treated as hard constraints and handled separately, thus potentially creating a more even search space for the remaining soft constraints. This would enable good solutions to be found more easily, and also make it easier for users to specify the necessary weightings since the remaining criteria will be relatively closer in importance.

The overall approach incorporates a number of heuristics. An initial consistency check calculates whether a feasible solution exists, and offers options for constraint relaxation if none is found. Multiple neighbourhoods relating to different criteria are searched using either tabu-search or steepest-descent heuristics, and a genetic algorithm introduces variety in order to escape local optima. Meta-heuristics are used to find improvements in a number of specific areas, and a shuffling algorithm mimics a manual process. A post-scheduling algorithm offers additional shifts which will either improve the coverage at the expense of schedule quality, or improve the schedule quality at the expense of overstaffing. This solution has evolved from a start point based on tabu search, with the introduction of additional heuristics at each stage designed to overcome problems identified at earlier stages. It is unclear now which of these heuristics are the primary drivers of solution quality, and it may be that there is some redundancy in the overall

7.1 Practical approaches to Nurse Rostering

process which could give rise to streamlining opportunities, leading in turn to faster execution times.

The main points arising are that a) the nature of the problem, characterised by a large number of constraints of varying degrees of importance, makes good solutions hard to find with heuristic approaches and b) in practical applications, it is probable that some degree of decomposition will be required even with heuristic methods. Our perception is that there is an opportunity to apply, or develop our Integer Programming models to address the nurse rostering problem.

7.1.2 Mathematical Programming Approaches

Most of the MP approaches to nurse rostering reporting in the literature are based on set-covering formulations using weekly or fortnightly “schedules” or tours, which are sequences of shifts and off-duty days covering the whole rostering period. There are two main approaches. One is to select a tour for each nurse from a working set of pre-constructed candidate tours. The other is to schedule weekly tours and then assign the tours to nurses as a separate step, in a similar way to our own call centre rostering model described in the previous two chapters. An alternative approach, reported in just one article, is to work with “stints” of shifts and off-duties. We begin by outlining these approaches, and then describe in detail the papers and models which apply them.

1. The first idea is to assign weekly (or fortnightly) tours direct to individual nurses. The benefit of this approach is that period-to-period continuity, fairness issues, and personal requests can be taken into account at the same time as the coverage constraints. The difficulty is that the number of possi-

7.1 Practical approaches to Nurse Rostering

ble tours can become impracticably large, particularly if there are multiple shift types, or if the number of shifts to be worked by each nurse is variable. Early methods of addressing this difficulty were:

- To restrict the number of tours under consideration for each nurse, based on their individual preferences (Warner (1976))
- To impose a structure on the solution based on policies such as a regular rotation of weekend shifts (Arthur & Ravindran (1981); Miller *et al.* (1976); Warner (1976))
- To assign specific shifts subsequent to the allocation of working days (Arthur & Ravindran (1981)).

More recent papers (Jaumard *et al.* (1998), Bard & Purnomo (2005)) use column generation to restrict the number of tours under consideration at any time, or goal programming to prioritise the scheduling objectives (Topaloglu & Ozkaran (2004)).

2. The second main approach is to schedule tours of one or two weeks, and then subsequently assign the tours to nurses as a separate step. This reduces the complexity of the problem since it allows the application of implicit modeling techniques (Ozkaran & Bailey (1988), Isken (2004)). However, it becomes difficult to comply with the full range of requirements in terms of period-to-period continuity, fairness and individual requests since the tour assignment is restricted to those tours already scheduled. Although this is the mainstream approach to call centre rostering, where such considerations have been viewed as being relatively unimportant, it is less practical for

7.1 Practical approaches to Nurse Rostering

nurse rostering.

3. Millar & Kiragu (1998) differ by basing their approach on stint scheduling. This avoids the need to explicitly define tours or even off-duty patterns. Rather, these are implicit in the sequences of working and off-duty stints. In modern benchmark nurse rostering instances such as those we will address in the following chapter, the rostering period covers several weeks, and the number of shifts to be worked by each nurse is not normally fixed each week. In such cases, the number of possible tours (over the whole rostering period) can become very large and so the stint based approach strikes us as offering an interesting alternative to tour scheduling.

We now discuss each of the papers in the three categories outlined above.

Integrated models

Warner (1976) describes a nurse scheduling system used in a number of units of a US hospital. Individual preferences are quantified in terms of off-duty patterns and shift sequences, and the objective is to minimise the overall preference cost, subject to daily shift coverage constraints. 14-day tours are constructed and the cost of assigning each tour to each nurse is calculated, and adjusted to allow for period-to-period continuity. Some hard constraints are imposed on the tour structures, for example backward rotations are not allowed, and full-time nurses work exactly five days per week. These constraints limit the number of the feasible tours, and typically a maximum of the 50 “best” tours are considered for each nurse. Cyclical “policies” are imposed on the allocation of weekend shifts, and the rotation between shifts. This manual step provides a basic skeleton around

7.1 Practical approaches to Nurse Rostering

which the remainder of the roster is constructed. The solution method is a two-stage process, whereby an initial search for a feasible solution is followed by the execution of an improvement algorithm. One problem with this model is that the skeleton schedule is inflexible, and is constructed manually. A second issue is that the number of tours for consideration is limited, and although it is claimed that this restriction did not markedly affect the solution quality, that may not continue to be the case if the scheduling rules are more flexible, for example if the number of shifts to be worked each week is allowed to become variable.

Miller *et al.* (1976) presented a rostering model for nurses belonging to 3 grades, for a single shift, with a two-week rostering period. The objective function trades off individual preferences for patterns of on and off-duty days, against daily coverage requirements. Hard constraints are used to restrict the number of tours in the feasible set, and soft constraints relating to tour and shift patterns are imposed through the application of weighted preference costs. These costs are inherent in the tour patterns, and are scaled up by individual employees "aversion factors", and again by the individual's rostering history (in order to maintain "fairness") to arrive at a cost of assigning each tour to each employee. The solution method is construction and improvement, where an initial feasible solution is iteratively improved using a co-ordinate descent algorithm. The solutions obtained are near optimal, and the practicality of the model demonstrated by implementation in a US hospital. However, the problem addressed is very simple. In particular, the consideration of only one shift, together with rigid schedule rules such as a fixed number of working days per period, results in a relatively small set of feasible tours. With each nurse working 10 days in a 14 day period, there are 1001 possible tours, but as the authors point out, this number

7.1 Practical approaches to Nurse Rostering

will reduce considerably in the light of feasibility constraints such as minimum and maximum length working stints, and special requests.

Arthur & Ravindran (1981) present a two-stage model in which the first stage assigns days-off patterns to nurses, and the second allocates a total of 5 shifts to each nurse, each week. Each nurse takes every other weekend off duty, and this policy is imposed in the problem structure. The first stage of the model is implemented using a goal programming approach. The model objectives are grouped into four sets of goals, namely, minimum staffing requirements, preferred staffing requirements, nurse preferences, special requests. The formulation is based on a single set of decision variables, x_{ij} , defined to be 1 if nurse i works on day j . Constraints are generated for each category of goal, using deviational variables to measure the variation from the ideal, and the objective is to minimise the sum of the deviations over all four goals. A priority order is assigned to the goals, and a zero-one goal-programming algorithm used to solve this first stage problem. At the second stage, specific shifts are assigned to nurses, ensuring a minimum of 16 hours off-duty between successive shifts. The method of equalising the fairness of distribution of shift types is to rotate the order of nurses for "tie-breaking" when multiple assignment options exist. As with earlier models, the problem structure is rather simple (a one week schedule with a fixed number of shifts per week and a rigid rotation of weekends worked), and specific constraints relating to night shift working are considered.

Warner *et al.* (1991) present an entire, practical nurse rostering system known as ANSOS (Automated Nurse Scheduling Office System), which was installed in 750 US hospitals at the time of publication (1990). The system includes a number of modules including a staffing module which computes period-by-

7.1 Practical approaches to Nurse Rostering

period nurse requirements based on “acuity” data. The scheduling module is developed from that of Warner (1976), extended to take account of a wider range of nurse preferences including days-off patterns and shift sequences. Possibly due to commercial sensitivities, a detailed formulation of the scheduling model is not presented in the paper, which instead focuses on a more general description of the nurse rostering problem and the ANSOS system as a whole.

Jaumard *et al.* (1998) present a flexible Linear Programming model which applies branch and bound together with column generation. The model generalises previous mathematical modeling approaches in the literature. The rostering period is two-weeks, with staffing requirements defined by staffing periods, with each shift covering one or more period. The workforce comprises a mixture of Full and Part-Time nurses. Uncovered demand may be covered from a centralised pool of reserve nurses. The master problem is to find a combination of schedules that satisfy the coverage constraints, in the same way as, for example, the early models of Miller *et al.* (1976) and Warner (1976). However, additional requirements are added which relate to care quality, quantified by the balance of experience during each period, and to individual preferences. Due to the problem size, relatively few schedules are considered by the master problem at any time. The auxiliary problem is to find schedules for each nurse, which are feasible in terms of the contractual and collective agreements relating to that nurse.

A similar approach is taken by Bard & Purnomo (2005). Starting with a “base” schedule for each nurse, additional columns are generated via a heuristic, which exchanges one or two working or off-duty days within a nurse’s schedule. Following such an exchange, the new schedule must be checked for feasibility in terms of the hard constraints relating to that nurse (for example, one nurse

7.1 Practical approaches to Nurse Rostering

may require every other weekend off). An integer program is then solved, using CPLEX, to minimise a multi-objective function based on individual preferences and coverage requirements, with agency nurses used to cover any shortages. Additional columns are generated based on the periods covered by the agency nurses, and the process repeats. The accuracy of the solution is highly dependent on the efficacy of the column generation heuristic, and much of the processing effort can be dedicated to the repair process, which ensures column feasibility.

Topaloglu & Ozkarahan (2004) apply a goal programming approach to assign shifts direct to nurses. A total of 11 goals relating to coverage and nurse preferences are identified. The goals are prioritised according to local preferences, and the model is first solved for the primary goal. The solution details are then added as additional constraints, and the model solved again for second priority problem, and so on. A set-covering formulation is used to generate daily shifts, including fixed breaks, which cover the interval staffing requirements. The coverage goal is to assign these shifts to the nurses, with deviational variables used to measure staffing errors beyond an acceptable range. Individual nurse preference goals relate to the number of assigned work hours, shift types, consecutive working days, and weekend working. Hard constraints ensure that each nurse receives a break during each shift, and that a minimum rest time is imposed between consecutive shifts. No other constraints are imposed on shift sequences (such as forward rotation), and no specific rules are applied to night shift working. It is not clear how an equitable distribution of the unpopular shifts is to be managed over a long period. For example, weekends are assigned based on individual requests, and it is unclear how conflict can be resolved if there are more requests than weekends off. Handling the multiple goals entirely separately could mean

7.1 Practical approaches to Nurse Rostering

that good compromise solutions may be not be found, and it would interesting to assess the extent to which the number of goals can be reduced by grouping them together so that some are solved simultaneously. The problem instance addressed is a fairly small one, with only 11 nurses.

Tour Scheduling and Assignment

Ozkarahan & Bailey (1988) develop the general tour-scheduling model of Bailey (1985) for application to the nurse rostering problem. Essentially, the model has two components. A time-of-day model ensures that the set of shifts scheduled for each day covers the staffing requirements for each period of that day (although understaffing is allowed, and is modeled as a goal using deviational variables). The second part of the model is the day-of-week component, which schedules the days-off patterns to be worked over a week. Each of the two components is modeled using a set-covering formulation. The overall complexity of the problem led the authors to treat each stage separately and develop an algorithm for assigning the daily shifts to the weekly patterns. It is suggested that the assignment of these “optimum” tours to individual nurses can be handled either using an assignment algorithm, or by allowing nurses to choose the tours themselves. The mathematical formulation of the standard two-dimensional assignment problem, and its application to the assignment of weekly tours to individual nurses, is described in Ozkarahan (1991).

Isken (2004) is another to take the approach of decomposing the rostering process into separate tour scheduling and assignment stages. A tour type is defined as a combination of a days-on pattern and a shift length. Daily shifts are implicitly matched with tour types by a set of control constraints. The variation

7.1 Practical approaches to Nurse Rostering

in shift start times that can appear in a single tour is limited by a start-time bandwidth, which is imposed using “upper and lower bound sweep constraints”, which operate on a similar basis to the forward and backward pass constraints which Bechtold & Jacobs (1990) used to restrict break-time variations within a shift. The author claims that this model complements Brusco & Jacobs (2000) by allowing a mixture of full and part-time workers, and there are clear similarities between the two models. It is recognised that week-to-week assignment may be problematic, that certain tour types may be more difficult to accommodate than others. It is suggested that the assignment process can be made easier by disallowing certain tour types, such as those including both Saturday and Sunday (the last two days of the week). However, since the staffing requirements for these two days must be covered one way or another, we would point out that the recommended course of action will have the effect of reducing the number of weekends off, which are generally considered to be highly desirable.

Our own view is that weekly tour scheduling and assignment in a nurse rostering context is problematic in that week-to-week continuity is difficult to manage. For example, a nurse’s aversion to a particular tour pattern will change in the light of the pattern that was worked in the previous week. Moreover, many of the constraints present in modern nurse rostering instances apply over periods longer than one week.

Stint Scheduling

Millar & Kiragu (1998) introduce the concept of a stint, namely a sequence of working days or off-duty days. Working stints include the specification of the type of shift to be worked on each day. The problem addressed is based on 12-

hr shifts, with only two shift types, namely day and night. It is assumed, that with such long shifts, working stints will be restricted to a maximum of 4 days. A network is constructed where each node corresponds to a working or off-duty stint, and transitional costs between stints are calculated and represented as arc costs. The intrinsic cost of each stint is represented as a node cost. The problem is then to find shortest-paths through the stint network, and hence minimise the total tour costs. A problem instance from a single hospital is solved using CPLEX. Although some additional side constraints are also specified (such as the coverage constraints, the total number of shifts per nurse, and the number of weekends), the problem addressed remains relatively simple, particularly in terms of the number of shift types and the length of stints.

7.2 The Nurse Rostering Model

Our initial intention in addressing the nurse rostering problem was to apply the same two-stage decomposition (namely, tour scheduling and tour assignment) as we successfully used to solve the call centre rostering problem. This approach turned out to be inapplicable in the nurse rostering context, because although we identified a number of similarities between the call centre and nurse rostering problems (described in the introduction to this chapter), there are also some important differences.

1. The number of shifts or hours that nurses work in a week can vary, whereas call centre agents typically work the same number of shifts and hours per week.

2. The continuation from one rostering period to the next is easy to manage in call centres because the number of shifts required on the last day of the week (Sunday) are very few in comparison with the number required on the first day of the next week (Monday). As we described in detail in chapter 6, feasible, low-cost, week-to-week continuations are normally available when weekly rosters are produced independently of each other. This is not the case in nurse rostering, where the Sunday staffing requirement is similar to Monday, and the continuity between rostering periods can be difficult to manage. We address this methodological issue in detail in the second part of our paper, which forms chapter 8 of this thesis.
3. Call centres employ dozens or sometimes hundreds of employees. The number and variety of tours produced allows a great deal of choice when assigning tours to employees. This makes it relatively easy to comply with individual requests for specific shifts or off-duty days. Again, this is not the case in nurse rostering, where typically only a handful of nurses are rostered together. So, although nurse rostering problem instances are typically smaller than those of call centres, the effect of the individual requirements in nurse rostering is to create a search-space where good solutions are hard to find. By contrast, the flexibility in the call centre tour scheduling problem typically leads to the existence of multiple solutions close to the optimum.

These factors led us to reconsider the structure of the model. Ideally, the whole problem would be solved in a single step, but the complexity of the problem makes an integrated approach computationally impractical. Our overall approach is to decompose the problem into two stages, both of which are solved using Integer

Programming models.

At the first stage, we assign stints of shifts to individual employees, according to constraints relating to the number of shifts to be worked, the patterns of off-duty days, and weekend working. Night shifts are superimposed onto these working stints in accordance with the constraints relating to night shift working. Coverage constraints for each day are also taken into account at this stage, as are any individual requests relating to days on or off-duty. We term this the “attendance rostering model”.

At the second stage, we take the partially completed roster and assign each working day (which is not already designated as a Night shift) to be either a Day, Early, or Late shift. We take account of constraints relating to the numbers of each type of shift required on each day of the roster, and to shift sequences. We address any requests for particular shift types on specific days. We term this the “shift allocation model”.

For ease of presentation, we have tailored the model formulations for the benchmark problem “ORTECO1”, which is described in detail in chapter 8. The model is applicable to other problem instances by changing some parameters or adding more constraints as required. The key to the constraints and penalties in our model reflect, for illustrative purposes, those of the problem instance “ORTECO1” as given in Tables 8.1, 8.2, 8.3 in section 8.2. In order to facilitate understanding of the model, we advise the reader to read section 8.2 (“benchmark problems”) of the next chapter before returning to this point.

7.2.1 Stage 1 - Attendance Rostering

The aim of this model is to find a set of assignments which cover the daily shift requirements, while minimising the penalties for violations of constraints relating to the numbers and patterns of shifts that can be worked by individual employees. Our approach is to impose those constraints which are associated with a penalty of 90 or more, as hard constraints.

We define a *stint* as a sequence of consecutive working days. Stints of length 2 to 6 are configured in our input data, such that no stints end on a Saturday or begin on a Sunday, since this would incur penalty S2. We allow a single stint of length 1, which covers the final day of the roster, in accordance with C1 and C5. Generally, stints of length 1 attract penalty S3 and are therefore disallowed. A single rest day is identified by the number of off-duties between successive stints, and hence a single off-duty day on the last day of the roster is not penalised, in accordance with C3. Night shifts are assigned in stints of 2 or 3, in accordance with H3 and S1. Note that we do not take advantage of the flexibility of an unpenalised single night shift at the end of the month (as allowed by C4) within the formulation as stated below, as the resulting optimal solutions for the benchmark instances reported in chapter 8 does not improve when a single night shift is allowed, and so an alternative formulation was not required.

Note that the working week is defined to run from Monday to Sunday. The first and last weeks of the roster may therefore contain fewer than 7 days. In ORTEC01, the one month rostering period begins on a Wednesday, so the first week in this instance only contains Wednesday to Sunday, and the last week contains only Monday to Friday.

Indices

i nurse, $i \in \mathcal{I}$

j stint, $j \in \mathcal{J}$

k day of roster, $k \in \mathcal{K}$

l week of roster, $l = 1, \dots, 5$

Data

a_{jk} =1 if stint j covers day k , 0 otherwise

b_{jk} =1 if stint j ends on day $k - 1$ and is therefore followed by an off-duty
on day k , 0 otherwise

c_{jk} =1 if day k is the first day of stint j , 0 otherwise

r_k the total number of shifts required on day k

Binary Decision Variables

x_{ij} =1 if nurse i works stint j , 0 otherwise

n_{ik} =1 if nurse i begins a stint of night shifts on day k

d_{il}^{+s} =1 if nurse i works s shifts more than the ideal in week l

d_{il}^{-s} =1 if nurse i works s shifts fewer than the ideal in week l

Sets

$\mathcal{J}$	the set of all nurses
$\mathcal{J}^F$	the set of Full-Time nurses
$\mathcal{J}^P$	the set of Part-Time nurses
$\mathcal{J}$	the set of all possible stints
$\mathcal{J}^s$	the set of stints of length s , where $s = 1, \dots, 6$
$\mathcal{J}^W$	the set of stints that include all or part of any weekend of the roster
$\mathcal{K}$	the set of days in the roster
$\mathcal{K}_l$	the set of days belonging to week l of the roster
$\mathcal{K}^{N2}$	the set of days on which a stint of 2 night shifts can begin
$\mathcal{K}^{N3}$	the set of days on which a stint of 3 night shifts can begin

Note that in ORTEC01, exactly one night shift is required for each day. In the following chapter, we demonstrate that in order to avoid high penalties, the night shifts must be scheduled in stints of 2 or 3 shifts in a specific pattern of Monday and Tuesday, (2 shifts), Wednesday and Thursday (2 shifts), then Friday, Saturday and Sunday (3 shifts). Hence, the stints of 2 night shifts can only start on Mondays and Wednesdays, and the stints of three nights begin on Fridays. These are the days which form sets $\mathcal{K}^{N2}$ and $\mathcal{K}^{N3}$.

Note also that the assignment to a full-time nurse of a stint of length 2 or 3 days is exempt (by C2) from penalty S7 if it ends on last day of the roster. These two specific stints are therefore excluded from $\mathcal{J}^2$ and $\mathcal{J}^3$ respectively.

7.2 The Nurse Rostering Model

Integer Programming Formulation for the Attendance Rostering Model

Minimize

$$\begin{aligned}
 & 40 \sum_{i \in \mathcal{I}^F} \sum_{j \in \mathcal{J}^2} x_{ij} + 10 \sum_{i \in \mathcal{I}^F} \sum_{j \in \mathcal{J}^3} x_{ij} + 40 \sum_{i \in \mathcal{I}^P} \sum_{j \in \mathcal{J}^5} x_{ij} + 10 \sum_{i \in \mathcal{I}^P} \sum_{j \in \mathcal{J}^4} x_{ij} + \\
 & 40 \sum_{i \in \mathcal{J}} \sum_{l=1}^4 (d_{il}^{-2} + d_{il}^{+2}) + 10 \sum_{i \in \mathcal{J}} \sum_{l=1}^4 (d_{il}^{-1} + d_{il}^{+1})
 \end{aligned}$$

subject to

$$\sum_{i \in \mathcal{J}} \sum_{j \in \mathcal{J}} x_{ij} \cdot a_{jk} = r_k \quad \forall k \quad (7.1)$$

$$\sum_{j \in \mathcal{J}} x_{ij} \cdot (a_{jk} + b_{jk}) \leq 1 \quad \forall i, k \quad (7.2)$$

$$\sum_{j \in \mathcal{J}} x_{ij} \cdot b_{jk} + \sum_{j \in \mathcal{J}} x_{ij} \cdot c_{jk+1} \leq 1 \quad \forall i, k \quad (7.3)$$

$$\sum_{j \in \mathcal{J}^W} x_{ij} \leq 3 \quad \forall i \quad (7.4)$$

$$\sum_{i \in \mathcal{J}} n_{ik} = 1 \quad \text{for } k \in \mathcal{K}^{N2} \cup \mathcal{K}^{N3} \quad (7.5)$$

$$\sum_{k \in \mathcal{K}^{N2} \cup \mathcal{K}^{N3}} n_{ik} \leq 1 \quad \forall i \quad (7.6)$$

$$\sum_{j \in \mathcal{J}} x_{ij} \cdot b_{jk+2} \geq n_{ik} \quad \text{for } k \in \mathcal{K}^{N2}, \forall i \quad (7.7)$$

$$\sum_{j \in \mathcal{J}} x_{ij} \cdot b_{jk+3} \geq n_{ik} \quad \text{for } k \in \mathcal{K}^{N3}, \forall i \quad (7.8)$$

7.2 The Nurse Rostering Model

$$\sum_{j \in \mathcal{J}} \sum_{k \in \mathcal{K}_l} x_{ij} \cdot a_{jk} - d_{il}^{+1} \leq 5 \quad \text{for } l=1, \dots, 5 \ i \in \mathcal{J}^F \quad (7.9)$$

$$\sum_{j \in \mathcal{J}} \sum_{k \in \mathcal{K}_l} x_{ij} \cdot a_{jk} + d_{il}^{-1} + 2d_{il}^{-2} \geq 4 \quad \text{for } l=1, \dots, 5 \ i \in \mathcal{J}^F \quad (7.10)$$

$$d_{il}^{-1} + d_{il}^{-2} \leq 1 \quad \text{for } l=1, \dots, 5 \ i \in \mathcal{J}^F \quad (7.11)$$

$$\sum_{j \in \mathcal{J}} \sum_{k \in \mathcal{K}_l} x_{ij} \cdot a_{jk} - d_{il}^{+1} - 2d_{il}^{+2} \leq 3 \quad \text{for } l=1, \dots, 5 \ i \in \mathcal{J}^P \quad (7.12)$$

$$\sum_{j \in \mathcal{J}} \sum_{k \in \mathcal{K}_l} x_{ij} \cdot a_{jk} + d_{il}^{-1} + 2d_{il}^{-2} \geq 2 \quad \text{for } l=1, \dots, 5 \ i \in \mathcal{J}^P \quad (7.13)$$

$$\sum_{j \in \mathcal{J}} \sum_{k \in \mathcal{K}} x_{ij} \cdot a_{jk} \leq 20 \quad i \in \mathcal{J}^F \quad (7.14)$$

$$\sum_{j \in \mathcal{J}} \sum_{k \in \mathcal{K}} x_{ij} \cdot a_{jk} \leq 11 \quad i \in \mathcal{J}^P \quad (7.15)$$

$$d_{il}^{+1}, d_{il}^{-1}, d_{il}^{+2}, d_{il}^{-2}, n_{ik}, x_{ij} \in \{0, 1\}$$

In this model, we only treat those constraints with penalties of less than 90 as being soft (S5-S8, with $n \leq 2$, S9-S10) with all other considerations (S1-S4) treated as hard constraints as mentioned above. There are no defined stints which are longer than 6 days, and the only single day stint is on the final day of the schedule and attracts no penalty. The objective function is made up of a number of terms which impose the penalties incurred for violations of the soft constraints. The penalty weights attached to the terms of the objective function are as specified for the problem instance ORTEC01, and the full set of these weights are given in Table 8.1 in section 8.2 of the following chapter. The exponentially increasing weights are typical of nurse rostering formulations and, generally speaking, are set so as to establish a priority order to different groups of rostering criteria. Such a hierarchy of criteria facilitates the distinction between,

7.2 The Nurse Rostering Model

for example, legal requirements which should always be complied with, and those preferences which should be met if possible, but make little difference to overall roster quality.

The first four terms of the objective function reflect the soft constraints relating to the length of stints to be worked (S7-S8). The first two of these relate to full-time nurses, who are expected to work stints of length 4-6 days (S7). The first term imposes a penalty of 40 for assigning a 2-day stint to a full-time nurse, and the second imposes a penalty of 10 for assigning a 3-day stint. The third and fourth terms are similar in nature to the first two, and relate to the ideal stint length of 2-3 days for part-time nurses (S8). Penalties of 40 and 10 are applied by the third and fourth terms for assigning stints of length 5 and 4 respectively. The last two terms relate to the number of days worked by each nurse in each week. Full-time nurses ideally work 4-5 days each week (S5), and part-timers 2-3 days (S6). The fifth term imposes a penalty of 40 whenever nurse i works two days fewer or more than the ideal in week l , with the number of such occurrences denoted by d_{il}^{-2} and d_{il}^{+2} respectively. The final term imposes a similar penalty when a nurse is assigned one shift too few or too many, denoted by d_{il}^{-1} and d_{il}^{+1} . Note that these penalties are only applied to weeks 1-4, and not to week 5, in line with C2.

Constraint (7.1) ensures that the total number of shifts of all the nurses provided on each day k , given by $\sum_{i \in J} \sum_{j \in \mathcal{J}} x_{ij} \cdot a_{jk}$, exactly equals the requirement for that day, since no over or under staffing is allowed in the solution for ORTEC01. Constraint (7.2) ensures that each nurse is scheduled to work a maximum of one shift per day and that a nurse does not begin a new stint on the day after finishing the previous stint (otherwise the two stints would form one single stint).

7.2 The Nurse Rostering Model

The data item b_{jk} marks the day, k , which follows the completion of stint j and which must be taken off duty. Constraint (7.3) avoids incurring the penalty for a single off-duty day (S4). A minimum of two off-duty days between successive stints is imposed by ensuring that the first day of a new stint, c_{ij} , cannot immediately follow the off-duty which followed the previous stint, b_{jk} . Constraint (7.4) guarantees that each nurse will work a maximum of three weekends during the rostering period (H11). This is imposed by ensuring that each nurse works no more than three of the stints which include either a Saturday or a Sunday in the roster.

Constraints (7.5) to (7.8) relate to night shifts. In order to avoid high penalties, we only allow night shift stints to be scheduled in a particular pattern of coverage, as described in section 8.2. This pattern is imposed by constraint (7.5), which ensures that night shift stints of two or three days (H3) are scheduled to begin on the appropriate days, thus providing one night shift on each day of the roster. Each nurse is allowed to work a maximum of 3 night shifts in a rolling 5 week period (H10). This condition is imposed by constraint (7.6), which limits each nurse to one night stint in the rostering period. Night shift stints are allowed to form part of a longer stint, but must be followed by two complete days off-duty (H5). This rule is imposed by constraints (7.7) and (7.8), which ensure that night shift stints must begin two or three days respectively before the end of the working stint.

Constraints (7.9) to (7.13) relate to the number of days to be worked in a week of the roster. The full-time nurses ideally work 4 or 5 shifts in each week (S5). For each full time nurse, the number of assigned shifts in week l is given by $\sum_{j \in \mathcal{J}} \sum_{k \in \mathcal{K}_l} x_{ij} \cdot a_{jk}$. Constraint (7.9) reflects the ideal upper limit of 5 shifts

7.2 The Nurse Rostering Model

in a week for a full-time nurse, and defines the binary deviational variable d_{il}^{+1} to be 1 if nurse i works 6 shifts in week l , which is the maximum possible as there are no stints of longer than six days (H2). Constraint (7.10) defines deviational variables d_{il}^{-1} and d_{il}^{-2} to take values of 1 to reflect a shortfall of 1 or 2 shifts respectively. The values of all deviational variables are minimised in the last two terms of the objective function. Constraint (7.11) prevents variables d_{il}^{-1} and d_{il}^{-2} from taking values of 1 simultaneously for full-time nurses. This imposes a lower limit of 2 shifts for a full time nurse in any week, in order to avoid the penalty of 90 which is incurred when fewer than 2 shifts are assigned (S5).

Constraints (7.12) and (7.13) similarly work to penalise deviations from the range of 2-3 shifts per week for the part-time nurses (S6). An equivalent to constraint (7.11) is not required, since a) a shortfall of 3 shifts in a week is not possible as the minimum of the ideal range is only 2, and b) a surplus of 3 shifts is not possible as none of the six day stints are included in the set of stints $\mathcal{J}^P$ which can be worked by part-time nurses, and any two stints are separated by two off duty days by constraint (7.3).

Finally, constraints (7.14) and (7.15) limit the total number of shifts that can be worked during the rostering period, by a full-time or part-time nurse to 20 and 11 respectively. In ORTEC01, one of the full-time nurses works fewer hours than the others, and so a similar constraint is required for that particular nurse.

7.2.2 Stage 2 - The Shift Allocation Model

The purpose of this model is to designate each of the daily shifts assigned by the stage 1 rostering model. Some of these shifts will already have been identified

7.2 The Nurse Rostering Model

as Night shifts in the solution generated by the rostering model. The remainder need to be allocated as either Early, Day, or Late shifts. This shift assignment model takes as input the roster created at the previous stage, and specifies the type of each shift in accordance with the relevant soft constraints S9-S11. We use the following notation to describe the shift allocation model.

Indices

- i nurse, $i = 1, \dots, 16$
- k day of roster, $k = 1, \dots, 31$
- s length of series of working days

Sets

- $\mathcal{J}^F$ the set of Full-Time nurses
- $\mathcal{J}^P$ the set of Part-Time nurses

Data (from solution of rostering model at stage 1)

- $\alpha_{ik} = 1$ if nurse i is assigned a Day shift from the solution at stage 1 on day k , or 0 if they are off-duty or working a Night shift
- $n_{ik} = 1$ if nurse i starts a stint of night shifts on day k , 0 otherwise

Data (from the problem instance)

- r_k^D the number of Day shifts required on day k
- r_k^E the number of Early shifts required on day k
- r_k^L the number of Late shifts required on day k

Binary Decision Variables

d_{ik} =1 if nurse i is assigned a Day shift on day k

e_{ik} =1 if nurse i is assigned an Early shift on day k

l_{ik} =1 if nurse i is assigned a Late shift on day k

p_{ik} =1 if nurse i is assigned a Day shift on day k followed by an
Early shift on day $k + 1$

q_{ik} =1 if nurse i is assigned an Early shift on day k followed by a
Night shift on day $k + 1$

y_{iks} =1 if nurse i begins a series of s Early shifts on day i , where $s=1,4,5,6$

z_{iks} =1 if nurse i begins a series of s Late shifts on day i , where $s=1,4,5,6$

The data element α_{ik} is set to zero on days k when nurse i is on a night shift in order to reduce the data requirements of the model. Similarly, n_{ik} denotes only the start of a stint of night shifts, rather than all night shifts worked by nurse i .

Note that in the rostering model, l was used to index the week of the roster and d to denote deviational variables related to the number of shifts worked each week. Since neither of these are required in this shift assignment model, for convenience we introduce the variables d_{ik} and l_{ik} to denote day and late shift assignments respectively. We use s in this model to index a length of working stint, or series of shifts. Since there are no penalties for assigning series of 2-3 early or late shifts, we only need define x_{iks} and y_{iks} variables for series length 1,4,5 or 6.

7.2 The Nurse Rostering Model

Integer Programming Formulation for the Shift Assignment Model

Minimize

$$\sum_{i \in \mathcal{I}} \sum_{k=1}^{31} \{90(y_{ik6} + z_{ik6}) + 40(y_{ik5} + z_{ik5}) + 10(y_{ik4} + z_{ik4} + y_{ik1} + z_{ik1}) + 5p_{ik} + q_{ik}\}$$

subject to

coverage

$$\sum_{i \in \mathcal{I}} d_{ik} \geq r_k^D \quad \text{for } k = 1, \dots, 31 \quad (7.16)$$

$$\sum_{i \in \mathcal{I}} e_{ik} \geq r_k^E \quad \text{for } k = 1, \dots, 31 \quad (7.17)$$

$$\sum_{i \in \mathcal{I}} l_{ik} \geq r_k^L \quad \text{for } k = 1, \dots, 31 \quad (7.18)$$

designate type of “day” shift

$$d_{ik} + e_{ik} + l_{ik} \leq \alpha_{ik} \quad \forall i, k \quad (7.19)$$

shift changes

$$l_{ik} + e_{ik+1} + d_{ik+1} \leq 1 \quad \forall i, \text{ for } k = 1, \dots, 30 \quad (7.20)$$

$$d_{ik} + e_{ik+1} - 1 \leq p_{ik} \quad \forall i, \text{ for } k = 1, \dots, 30 \quad (7.21)$$

$$e_{ik} + n_{ik+1} - 1 \leq q_{ik} \quad \forall i, \text{ for } k = 1, \dots, 30 \quad (7.22)$$

7.2 The Nurse Rostering Model

set y and z variables for $s = 4, 5, 6$

$$\sum_{k'=k}^{k+5} e_{ik'} - 5 \leq y_{ik6} \quad \forall i, \text{ for } k = 1, \dots, 26 \quad (7.23)$$

$$\sum_{k'=k}^{k+4} e_{ik'} - 4 \leq y_{ik6} + y_{ik-1,6} + y_{ik5} \quad \forall i, \text{ for } k = 1, \dots, 27 \quad (7.24)$$

$$\begin{aligned} \sum_{k'=k}^{k+3} e_{ik'} - 3 \leq y_{ik6} + y_{ik-1,6} + y_{ik-2,6} \\ + y_{ik5} + y_{ik-1,5} + y_{ik4} \quad \forall i, \text{ for } k = 1, \dots, 28 \end{aligned} \quad (7.25)$$

$$\sum_{k'=k}^{k+5} l_{ik'} - 5 \leq z_{ik6} \quad \forall i, \text{ for } k = 1, \dots, 26 \quad (7.26)$$

$$\sum_{k'=k}^{k+4} l_{ik'} - 4 \leq z_{ik6} + z_{ik-1,6} + z_{ik5} \quad \forall i, \text{ for } k = 1, \dots, 27 \quad (7.27)$$

$$\begin{aligned} \sum_{k'=k}^{k+3} l_{ik'} - 3 \leq z_{ik6} + z_{ik-1,6} + z_{ik-2,6} \\ + z_{ik5} + z_{ik-1,5} + z_{ik4} \quad \forall i, \text{ for } k = 1, \dots, 28 \end{aligned} \quad (7.28)$$

set y and z variables for $s = 1$

$$e_{ik} - e_{ik-1} - e_{ik+1} \leq y_{ik1} \quad \forall i, \text{ for } k = 1, \dots, 30 \quad (7.29)$$

$$l_{ik} - l_{ik-1} - l_{ik+1} \leq z_{ik1} \quad \forall i, \text{ for } k = 1, \dots, 30 \quad (7.30)$$

$$d_{ik}, e_{ik}, l_{ik}, p_{ik}, q_{ik}, y_{iks}, z_{iks} \in \{0, 1\}$$

The objective function penalises the assignment of series of early or late shifts which are longer or shorter than the preferred length of 2-3 (in accordance with

7.2 The Nurse Rostering Model

soft constraint S9), and also any undesirable shift changes on consecutive days. The first term in the summation, $90(y_{ik6} + z_{ik6})$, applies a penalty weight of 90 if nurse i is assigned a series of 6 consecutive early or late shifts starting on day k , in accordance with S9. The second term applies a penalty of 40 for a series of 5 early or late shifts, and similarly the third term applies a penalty of 10 for a series of 4 or 1, again as per S9. The final two terms relate to undesirable shift changes. A penalty of 5 is applied if nurse i is assigned a backward rotation from a day shift on day k to an early shift on day $k + 1$, denoted by p_{ik} (S10). Similarly, the assignment of an early shift on the day before commencing a stint of night shifts q_{ik} attracts a penalty of 1 (S11). Again, the penalty weights are taken directly from the problem instance ORTEC01, although they could be changed to reflect local priorities.

The first three constraints are coverage constraints. Constraints (7.16), (7.17) and (7.18) respectively ensure that the requisite number of day shifts, r_k^D , early shifts, r_k^E , and late shifts r_k^L are assigned on each day k . Constraint (7.19) ensures that a nurse can only be assigned a day, early or late shift if they were scheduled, by the stage 1 rostering model, to be both on duty and not working a night shift.

The next three constraints relate to undesirable shift changes. Constraint (7.20) prevents the assignment of either a day shift, d_{ik+1} , or an early shift, e_{ik+1} on the day following a late shift assignment l_{ik} , thus imposing the minimum rest period of 11 hours between shifts (H4). Constraint (7.21) sets the penalty counter p_{ik} to be 1 whenever an early shift is assigned to immediately follow a day shift, i.e. when $d_{ik} + e_{ik+1} = 2$. Similarly, constraint (7.22) sets variable q_{ik} to be 1 when an early shift precedes a night shift. The variables p_{ik} and q_{ik} are used to impose penalties in line with soft constraints S10 and S11.

The next six constraints, (7.23) to (7.28), all deal with the length of a sequence of early or late shift assignments, the ideal length being 2-3 days (S9). Constraint (7.23) sums the number of early shifts assigned to nurse i in each period of six consecutive days in the roster, $\sum_{k'=k}^{k+5} e_{ik'}$, and defines variable y_{ik6} to be 1 if nurse i works 6 consecutive early shifts, beginning on day k . Constraint (7.24) similarly identifies sequences of 5 early shifts. However, we need to take into account the fact that such a sequence may be part of an even longer one. Constraint (7.25) similarly identifies the sequences of 4 early shifts, y_{ik4} . The next three constraints (7.26) to (7.28) are identical to (7.23) to (7.25), except that they relate to late rather than early shifts.

The final two constraints identify isolated early or late shifts. Constraint (7.29) sets the value of variable y_{ik1} to be 1 if nurse i works an early shift on day k , (and hence, $e_{ik} = 1$), but does not work an early shift on either day $k - 1$ or day $k + 1$ (so that $e_{ik-1} + e_{ik+1} = 0$). Constraint (7.30) is identical but relates to late rather than early shifts.

7.3 Chapter Summary

We have identified a number of similarities between the nurse rostering and call centre rostering problems. These similarities, and the availability of benchmark instances for the nurse rostering problem motivated us to consider the applicability of our call centre rostering models to the domain of nurse rostering. We presented an analysis of the heuristic techniques and mathematical models used by leading researchers in the field. We found the decomposition of the nurse rostering problem into tour scheduling and tour assignment stages to be imprac-

tical for a number of reasons, one of which was a methodological issue relating to continuity between rostering periods, an issue which we will address in the next chapter, which forms our nurse rostering paper.

We presented our nurse rostering model, based on an alternative two-stage decomposition of the problem. The first stage Integer Programming model determines the days to be worked or off-duty for each employee, and superimposes stints of 2-3 night shifts where appropriate. The second stage Integer Programming model allocates specific shift types to those working days not already designated as night shifts. The formulation for both models is tailored to the benchmark problem instance “ORTEC01”. Taken together, this chapter and the next form the final contribution of this thesis.

Chapter 8

The Nurse Rostering Problem

This chapter presents a paper submitted to the European Journal of Operational Research (Glass & Knight (2008b)). For ease of continuity in the context of this thesis, the paper has been slightly restructured so that information from the paper's Appendix is presented here in the body of the chapter.

This paper is concerned with the problem of nurse rostering within hospitals. We analyse two benchmark problems from the nurse rostering literature to provide insight into the nature of the problem. By highlighting the structure of the problem we are able to reduce the relevant solution space. A Mixed Integer Linear programme is then able to find an optimal solution to each of the two single month problem instances in only a few minutes. Our second contribution is to extend current mathematical approaches to nurse rostering to take better account of the practical considerations. We provide a methodology for handling rostering constraints and preferences arising from the continuity from one scheduling period to the next.

8.1 Introduction

The nurse rostering problem is to decide which members of a team of nurses should be on duty at any time, during a rostering period of, typically, one month. This involves determining the days on and off duty for each nurse, as well as their shift start and finish times on each working day. The quality of a roster can be judged in terms of compliance with legal requirements (for example, the European Working Time Regulations¹), employment contracts, and the collective and individual preferences of the nurses themselves. There are a wide range of considerations which constitute rostering “good practice”. The UK Government’s Health and Safety Executive (HSE) published a number of guidelines in HSE (2006), and a set of 10 ergonomic principles was presented by Kundi (2003). Poorly designed rosters can lead to inadequate periods of rest and recovery time, and to socially unacceptable patterns of working. Such outcomes are stressful to the nurses, and the negative impacts upon their productivity and levels of engagement are well documented in the literature on the ergonomics of shift working (see, for example, Totterdell (2005)). The complexity of the task has led to the use of automated rostering systems in some European hospitals, although in the UK the task is usually carried out manually. For these reasons, there is a need to develop nurse rostering methodologies which have the potential for practical application.

There is a large body of research literature relating to nurse rostering. The survey papers by Ernst *et al.* (2004a,b) and Burke *et al.* (2004b) give a detailed overview of the literature up to 2005. However, the latter conclude that only a

¹European Working Time Directive (93/104/EC)

few papers are based on real world data, or address the development of rostering systems for implementation in hospitals. A recent paper, Burke *et al.* (2008), was jointly authored by researchers at Nottingham University and at ORTEC¹, a major supplier of rostering software. That paper presents a practical, heuristic based methodology and compares results from variable neighbourhood local search heuristics with those obtained using ORTEC's Harmony software on a range of problem instances. Since the context is the same of the recent article by Burke *et al.* in this journal (EJOR), we omit a detailed literature review in order to avoid unnecessary repetition.

Our main research interest is in employee rostering for call centres. In this domain, the application of Mixed Integer Programming (MIP) techniques is more common than the use of heuristics. The traditional approach to call centre rostering is to decompose the problem into two stages. First, a set of weekly *tours* are scheduled, where each tour consists of a set of attendance details; days on and off duty, daily shift start and finish times, and even lunch and coffee break timings. At a second stage, each tour is assigned to an individual employee, taking account of availability and preferences. The problem is complicated by the large number of feasible shifts, which may start or finish in any reasonable quarter-hour interval, unlike in nurse rostering where generally only a handful of shifts are considered. The most comprehensive tour scheduling model in the call centre literature is that of Brusco & Jacobs (2000). This model allows the specification of a *bandwidth* parameter which limits the variation in shift start-times within a working week. A major problem is that this approach does not distinguish between forward and backward rotation. For example, a typical restriction in

¹ORTEC, Groningenweg 6-33, 2803 PV Gouda, Holland, The Netherlands

practice is that a day shift cannot immediately follow a night shift, as indicated by H5 in the list of hard constraints in Table 8.3. If this condition is imposed through a bandwidth restriction, then the same constraint prevents a night shift from appearing *anywhere* in the same tour as a day shift. Hence, bandwidth has little relevance in the context of nurse rostering, where specific rules relating to night shift working (for example, constraints H3,H5,H10,S10) are typically imposed. In the rostering literature relating to healthcare, Isken (2004) developed a tour scheduling model which treated start time variations in the same way as Brusco & Jacobs (2000), albeit with a different formulation. The Isken model therefore suffers from the same drawback.

We have developed an alternative tour scheduling model, presented in our paper Glass & Knight (2008a). We recognise that while it is preferable to schedule the same shift start time (or later) on consecutive working days, there is usually no objection to any shift change following a period of one or two days off-duty.

An important issue in both call centre and nurse rostering is the continuity from one rostering period to the next. The nurse rostering benchmark instances are designed only to produce rosters for an isolated period, applying penalties in accordance with the convention that all potential violations are counted at the beginning of the period, and ignored at the end. We recognise that the benchmark instances are intended as a basis for comparison between alternative rostering methodologies, and that the consideration of an isolated rostering period serves this purpose. However, in a practical environment, information relating to one rostering period is carried forward to the next, creating additional considerations of “continuity”.

Our contribution to the mathematical nurse rostering literature is as follows.

1. Analysis of a benchmark problem instance which identifies a lower bound on the solution, and which brings to light aspects of the problem structure which we are later able to exploit in order to radically reduce the size of solution space.
2. Improved results for two benchmark problems. We have optimally solved these two problems with only a few minutes execution time, using an MIP approach implemented in MPL/CPLEX¹ on a standard desktop PC.
3. Improvements in handling shift changes within an MIP, extended from the context of call centre rostering to nurse rostering.
4. A methodology for handling continuity between rostering periods.

The paper is organised in two parts. In the first part we analyse two benchmark problems initially specified by ORTEC, each of which involves a single rostering period. We give a structural analysis of the interrelationship between constraints leading to a lower bound on penalty costs for one of the problems. We also give optimal solutions for both problems and compare our results to those obtained previously.

In the second part of this paper we propose a more flexible approach, handling those requirements which relate to the continuity between rostering periods. We are aware that these issues are well understood and taken account of in practice. Indeed, the benchmark problems originate from practice. Our contribution is to formalise the continuity goals. We illustrate how our approach achieves these objectives by producing a solution which continues from one month to the next.

¹Maximal Software Inc.

8.2 Benchmark Problems

In this section, we describe two benchmark problems originally provided by ORTEC. Both of these problems are among a number of nurse rostering benchmark instances which are maintained and made available by Nottingham University¹. The first problem is referred to as “GPost” and the second as “ORTEC01”. Each problem instance consists of a set of daily staffing requirements, a set of rostering “rules”, and a set of penalties for rule violations.

“GPost” is a small introductory problem for 8 nurses across a rostering period of exactly 4 weeks. There are only two nurse contracts, full and part time, and two shift types, day and night. However, this problem contains many of the structural elements of the larger, more complex nurse rostering examples.

“ORTEC01” is a larger problem for 16 nurses, for the 31 days of January 2003. There are 4 shifts: Early, Day, Late and Night (E,D,L and N). The problem is fully described in Burke *et al.* (2008), except for a couple of individual rostering constraints which are not explicitly listed in the academic papers but are embedded in the solutions presented on the Nottingham University website. The full set of rostering rules for both problems are as follows.

8.2.1 Problem Descriptions and Rostering Constraints

Instance “GPost”

Staff availability and shift patterns. This problem has 8 nurses. A, B, C and D are full-time and work 36hrs/week, and E, F, G, H are part time and work 20

¹<http://www.cs.nott.ac.uk/~tec/NRP/>

hrs/week. There are only Day (D) and Night (N) shifts.

Staffing requirement. The requirement is for three day shifts and one night shift per day. The rostering period is for four weeks from Monday 2nd January to Sunday 29th January 2006.

Pre-assigned shifts (hard constraints)

Nurse A: Day shift on on 2nd and 3rd January.

Nurse C: Day shift on on 2nd and 3rd January.

Nurse D: Night shift on on 2nd and 3rd January.

Nurse E: Day shift on on 2nd and 3rd January.

Instance “ORTEC01”

Staff availability and shift patterns. This problem has 16 nurses. 12 nurses, A,B,D,E,F,G,J,K,L,M,N,P are full time and work 36hrs per week. One nurse, O, works 32 hrs/week and the other 3, C,H,I, are part time and work 20hrs/week. There are four shift types, Early (E), Day (D), Late (L) and Night (N).

Staffing requirement. The requirement is for three of each of E, D and L and one N shift on each weekday. On Saturday and Sunday the requirement is for two each of E,D and L and one N. The rostering period is for the 31 days from Wednesday 1st January to Friday 31st January 2003.

Individual requests (soft constraints)

Nurse *A*: Day shift on Monday 6th January (penalty 100 for violation).

Nurse *M*: No Late shifts for the whole month (penalty 10000 for violation).

Common structure and constraints

Some aspects of the problem structure are common to both Gpost and ORTEC01.

- No under or over-staffing is allowed.
- The night shift starting on a Friday night is considered a weekend shift.
- The working week runs from Monday to Sunday.

The constraints for the two problems are given in Tables 8.1, 8.2, and 8.3.

8.2 Benchmark Problems

	Soft Constraint	Preferred range	Gpost penalty	ORTEC01 penalty
S1	Single night shift		100	1000
S2	Single weekend shift		100	1000
S3	Standalone shift		100	1000
S4	Single day off		10	100
S5	Number of shifts per week (FT)	4-5	$1 \times n^2$	$10 \times n^2$
S6	Number of shifts per week (PT)	2-3	$1 \times n^2$	$10 \times n^2$
S7	Stint length (FT)	4-6	$1 \times n^2$	$10 \times n^2$
S8	Stint length (PT)	2-3	$1 \times n^2$	$10 \times n^2$
S9	Series of Early or Late shifts	2-3	N/A	$10 \times n^2$
S10	Early shift follows Day shift		N/A	5
S11	Night shift follows Early shift		N/A	1

Table 8.1: Soft constraints for Nurse Rostering benchmarks (where n is the deviation from the preferred range).

8.2.2 General Approach

Our general approach to both problems is to initially treat those soft constraints with high penalties as being hard constraints. Should this lead to infeasibility, we relax the constraints incrementally in order of their cost, until a feasible solution is found. This approach is fairly standard and aligns that of Burke *et al.* (2008), where a rule relating to the number of night shifts to be worked in a 5 week period is listed both as a hard constraint, and also as a soft constraint with a high penalty. All our experiments were performed using a desktop PC with a P4 2.67 GHz processor.

We observe that, in both problems, the staffing requirement is fairly consistent on a daily basis, and throughout the rostering period. This imposes a high degree of structure on the problem which we can exploit. We start by analysing the problem structure before addressing each of the two instances individually.

8.2 Benchmark Problems

	Continuity	Constraint Convention	
		Start period	End Period
C1	Too short stints	Penalised	Allowed
C2	Too few days in week	Penalised	Allowed if partial week Penalised if full week
C3	Single rest day	Penalised	Allowed
C4	Single night shift	Penalised	Allowed
C5	Standalone shift	Penalised	Allowed

Table 8.2: Penalty exemptions for Nurse Rostering benchmarks

	Hard Constraint	Limit	
		GPost	ORTEC01
H1	Maximum shifts per nurse per day	1	1
H2	Maximum number of consecutive shifts	6	6
H3	Maximum number of consecutive nights	3	3
H4	Minimum time between shifts	N/A	11 hrs
H5	Minimum time off after night shifts	48 hrs	42 hrs
H6	Total shifts for 36 hr/wk contract	18	20 per week
H7	Total shifts for 32 hr/wk contract	N/A	18 per week
H8	Total shifts for 20 hr/wk contract	10	11 per week
H9	Maximum number of overtime hours	N/A	4 hrs per month
H10	Maximum number of night shifts worked (rolling)	4	3 in 5 weeks
H11	Maximum proportion of weekends worked (rolling)	2 of 3	3 in 5 weeks
H12	Maximum average hrs/week (rolling) excluding night shifts	N/A	36 hrs over 13wks

Table 8.3: Hard constraints for Nurse Rostering benchmarks

8.3 Analysis of Problem Structure of “GPost”

The rules relating to night shifts (H3, H5 and the fact that Friday night is treated as a weekend shift) impose a strong structure on the optimum solution, described as follows.

Proposition 1 *For any problem where the rostering period is a number of complete weeks, and the same number of night shifts are required on each night of the week, any solution which avoids incurring the penalty for a single night shift ($S1$), or for a single weekend shift ($S2$), has the same sequence of night shift stints, namely 2-2-3, in each week of the roster.*

PROOF. We begin with the following observations regarding night shift working.

(1) If a weekend with one shift is not allowed (constraint $S2$), and Friday night is counted as a weekend shift, and a series of night shifts is to be followed by 48 hours off duty ($H5$), then a series of night shifts cannot end on Friday.

(2) Given that the same staffing level must be provided on Saturdays and Sundays, then if a series of night shifts ends on a Saturday, one of the Sunday shifts must be covered by a nurse beginning a new stint of duties, giving an isolated weekend day, which is not allowed. (Otherwise, one more Saturday than Sunday duty will be provided).

Therefore, since the night shifts are to be worked in series of 2-3 days, the only valid pattern of coverage over the weekend is to cover Friday, Saturday and Sunday night with a single series. Since a single night shift is required on each of the four other days of the week, this means that the only valid pattern for

8.3 Analysis of Problem Structure of “GPost”

covering the night shifts over whole of the 4 week rostering period is in a weekly pattern of Monday + Tuesday, Wednesday + Thursday, Friday + Saturday + Sunday. $\square$

Corollary 2 *Any solution to problem “ORTEC01” which avoids incurring the penalty for a single night shift (S1) will include at most one stint of either 2 or 3 night shifts for each employee.*

PROOF. Hard constraint H10 imposes a limit of 3 on the total number of night shifts to be worked by any employee in a rolling 5 week period. Since the rostering period in ORTEC01 is less than 5 weeks, there is an upper limit of 3 night shifts per employee in the solution. In a solution which avoids single night shifts, all night shift stints will be of length 2 or 3. Hence only one stint per employee is possible. $\square$

We can now show that the “GPost” problem instance has a lower bound of 5 and that these penalty costs arise in the manner presented in our solution below. We now consider the interrelationship between the pre-assignments of nurses to specific shifts at the start of the roster, and the night shift pattern derived earlier, in order to prove a lower bound on the optimal solution.

Proposition 3 *Problem “GPost” has no feasible solution with penalty less than 5. In any solution, Nurse “D” attracts a minimum penalty of 4, and Nurse “B” attracts a minimum penalty of 1.*

8.3 Analysis of Problem Structure of “GPost”

PROOF. Nurse “D” is pre-assigned to night shifts on days 1 and 2 (Mon, Tue). Since all night shifts are in blocks of Mon+Tue, Wed+Thu, Fri+Sat+Sun, and a night shift cannot be followed a day shift, Nurse “D” must work a two-day stint on days 1 and 2, attracting a penalty cost of 4 and proving the first part of the proposition. We now prove by contradiction that Nurse “B” attracts a minimum penalty of 1.

Suppose that there is an optimal roster in which Nurse “B” incurs no penalty. Consider the first week. All four shifts required for each of days 1 and 2 are pre-assigned. Since there can be no overstaffing in the solution, Nurse “B” must be off duty on these two days. Thus, in order to avoid penalties, 4 or 5 shifts including the whole weekend must be worked in week 1, including days 4-7. This partial roster is shown diagrammatically below, where “O” denotes an off-duty, and “X” a working day.

DAY OF ROSTER																												
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28
B	O	O		X	X	X	X																					

Now consider week 2. Were we to assign the weekend off (days 13-14), then in order to avoid violating soft constraint S5, 4-5 shifts would need to be worked between Monday and Friday (days 8-12). A 4 or 5-day stint starting on Monday would result in 8 or 9 consecutive days being worked across weeks 1 and 2, violating hard constraint H2. Alternatively, a 4-day stint starting on Tuesday would result in a single off duty day on Monday violating S4. Therefore our only cost free option is to assign the second weekend (days 13-14) as working days, which forces us, by constraint H11, to assign the third weekend (days 20-21) as off duty, as illustrated below.

8.4.1

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28
B	O	O		X	X	X	X						X	X						O	O							

Our two benchmark programs are shown as optimal solutions to the

Now consider week 3. To avoid penalties, we need to assign 4 or 5 working days between Monday and Friday (days 15-19). A five-day stint gives rise to 7 consecutive working days (13-19) and a four-day stint from Tuesday to Friday (days 16-19) gives a single day off on Monday. Therefore our only cost-free option is to assign four shifts to complete a six-day stint on days 13-18. This maximum length stint must be preceded by two off-duty days (11-12).

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28
B	O	O		X	X	X	X				O	O	X	X	X	X	X	X	O	O	O							

At this point, week 2 still has only two working days (13 and 14). We need to assign a minimum of two more shifts between days 8 and 10. Assigning 3 days gives too long a stint (days 4-10), and a two-day stint on days 9 and 10 leaves an isolated day off, so we must assign two working days on days 8 and 9. This gives another six-day stint from days 4-9, so day 3 must now be off duty.

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28
B	O	O	O	X	X	X	X	X	X	O	O	O	X	X	X	X	X	X	O	O	O							

We have now assigned a total of 12 shifts, leaving 6 more to complete the rostering period. All six of these shifts must be worked during week 4, incurring a minimum penalty cost of 1 and proving the required contradiction. $\square$

8.4 Optimal Solutions

We now present the optimal solutions to the two benchmark instances.

8.4.1 Optimal Solution for “GPost”

Our mathematical programme produced an optimal solution to the “GPost” problem with a penalty cost of 5. The previous best-known solution was 7, an improvement of 29%. The complete solution is shown in Figure 8.1. The program took 22 minutes to find the solution. The breakdown of the total penalty cost is shown in Table 8.4, where labels S5 and S7 refer to the list of soft constraints in Table 8.1.




Figure 8.1: Optimal solution for problem “GPost”

Nurses	Constraint Violation	Unit Penalty	Total Penalty
B	S5 (6 shifts in week 4)	1	1
D	S7 (2-day stint in week 4)	4	4
Total Penalties			5

Table 8.4: Penalties for problem “GPost”

8.4.2 Optimal Solution for “ORTEC01”

Our model reports the optimal solution to the “ORTEC01” problem as 270, after an execution time of only 2 minutes. As before, we constrained the night shifts to fit the only valid pattern, and generally treat those constraints with penalties of 90 or higher as hard constraints initially. The full solution is shown in Figure 8.2.

DAYOF MONTH (JAN 03)																																	
Week 1					Week 2					Week 3					Week 4					Week 5													
W	Th	F	Sa	Su	M	Tu	W	Th	F	Sa	Su	M	Tu	W	Th	F	Sa	Su	M	Tu	W	Th	F	Sa	Su	M	Tu	W	Th	F			
1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31			
A	L	L	L			D	D	D	L	L					E	E	L	L	N	N			D	D	D	L	L				D		
B	L	L	L				E	E	E	D			E	E	E	L	L			E	E	E	D				D	D	D	N			
C	D	D							N	N	N				E	E	D						E	E	E								
D	L	L	L			E	E	E	L	L			E	E	L	L	L					D	N	N	N			L	L	L			
E	E	E	E						D	D	D	D				D	L	L	L					E	E	E	L	L	L				
F			E	E	E	L	L	L			E	E	L	L	N	N						E	E	L	L	L				E	E		
G				D	D	D	L	L	L						D	D	N	N	N			E	E	D	D	D	D			E	E		
H	N	N					L	L							E	E			L	L	L						D	D					
I	D	D				N	N						L	L					D	L	L						L	L					
J	E	E	E				E	E	E	L	L	L			E	E	E	D			L	L	L	N	N					D			
K	D	D	N	N	N			E	E	D	D	D	D			E	E	E	D	L	L						E	E	E				
L			D	L	L	L			E	E	E	L	L	L			D	D	L	L	N	N				E	E	E	D				
M	E	E	D			E	E	D	D	D			D	D	D	D			E	E	D	D				D	D	N	N				
N				E	E	E	D	N	N				E	E	E	L	L			D	D	L	L			E	E	D	L	L			
O				D	D	D	D	D	D				D	D	D	D			E	E	E	L	L							D	D	D	
P		D	L	L	L					L	L	L	N	N			D	D	D	D	D	D				E	E	E	L	L			

Figure 8.2: Optimal solution for problem “ORTEC01”

The penalties incurred by this optimal solution are shown in Table 8.5. It is interesting to note that the entire penalty of 270 is incurred during the first (partial) week of the month. The relevant soft constraints S5 and S7 relate to the number of shifts to be worked by full-time nurses during a week, and to the length of stints to be worked. These both relate to the shift content of a working week, which is counted from Monday to Sunday. However, the first week contains

only Wednesday to Sunday.

Nurses	Constraint Violation	Unit Penalty	Total Penalty
G,N,O	S5 (2 shifts in week 1)	40	120
A,B,D,E,F,J,L,M,P	S5 (3 shifts in week 1)	10	90
A,B,D,E,J,M	S7 (3 day stint in week 1)	10	60
Total Penalties			270

Table 8.5: Penalties for problem “ORTEC01”

Previous best-known solutions to “ORTEC01” have been found by ORTEC themselves, using their Genetic Algorithm (GA), and by Burke *et al.* (2006, 2008) using the Variable Neighbourhood Search methodology (hybrid VNS). These solutions are listed in Burke *et al.* (2008), and are summarised in Table 8.6 below. The execution times were achieved using comparable computers. Burke *et al.* used a P4, 2.4 Ghz processor PC, whereas ours has a clock speed of 2.67 Ghz. The timeline of the development is, however, not mentioned by the authors.

We have thus achieved an improvement in quality, as measured by the penalty function, of 10 which is equivalent to over 3%. In the practical setting where results may be required within an hour, we incur a penalty cost of not less than

Penalty	Found By	Execution Time
775	GA	1 hr
681	GA	24 hrs
587	GA (iterative)	“long”
706	Hybrid VNS	1 hr
541	Hybrid VNS	12 hrs
280	Hybrid VNS	16+ hrs
270 *	IP	2 minutes
* optimal solution		

Table 8.6: Historical solutions to ORTEC01

40% of the current best of 706.

We recognise that our short execution time is partly due to the availability of a feasible solution with no penalties of 90 or more. Should this have proved infeasible, we would have relaxed one or more of the higher penalty constraints and would expect the run time to increase as a result. On the other hand, the previous best-known solution of 280 took over 16 hours to achieve and may have required the Hybrid VNS model to be highly tuned to this particular problem instance.

We are aware that following some sharing of information, the Nottingham group have also identified the optimal solution to ORTEC01, using a form of relaxation. This supports our view that the methodology used is less important than the problem analysis. In other words, the analysis presented in this first part of the paper makes the two benchmark instances accessible to both heuristic and MIP approaches. We have applied an MIP approach adapted from our call centre rostering model, but simply based upon the constraints which are fully elaborated in section 8.2. Our particular implementation is an adaptation of Glass & Knight (2008a). We recognise that problems of higher levels of complexity, such as those involving multiple nurse grades, may be computationally impractical for IP methods.

8.5 Month to Month Continuity

In this second part of the paper we address the issue of continuity between rostering periods. One issue with the ORTEC01 benchmark is that although the rostering period is one month in length, the constraints do not primarily relate to

a one month period. Of those constraints which relate to periods of time, some relate to one week (S5-S6 , H6-H8), two to a rolling five week period (H10, H11), and one to a rolling thirteen-week period (H12). Only one constraint, H9, applies to a calendar month. Part of the difficulty of producing nurse rosters is the handling of rules relating to various time periods. Our approach to this problem is to regard the one-month rostering period as being primarily a *reporting* period, rather than the basis for roster calculation. The idea behind our methodology for continuous rostering is to calculate the roster across a number of whole weeks, including the transition week from the current rostering period to the next one.

8.5.1 Advantages

Before describing our approach we highlight the significance of the continuity issue for the ORTEC01 solution in Figure 8.2. The rostering period begins on a Wednesday, but the working week runs from Monday to Sunday. All of the penalties for our optimal solution to ORTEC01 relate to nurses working too few shifts during the first (shortened) week, or too short stints at the start of it. In total, 12 nurses attract such penalties. However, the penalties are applied on the assumption that previous roster was empty. If, in fact, 10 nurses worked on 1st and 2nd January (the standard weekday requirement), then many of the penalties would not be incurred. The penalties as applied at the start of the rostering period are therefore highly pessimistic.

At the other end of the rostering period the convention is not to penalise too-short stints, the optimistic assumption being that these will be extended into the next rostering period. However, this will not always be possible. Consider, for

example, the roster for Nurse “J”. This nurse has been assigned three consecutive weekends and will therefore be hard-constrained to take the next weekend off (by constraint H10). This next weekend falls on the first two days of the next rostering period, thus leaving a single working day on 31st January, with penalty of 1000 (by constraint S3). Similar situations are also apparent in some of the solutions to ORTEC01 published on the Nottingham University website. So, although our roster presented in Figure 8.2 is optimal with regard to the conventional application of penalties at the beginning and end of an isolated rostering period, it is not a very good solution in the context of continuous rostering.

We therefore conclude that the minimum total penalty on our apparently optimal solution to ORTEC01 is 1270 when formulated in the current fashion and attention is restricted to the single month. We show in section 8.5.3 that a more realistic approach to the treatment of month-to-month continuity produces a higher quality solution. The associated penalty is reduced to 20, in other words to less than 2%.

8.5.2 Approach

Our overall approach to the nurse rostering problem is to apply an adaptation of our MIP model for call centre rostering, presented in our paper Glass & Knight (2008a). We have further extended this model to handle continuity from one month to the next based on the following two component ideas. (1) At the start of each rostering period, *continuity constraints* are generated for each nurse in order to accurately reflect the costs incurred during the transitional week from the *previous* rostering period to the current period. (2) *Implied penalties* are

introduced. These are standard penalties as specified in the problem definition, but which will necessarily be incurred during the *following* rostering period, due to decisions made during the current period (such as the penalty of 1000 for a single working day in ORTEC01, as described above). Our idea is to identify these and account for them in the current period, giving a more realistic measurement of roster quality. Such penalties are easier to identify when rostering whole weeks, since a number of constraints relate to the shift content of a week. The implied penalties give a lower bound cost for the following rostering period.

Note that although we are producing a roster which extends beyond the end of the one-month rostering period, we are not advocating that the rostering period be changed. We recognise that the shorter rostering period can be beneficial to nurses in that it allows a shorter notice period within which individual requests can be accommodated. Our aim is to support the monthly notification of rosters, and the purpose of extending the roster calculations is to establish a lower bound for the following period. The unpublished part of the 5-week roster (i.e. that part of the transition week extending into the next period) can be re-designed in the light of changing circumstances, for example personal requests. We now demonstrate how this process works.

8.5.3 Empirical Example

Staying with the problem ORTEC01, we first generate a roster for the previous month (December 02). The aim of this exercise is to give us an example of late December assignments, so that we can take account of these when producing the January roster. Observe that we can impose all of the soft constraints and ten of

the twelve hard constraints (H1-H10) at the beginning of January by taking account at most one weeks roster from the end of December. To impose constraints H11 and H12, which specify a maximum number of weekends and night shifts (respectively) that an employee can work in a rolling five week period, we would need to take account of 4 weeks prior information relating to night and weekend assignments. H13 relates to a maximum number of shifts to be worked in a rolling 13 week period, and would therefore require us to maintain information for 12 previous weeks.

For ease of illustration, we will produce the January roster taking account of only those assignments shown in Figure 8.3. Only those late December assignments which have a direct bearing on the January roster are indicated. In order to impose some constraints such as those relating to the length of stints worked, it is not necessary to know whether a shift was an Early (E), Day (D) or Late (L), it is enough to know which days were working days. Where this is the case, we have used “X” to denote these unspecified shifts.

The December solution reports an implied penalty of 10 for Nurse A, who is assigned a 3-shift week in the transition week (30th December to 5th January). This penalty would not have been apparent from the published December roster (1st-31st December), and as such is an implied penalty which will be incurred when rostering the January period. Hence, this set of December assignments implies a lower bound of 10 on January’s roster.

We can now produce the January roster. Essentially, we produce a five week roster which covers the whole month, i.e. a roster for the period 30th December to 2nd February. In doing so, we first need to re-address the implications of the “hours-per-week” aspect of the employee contracts. Recall that the full time




Figure 8.3: December-January continuity, “ORTEC01”

nurses work a 36-hour week, which equates to 4.5 shifts per week. This implies that they work, on average, 22.5 shifts per 5 week period. We impose an upper limit of 23 shifts for the full-time nurses. Similarly, the part-time nurses on a 20-hr per week contract work an average of 12.5 shifts in a 5 week period, which we treat as an upper limit of 13. The 36-hour contract (Nurse “O”) works an average of 4 shifts per week and so we impose a constraint of 20 shifts.

Since the nurses have already been notified of their shifts for 30th and 31st December, we constrain the model to include in the solution these known assignments. Additionally, we generate constraints for each nurse to reflect the costs of extending working stints across the rostering periods. For example, Nurse “P” has worked a six-day stint up to 31st December, including Late shifts on the last two of these days. The January roster will therefore be constrained to place Nurse “P” off-duty on 1st January, and will apply the “single off-duty day” penalty if a shift is assigned on 2nd January. Those nurses who have recently worked night

shifts are not available for further night shifts until late in January (by constraint H10). Similarly, the weekend assignments on December 28th and 29th are taken into account when imposing constraint H11.

The complete January roster, together with those late December assignments affecting the continuity between the rostering periods, is shown in Figure 8.4. Note that the previously notified assignments (Dec 30th and 31st) are unaltered, but the remainder of the first week of the January roster has been substantially changed from the continuation identified as part of the December roster.

The solution has a cost of 20, and was found in 8min 15secs. The breakdown of penalties is given in Table 8.7. The solution is optimal given the data relating to the late December assignments. The two 10-point penalties are incurred by Nurse “A” (3-day week, from 30th December) and Nurse “L” (3-day week, from 13th January). The first of these was expected and is in accordance with the implied penalty identified in the solution for the previous rostering period. This time, no implied penalties are identified at the end of the rostering period and the continuation from January into February would appear to be straightforward. This solution is of a much higher quality than when ORTEC01 is solved for an isolated month, and we do not believe this gain to be atypical. We are aware that there are ad-hoc methods for handling the continuity issue in practice, however we provide a systematic means of incorporating month on month constraints within mathematical rostering approaches.




Figure 8.4: Optimal solution for “ORTEC01” with month-to-month continuity

Nurses	Constraint Violation	Unit Penalty	Total Penalty
A	S5 (3 shifts in week 1) (implied from Dec roster)	10	10
L	S5 (3 shifts in week 3)	10	10
Total Penalties			20

Table 8.7: Penalties for “ORTEC01” with continuity

8.6 Chapter Summary

We have solved two Nurse Rostering benchmark problems, “GPost” and “ORTEC01”, to optimality with only a few minutes execution time. We have demonstrated that our MILP model can take account of a set of realistic rostering constraints within a rostering period of one month. We have introduced the idea of calculating implied penalties in order to avoid accepting small improvements in the current rostering period at the cost of implicitly imposing larger penalties on

the next. We applied our MILP model to very quickly produce highly accurate results in a continuous rostering environment. The key to these developments has been the mathematical analysis of the problem structure. This lead to the identification of a restricted solution space which made the benchmark problems accessible to MIP solution methods as well as heuristics.

In terms of future development, we envisage the inclusion of considerations relating to personal preferences, some of which may be at odds with what would otherwise be seen as general preferences. For example, individuals with responsibilities as carers may need to work mostly night shifts or at weekends. The benefits (and dangers) of such prioritisation is discussed by Silvestro & Silvestro (2008). A related research direction is the development of additional mechanisms for ensuring equitable distribution of generally unpopular shifts, with the aim of improving the perceived “fairness” of the roster.

Chapter 9

Conclusions and Future Research Direction

In this final chapter we present a summary of the key points and conclusions of thesis, point to some some remaining gaps, and offer some direction for future research in these areas.

9.1 Thesis Summary and Conclusions

The concern of this thesis has been employee rostering, with particular reference to call centre telephone agents and hospital nurses. Motivated by the high cost of employee turnover in the UK call centre industry, and reports of poor levels of customer service, we began the thesis by questioning the role of call centre operations management, and in particular, the methods and models used in employee rostering.

In chapter 2 we identified two quite different sets of call centre management methodologies, each designed to support a different competitive strategy. The older of these two strategies dates back to the 1980's and seeks to establish competitive advantage by leveraging the scale economies which arise from efficiencies

9.1 Thesis Summary and Conclusions

inherent in call centre operations, in order to reduce operating costs. This strategy is supported by the mass-production model, characterised by standardised transaction processes, throughput based metrics and targets, and high levels of monitoring. In practice, this production model is implemented through a range of computer systems. Scripting systems lead the agent through each customer transaction, Automatic Call Distribution (ACD) systems report a wide range of statistics at both the team and individual employee level, and Work Force Management (WFM) systems generate employee rosters and also include functions such as “Agent Adherence”, which flags up occasions where employees are not doing (according to the ACD) what they should be doing (according to the rosters).

The more modern call centre culture, emerging over the last few years and increasing in popularity, is based not on cost reduction but on revenue generation and quality. The goals of this strategy are to increase the size of the customer base and market share, and to cross-sell alternative products to existing customers. This is achieved by deploying a knowledgeable, well-motivated, and unstressed workforce to deliver high quality services. Employee retention is thus an important factor in the success of this strategy, and hence, the high cost of employee turnover and the recent attention in the HR literature to “employee-friendly”, High Commitment Management (HCM) policies. At the Operations Management level, one of the key call centre systems used to implement the revenue-generation strategy is the Customer Relationship Management (CRM) system, which replaces the script-based transaction processing system used in the low-cost strategy. CRM presents the agent with a database of details for each customer, including their sales and service history, enabling the agent to use dis-

cretion in deciding, for example, whether to engage a customer in cross-selling. It is our perception, however, that in practice the same generation of WFM systems is used to support both strategies. In other words, the same methods and models are used to produce employee rosters in the quality-driven, revenue generation strategy as are used in the efficiency-driven, cost reduction strategy. The identification of the systems and models used to support the alternative management strategies and policies, and the requirement for an “employee-friendly” rostering model, moves toward the integration of call centre HR and Operations Management, and forms the first contribution of this thesis.

In chapter 3, we developed a set of 24 goals for an employee-friendly rostering system. These goals were identified from legislative requirements, best-practice guidelines and principles, and generalised preferences for shift sequences and patterns of off-duty. The best-practice guidelines are based on those presented by the UK Health and Safety Executive, supplemented and clarified by the academic literature on the ergonomic aspects of employee rostering. The generalised preferences reflect those criteria by which employees *as a group* judge the quality of a roster. While there is some overlap between these preferences and the ergonomic guidelines, others are added from experience of the UK call centre industry, which includes discussion and negotiation with trade unions at local and national level. Some of our goals relate to the tour scheduling stage of the rostering process, others apply to tour assignment, and many apply to both stages, indicating the desirability of the integration of the two stages. The identification of the set of employee-friendly rostering goals forms the second contribution of this thesis.

A review of the call centre tour scheduling and tour assignment literature in chapter 4 indicated that current models do not take account of the goals we

9.1 Thesis Summary and Conclusions

identified, and are somewhat one-dimensionally based on a single objective of minimising staff costs, within a framework of hard constraints which reflect only legal and contractual obligations. Commercially available software packages take some account of employee preferences when assigning tours to employees, but the choice of tour is restricted to those scheduled by a single-objective model. This overall approach supports the low-cost, efficiency driven strategy outlined earlier, confirming our view that there is a need for an employee-friendly rostering model in support of the modern, revenue generation strategy. Such a model must balance multiple objectives relating to staffing costs, customer service, and the needs of the employee.

In chapter 5, we presented our tour scheduling model. This is the first tour scheduling model in the call centre rostering literature to take account of general preferences for patterns of off-duty days, and for similar shift start times within a working week. We use a Mixed Integer Programming (MIP) approach, with the shift pattern preferences represented as soft constraints with weighted penalties in the objective function. The order of preference is generally for full weekends off duty, followed by consecutive-paired off-duties, then non-consecutive off duties around working stints of at least 2 days, and as a last resort, isolated work days. However, the model is flexible enough to allow for an alternative prioritisation (by adjusting the penalty weights) should this be desirable. The second important aspect of the tour scheduling model is the way in which we handle shift start time flexibility within tours. Current MIP models deal with this issue by allowing a “bandwidth” range within which start times can vary. This has two drawbacks. 1) Even a bandwidth of 1 hour would allow 5 different shift start times when working with (as is usual) quarter-hourly staffing intervals. This would be very

unpopular. 2) There is no distinction between forward and backward rotation, so that, for example, preventing a day shift from following a night shift would also prevent a night shift from following a day shift, which is to over-constrain the model. Our approach is to generate an initial set of tours, each with only a single shift start time, but potentially having fewer shifts than is strictly necessary to complete a full week. We then recombine partial tours in order to reduce the number, while maintaining forward rotation. The third aspect of our tour scheduling model is the way in which we handle the scheduling of lunch and coffee breaks, which results in an improved distribution of breaks across the working day. The results show that our model is able to produce higher quality tours in terms of the quantified employee-friendly criteria, while the method for handling start time variations may actually require fewer tours in total. This chapter of the thesis forms the basis of our working paper Glass & Knight (2008a) and constitutes the third contribution of this thesis.

Once a set of weekly tours has been scheduled, the tours must be assigned to the employees. Our tour assignment model is described in chapter 6. We first perform some pre-processing to redistribute the night shifts across the set of tours, in accordance with our rostering goals relating to night shift working. We then assign the tours to the employees using the standard Hungarian method. The cost function of the assignment model represents aspects of continuity between rostering periods, and fairness of distribution of unsociable night and weekend shifts in line with our ergonomic rostering goals. As a post processing step, we attempt to meet any individual requests by exchanging shifts or stints between employees, and finally we assign the lunch and coffee breaks, again taking account of personal requests. The pre- and post-assignment stages mean that the tours

9.1 Thesis Summary and Conclusions

which are eventually worked may be quite different from those originally scheduled, thus indicating an integrated rostering method. Results show that good quality rosters can be produced over a period of several weeks, although there are no benchmark instances or even comparable models to allow subjective evaluation. Ours is a more comprehensive tour assignment process than is currently found in the literature, and thus forms the fourth contribution of this thesis.

The remainder of the thesis addresses nurse rostering. We have identified a number of similarities in the nature of the nurse and call centre agent rostering problems. For example, both deal with the need to cover a varying level of staffing requirement within a 24-hr, 7-day operation, with a mixture of full and part-time employees working a range of contracts, and so on. Moreover, many of the rostering goals that we identified in chapter 3 are to be found in nurse rostering models, as well as the call centre rostering models we have developed in this thesis. Our initial idea of applying our weekly tour scheduling and assignment models directly to the nurse rostering problem was unsuccessful, largely because of week-to-week continuity considerations, and the lack of choice of assignment due to the small number of nurses to be rostered. We have developed a new approach, decomposing the problem into two stages. At the first stage, we roster the nurses to be either on or off duty for each day of the roster, taking account of the constraints relating to stint length, number of days to be worked each week and in total, the incidence of weekend working, and so on. We superimpose night shift stints onto this roster, such that the night shift stints form a particular sequence which repeats each week, and which analysis proved was the only pattern which avoided heavy penalty costs for those benchmark instances which have influenced our formulation. In the second stage of our rostering model, we take the roster from the

first stage and allocate a shift type (e.g. early, day, late) to each of the rostered days which is not an off-duty or a night shift. Our two-stage rostering model found optimal solutions to two benchmark nurse rostering instances, improving upon the best-known solutions at the time. In addition, we have addressed the issue of period-to-period continuity which, although well understood in practice, has not been previously addressed in the nurse rostering literature. Our analysis of the nurse rostering problem which led to the insights relating to night shift sequences, the application of our models to the benchmark instances, and our methodology for handling period-to-period continuity are described in our paper Glass & Knight (2008b), which forms chapter 8. Together with chapter 7, which describes our approach to the nurse rostering problem and gives the mathematical formulations, this constitutes the fifth and final contribution of this thesis.

9.2 Future Directions

The rostering goals that we have identified in this thesis are strongly influenced by legislative requirements and by ergonomic theory. These goals should therefore be applicable to *any* employee rostering domain, and not just call centres. We have described the broad similarities between our own set of rostering goals and those traditionally implemented in nurse rostering models. An interesting idea will be to formally study the similarities and differences between the rostering goals in different domains of employee rostering.

The main gap that we have identified in call centre rostering in it's widest sense, is the need for an extension of the Quality-Driven staffing regime to the smaller call centre. It is our own intention to explore, using simulation models,

the impact of alternative models on customer queuing times and agent occupancy levels in such call centres. Linked to this issue is the shortage of empirical studies which model the relationship between call centre staffing levels and customer satisfaction. There is currently little evidence to support the traditional assumption that customer satisfaction is strongly linked to the length of time spent queuing for answer, and further research is needed in this area.

In this thesis, we have considered only those call centres which provide a single type of service, in other words, we have assumed that the employees are interchangeable in terms of who should work which tour. An important development will be to extend the tour scheduling and tour assignment models to roster multi-skilled agents against staffing requirements for a number of discrete call types. Similarly, our nurse rostering model will need to be extended to handle the rostering of different grades of nurse. Our nurse rostering model assumed that the requirement for night shifts was the same on each day of the rostering period, and alternative patterns of night shift stints may need to be identified and allowed for when this is not the case.

An important future development will be to enhance the various models presented in this thesis to handle a wider range of individual preferences and personal requests, in addition to the ergonomic principles and generalised preferences which form the basis of the rostering system. Such personal requirements may include job sharing, compressed work hours, split-shifts, and so on. The types or range of contract offered to call centre agents may also change in the future, and such changes may require further model developments, particularly if the traditional structure of a fixed number of working days per week changes. There may be a need to take additional rostering requirements into account due to new

HR policies, such as the emerging focus on team working. Our aim is toward a flexible rostering system, supportive of family-friendly employment policies, and thus integrating the functions of HR and Operations management. It is our intention to conduct one or more case studies, working with call centre partners in order to establish “proof-of-concept” of our rostering methodology, to refine the models in the light of practical considerations, and to keep abreast of further developments.

Appendix A

The Brusco-Jacobs Tour Scheduling model

In this Appendix we present a simplified version of the “P2” formulation presented in Brusco & Jacobs (2000). Our aim is to offer an explanation of the P2 model, and thereby assist the reader to compare the functionality and results of P2 and our own model, KG1. To this end, the P2 formulation presented here is geared specifically to the problem instances addressed in chapter 5.

The general conditions regarding shift types and break requirements in P2 are replaced by the specific details of the chapter 5 problem, as follows. Shifts are of 8 hours in length and are contained wholly within a discontinuous 24-hr period, running from 07:00am to 07:00am. The staffing interval is 15-minutes, and shifts can begin in any interval. Thus, there are 96 quarter-hour staffing intervals in a day, and 65 possible shift types. We assume that each shift has 3 scheduled breaks of 1,4,1 intervals in length (respectively), to be scheduled in the windows illustrated in Figure 5.1.

A.1 Shift start times

The Brusco-Jacobs model controls the range of shift start time variations within a tour, using the concept of “bandwidth”. A bandwidth of 1 means that all shifts within a tour must start at the same. At the other extreme, a value of 65 means that all shift start time combinations are valid. In our examples in chapter 5, we used a bandwidth setting of 5. This implies that there are 61 start-time bands, each containing 5 shift types as follows:

Band	Shift Types
1	1,2,3,4,5
2	2,3,4,5,6
3	3,4,5,6,7
	etc
60	60,61,62,63,64
61	61,62,63,64,65

Table A.1: Shift types per band, $b=5$

Bandwidth is implemented in the model by explicitly representing each shift type, within each band, in the 0-1 shift type matrix (so that $a_{jlt} = 1$ if shift type j in band l covers interval t). Thus, with a bandwidth of 5, each shift type can appear up to five times in the shift matrix, since one shift can be present in up to five bands.

A.2 Break Scheduling

The Brusco-Jacobs model has a very different approach to break scheduling to our own, which is based on that of Aykin (1996). In order to aid understanding

of the Brusco-Jacobs model, we now give a brief explanation of their method of *forward pass* and *backward pass* constraints.

We illustrate the forward and backward pass constraints using the lunch break windows depicted in Figure 5.1. Applying the relative break intervals to a shift starting in interval 1 would give an earliest lunch start interval of 13, and a latest start of 17 for that shift type. Hence, we can surmise that the total number of lunches scheduled to start between intervals 13 and 17 on any day, must at least equal the number of shift type 1 scheduled for that day, otherwise there would be insufficient lunches to allocate to the scheduled shifts. This constraint forms the first of a *forward pass* set. The forward pass constraints build cumulatively through the day, as far as interval 81, which is the *latest start time for the lunch break* of shift type 65, which is the latest starting shift. If we define b_t to be the number of lunch breaks scheduled to start in interval t , and x_j as the number of scheduled shifts of type j (on any day), then the forward pass constraints for our particular configuration are as follows:

Lunch Breaks		Shifts
$b_{13} + b_{14} + .. + b_{17}$	$\geq$	x_1
$b_{13} + b_{14} + .. + b_{17} + b_{18}$	$\geq$	$x_1 + x_2$
$b_{13} + b_{14} + .. + b_{17} + b_{18} + b_{19}$	$\geq$	$x_1 + x_2 + x_3$
		etc
$b_{13} + b_{14} + .. + b_{17} + b_{18} + b_{19} + .. + b_{80} + b_{81}$	$\geq$	$x_1 + x_2 + .. + x_{64} + x_{65}$
Forward pass constraints for lunch breaks		

The effect of these constraints will be to ensure that breaks are not scheduled *too late* for feasible assignment. However, it is clear that the forward pass constraints will be satisfied if *all* the breaks are scheduled in interval 13. So a second

A.3 The Mathematical Model

set of constraints are required which complement the forward pass by ensuring that breaks are not scheduled *too early*. This is the *backward pass*, and in our example, is as follows:

Lunch Breaks	Shifts
$b_{77} + \dots + b_{80} + b_{81}$	$\geq x_{65}$
$b_{76} + b_{77} + \dots + b_{80} + b_{81}$	$\geq x_{65} + x_{64}$
$b_{75} + b_{76} + b_{77} + \dots + b_{80} + b_{81}$	$\geq x_{65} + x_{64} + x_{63}$
	etc
$b_{13} + b_{14} + \dots + b_{75} + b_{76} + b_{77} + \dots + b_{80} + b_{81}$	$\geq x_{65} + x_{64} + \dots + x_2 + x_1$
Backward pass constraints for lunch breaks	

Note that the last of the backward pass constraints is the same as the last of the forward pass constraints, and so one of these two is redundant. In fact, the pair of them are replaced in the P2 formulation with a single equality constraint which ensures that the total number of lunches exactly equals the total number of scheduled shifts.

A.3 The Mathematical Model

We now give the details of the Brusco-Jacobs model, formulated for our specific problem. Where possible, we have kept the notation in the formulation below similar to that used in KG1 (as described in section 5.2), in order to facilitate comparison between the two models.

Data

a_{jlt} = 1 if shift type j in start band l covers interval t , 0 otherwise.

c_{ik} = 1 if day i is a working day in pattern k , 0 otherwise.

r_{it} the number of employees required in interval t of day i .

Indices

i	day of week,	$i = 1, 2, \dots, 7$
j	shift type,	$j \in \mathcal{J}$
k	tour pattern,	$k \in \mathcal{K}$
l	start-time band,	$l = 1 \text{ to } 66 - b$
(where b is the bandwidth in intervals)		

Sets

$\mathcal{J}$	the set of all shift types
$\mathcal{J}^l$	the set of shift types in start band l
$\mathcal{K}$	the set of all tour patterns
$\mathcal{L}$	the set of all start-time bands

Decision Variables

- x_{ijl} the number of employees working shift type j on day i ,
in start-time band l .
- z_{kl} the number of employees working pattern k in band l .
- $b1_{it}$ the number of employees starting break 1 in interval t on day i
- $b2_{it}$ the number of employees starting break 2 in interval t on day i
- $b3_{it}$ the number of employees starting break 3 in interval t on day i

A.3.1 Formulation for P2, for specific shift types and breaks

Minimise

$$\sum_{l \in \mathcal{L}} \sum_{k \in \mathcal{K}} z_{kl} \tag{A.1}$$

subject to

$$\begin{aligned} & \sum_{j \in \mathcal{J}} \sum_{l \in \mathcal{L}} x_{ijl} \cdot a_{jlt} - b1_{it} - b3_{it} \\ & -b2_{it} - b2_{it-1} - b2_{it-2} - b2_{it-3} \geq r_{it} \quad \forall i, t \end{aligned} \tag{A.2}$$

$$\sum_{\nu=13}^t b2_{i\nu} - \sum_{l \in \mathcal{L}} \sum_{j \in \mathcal{J}^l: j \leq t-16} x_{ijl} \geq 0 \quad \forall i, \text{ for } t=17, \dots, 81 \quad (\text{A.3})$$

$$\sum_{\nu=t}^{81} b2_{i\nu} - \sum_{l \in \mathcal{L}} \sum_{j \in \mathcal{J}^l: j \geq t-12} x_{ijl} \geq 0 \quad \forall i, \text{ for } t=13, \dots, 77 \quad (\text{A.4})$$

$$\sum_{\nu=6}^t b1_{i\nu} - \sum_{l \in \mathcal{L}} \sum_{j \in \mathcal{J}^l: j \leq t-8} x_{ijl} \geq 0 \quad \forall i, \text{ for } t=9, \dots, 73 \quad (\text{A.5})$$

$$\sum_{\nu=t}^{73} b1_{i\nu} - \sum_{l \in \mathcal{L}} \sum_{j \in \mathcal{J}^l: j \geq t-5} x_{ijl} \geq 0 \quad \forall i, \text{ for } t=6, \dots, 70 \quad (\text{A.6})$$

$$\sum_{\nu=24}^t b3_{i\nu} - \sum_{l \in \mathcal{L}} \sum_{j \in \mathcal{J}^l: j \leq t-26} x_{ijl} \geq 0 \quad \forall i, \text{ for } t=27, \dots, 91 \quad (\text{A.7})$$

$$\sum_{\nu=t}^{91} b3_{i\nu} - \sum_{l \in \mathcal{L}} \sum_{j \in \mathcal{J}^l: j \geq t-23} x_{ijl} \geq 0 \quad \forall i, \text{ for } t=24, \dots, 88 \quad (\text{A.8})$$

$$\sum_{j \in \mathcal{J}} x_{ijl} - \sum_{k \in \mathcal{K}} c_{ik} \cdot z_{kl} = 0 \quad \forall i, l \quad (\text{A.9})$$

$x_{ijl}, z_{kl}, b1_{it}, b2_{it}, b3_{it}$ integer and non-negative.

The objective function minimises the total number of tours scheduled across each days-on pattern k , and each start-time band l . The coverage constraints (A.2) are equivalent to constraint (8) in the P2 formulation in Brusco & Jacobs (2000). These constraints ensure that the total number of agents provided in each interval, less those at break, are at least equal to the number required.

The next pair of constraints (A.3) and (A.4) are, respectively, the forward and backward pass constraints for the lunch breaks, and equivalent to (9) and (10) in P2. The ranges of the intervals and shift types to which each of these constraints applies are explained in section A.2 above. Similarly, constraints (A.5) and (A.6) are the forward and backward pass for break 1, and constraints (A.7) and (A.8) apply to break 3. Since P2 assumes only a single break, P2 contains

no direct equivalents to these two pairs of constraints. In P2, the final constraint of each forward and backward pass pair is replaced by an equality constraint to ensure that the total number of breaks exactly matches the number of scheduled shifts. This forms constraint (11) in P2. The formulation as given here allows the number of breaks to be greater than than are strictly required, subject to coverage constraints (A.2).

Finally, constraint (A.9) ensures that on each day i , for each start-time band l , the sum of the scheduled shifts of each type j , x_{ijl} , is exactly equal to the total number of scheduled tours in band l having day i as a working day. This is equivalent to constraint (12) in P2.

Note that in the actual implementation, the number of decision variables x_{ijl} can be kept to a minimum by defining a limited index for j with respect to the band l , as indicated in table A.1 above.

Appendix B

Tour Reduction for Tour Scheduling examples

In chapter 5, we summarised the results obtained by our tour scheduling model for two call centre instances (referred to as “Telecomms” and “Bank”) and compared them to the results obtained from the Brusco-Jacobs model. We explained that the number of tours initially scheduled (171 and 115 respectively) could be reduced to 169 and 112 by combining partially completed tours into whole tours containing more than one shift type. At the time of submission, we are developing an automated algorithm to perform this function, and so here give a description of the manual process used.

The output from our mathematical program specifies (among other things):

- The number of tours of each on-duty pattern, for each shift start time interval.
- The number of shifts of each type scheduled for each day.

This information constitutes an implicit definition of the set of scheduled tours. Generally, the output of any implicit tour scheduling model requires some post-

B.1 Example 1 - “Telecomms” call centre

processing, and ours is no different in that the construction of the scheduled tours requires the daily shifts to be matched with the weekly off-duty patterns.

Irrespective of any need to reduce the number of tours, we must first carry out this tour construction step. If there are no “spare” shifts present in the roster, the daily shifts and weekly patterns will match precisely. Where the number of daily shifts (of some type) are fewer than number of on-duty days indicated by the tour patterns for that shift type, we allocate spare shifts to make up the difference. Potentially, the tours can subsequently be completed by replacing the spares with any valid shift type, or absences can be allocated. Alternatively, the number of tours can be reduced by re-combining the partial tours. We now illustrate this reduction process for the two examples.

B.1 Example 1 - “Telecomms” call centre

In this example, 171 tours were initially scheduled, containing a total of 841 shifts. As the expectation is for 5 shifts per tour, the implication is that there are 14 spare shifts. After completing the step of assigning shifts to on-duty patterns, 13 tours remained incomplete. Two of these were night shift tours, each containing a single spare. As we handle night shift tours separately at the tour assignment stage (see chapter 6), we leave these two incomplete tours undisturbed for now. This leaves 11 tours with 12 spare shifts, and the aim is to reduce these to 9 tours containing only 2 spares. These 11 tours are shown in Figure B.1, where the numbers in the body of the table are shift types, 0 denotes an off-duty and “sp” denotes a spare shift.

B.1 Example 1 - “Telecomms” call centre

tour	M	Tu	W	Th	F	Sa	Su
1	0	0	5	5	5	5	sp
2	5	0	0	5	5	5	sp
3	5	5	5	0	0	5	sp
4	10	10	10	0	0	10	sp
5	10	10	10	0	0	10	sp
6	sp	10	10	0	0	10	sp
7	12	12	12	0	0	12	sp
8	0	0	13	13	13	13	sp
9	0	0	13	13	13	13	sp
10	0	0	13	13	13	13	sp
11	13	13	13	0	0	13	sp

Figure B.1: Incomplete tours for “Telecomms” (1)

We can observe that it will not immediately be possible to reduce the number of tours to 9, since we have 11 shifts present on Saturday and 10 on Wednesday. In order to overcome this difficulty, we need to include 2 more tours with Saturday off duty, and one more with Wednesday off. These can be selected from the 156 tours which were initially complete. In aiming to keep the number of tours with shift changes as low as possible, we select completed tours of the same type of some of our incomplete tours. Adding 3 such tours, we now have 14 tours as shown in Figure B.2.

tour	M	Tu	W	Th	F	Sa	Su
1	0	0	5	5	5	5	sp
2	5	0	0	5	5	5	sp
3	5	5	5	0	0	5	sp
4	10	10	10	0	0	10	sp
5	10	10	10	0	0	10	sp
6	sp	10	10	0	0	10	sp
7	12	12	12	0	0	12	sp
8	0	0	13	13	13	13	sp
9	0	0	13	13	13	13	sp
10	0	0	13	13	13	13	sp
11	13	13	13	0	0	13	sp
12	10	10	10	10	10	0	0
13	10	10	10	10	10	0	0
14	5	0	0	5	5	5	5

Figure B.2: Incomplete tours for “Telecomms” (2)

B.1 Example 1 - “Telecomms” call centre

These 14 tours can then be reduced to the 12 tours shown in Figure B.3. Quality defects (i.e. shift changes on consecutive working days) are highlighted. The method of removing tours is as follows. First, we treat all the spare shifts as off-duties, so that the incomplete tours contain more off-duty days then are required (that is, two). The idea is then to perform a series of exchanges of one or more days between tours, with the aim of constructing two tours which contain nothing but off-duties, and which can therefore be removed from the solution. Algorithms for performing this type of exercise are widely used in Nurse Rostering, and the discussion of such heuristics in chapter 7 gives some insights into the methods used.

tour	M	Tu	W	Th	F	Sa	Su
1	0	0	5	5	5	5	5
2	5	0	5	5	5	5	0
3	5	10	10	0	5	5	0
4	10	10	10	10	0	5	0
5	10	10	10	10	0	10	0
6	10	10	10	0	10	10	0
7	10	10	10	5	0	10	0
8	12	12	12	0	10	12	0
9	0	13	13	13	13	13	0
10	13	0	13	13	13	13	0
11	0	0	13	13	13	13	0
12	5	5	13	0	0	13	0

Figure B.3: Reduced tours for “Telecomms”

Note that tour 7 contains a backward rotation from shift type 10 (which corresponds to a start time of 9:15 am) to one of 5 (which is 8:00am). This backward rotation can be “repaired” by re-combining tour 7 with a second tour as shown in Figure B.4. The net result of the reduction exercise is shown in Table 5.5.

B.2 Example 2 - “Bank” call centre

Before:

tour	M	Tu	W	Th	F	Sa	Su
7	10	10	10	5	0	10	0
15	5	5	5	0	0	5	5

After:

tour	M	Tu	W	Th	F	Sa	Su
7	5	5	5	5	0	10	0
15	10	10	10	0	0	5	5

Figure B.4: Repairing a backward rotation

B.2 Example 2 - “Bank” call centre

We now give the partial tours from the “Bank” call centre exercise described in section 5.3.3. The initial schedule contains 115 tours and 556 shifts, implying that there are 19 spare shifts in the solution, thus offering the potential to reduce the number of tours to 112, by removing 15 of the spare shifts. After matching the scheduled shifts to the tour patterns, 17 tours contained one or more spare shifts. These tours are shown in Figure B.5, and can be reduced to the 14 tours shown in Figure B.6. The effects of this exercise in terms of roster quality are shown in Table 5.6. Only three tours contained a forward shift rotation on consecutive work days, namely tours 3,10 and 11. No tours contained a backward rotation. All other tours with two shift types had the shift change following a period of 1 or 2 days off duty, which is considered to be perfectly acceptable.

B.2 Example 2 - “Bank” call centre

tour	M	Tu	W	Th	F	Sa	Su
1	4	0	0	4	4	4	sp
2	4	4	4	0	0	4	sp
3	sp	5	5	5	5	0	0
4	0	0	5	5	5	5	sp
5	0	0	5	5	5	5	sp
6	5	5	0	0	5	5	sp
7	7	0	0	7	7	7	sp
8	0	0	8	8	8	8	sp
9	8	8	8	0	0	8	sp
10	9	9	9	0	0	9	sp
11	11	11	11	0	0	11	sp
12	11	11	11	0	0	11	sp
13	0	0	24	24	24	24	sp
14	29	0	0	29	sp	29	29
15	32	32	32	32	sp	0	0
16	32	32	sp	32	sp	0	0
17	0	0	sp	sp	53	53	53

Figure B.5: Incomplete tours for “Bank”

tour	M	Tu	W	Th	F	Sa	Su
1	4	0	4	4	4	4	0
2	9	9	9	0	5	5	0
3	0	5	5	5	5	8	0
4	8	0	5	5	5	5	0
5	0	0	5	5	5	5	0
6	4	4	0	7	7	7	0
7	7	0	8	8	8	8	0
8	11	11	11	0	0	11	0
9	11	11	11	0	0	11	0
10	0	8	8	24	24	24	0
11	29	0	24	29	0	29	29
12	32	32	32	32	0	4	0
13	32	32	0	32	0	9	0
14	5	5	0	0	53	53	53

Figure B.6: Reduced tours for “Bank”

References

- ANTON, J. (1997). *Call Center Management by the Numbers*. Purdue University Press / Call Center Press, Annapolis, MD. 16
- ARTHUR, J. & RAVINDRAN, A. (1981). A multiple objective nurse scheduling model. *IIE Transactions*, **13**, 55–60. 109, 112
- ASHBY, W. (1956). *Introduction to Cybernetics*. Chapman and Hall, Annapolis, MD, ISBN 0-416-68300-2. 74
- AYKIN, T. (1996). Optimal shift scheduling with multiple break windows. *Management Science*, **42**, 591–602. 42, 54, 171
- AYKIN, T. (2000). A comparative evaluation of modelling approaches to the labor shift scheduling problem. *European Journal of Operational Research*, **125**, 381–397. 69
- BAILEY, J. (1985). Integrated days off and shift personnel scheduling. *Computers and Industrial Engineering*, **9**, 395–404. 41, 115
- BAKER, A., ROACH, G., FERGUSON, S. & DAWSON, D. (2003). The impact of different rosters on employee work and non-work time preferences. *Time & Society*, **12**, 315–332. 30

REFERENCES

- BAKER, A., ROACH, G., FERGUSON, S. & DAWSON, D. (2004). Shiftwork experience and the value of time. *Ergonomics*, **43**, 307–317. 34
- BAKER, K. (1976). Workforce allocation in cyclical scheduling problems: A survey. *Operational Research Quarterly*, **27**, 155–167. 41
- BARD, J. & PURNOMO, H. (2005). Preference scheduling for nurses using column generation. *European Journal of Operational Research*, **164**, 510–534. 90, 109, 113
- BARD, J., BINICI, C. & DESILVA, A. (2003). Staff scheduling at the United States Postal Service. *Computers and Operations Research*, **30**, 745–771. 43
- BATT, R. (2002). Managing customer services: Human resource practices, quit rates and sales growth. *Academy of Management Journal*, **45**, 587–597. 18
- BATT, R. & MOYNIHAN, L. (2002). The viability of alternative call centre production models. *Human Resource Management Journal*, **12**, 14–34. 15, 18, 19, 25
- BECHTOLD, S. & JACOBS, L. (1990). Implicit modelling of flexible break assignments in optimal shift scheduling. *Management Science*, **36**, 1339–1351. 42, 43, 54, 69, 116
- BORST, S., MANDELBAUM, A. & REIMAN, M. (2004). Dimensioning large call centers. *Operations Research*, **52**, 17–34. 6, 20
- BRISTOW, G., MUNDAY, M. & GRIPAIS, P. (2000). Call centre growth and location: corporate strategy and the spacial division of labour. *Environment and Planning*, **32**, 519–538. 16

- BRUSCO, M. & JACOBS, L. (2000). Optimal models for meal-break and starting-time flexibility in continuous tour scheduling. *Management Science*, **46**, 1630–1641. 7, 40, 42, 43, 49, 53, 67, 116, 137, 138, 170, 176
- BUFFA, E., M.J.COSGROVE & B.J.LUCE (1976). An integrated shift scheduling system. *Decision Sciences*, **7**, 620–630. 6, 39
- BURKE, E., CAUSEMAECKER, P.D. & BERGHE, G.V. (1999). A hybrid tabu search algorithm for the nurse rostering problem. *Lecture Notes in Artificial Intelligence*, **1585**, 187–194. 101
- BURKE, E., CAUSEMAECKER, P.D., PETROVIC, S. & VANDEN BERGHE, G. (2001a). Fitness evaluation for nurse scheduling problem. In *Proceedings of Congress on Evolutionary Computation, CEC2001*, 1139–1146, IEEE Press. 101, 103
- BURKE, E., COWLING, P., CAUSEMAECKER, P.D. & VANDEN BERGHE, G. (2001b). A memetic approach to the nurse rostering problem. *Applied Intelligence*, **15**, 199–214. 101, 102
- BURKE, E., CAUSEMAECKER, P.D., PETROVIC, S. & VANDEN BERGHE, G. (2002). A multi-criteria meta-heuristic approach to nurse rostering. In *Proceedings of Congress on Evolutionary Computation, CEC2002*, 1197–1202, IEEE Press. 101, 104
- BURKE, E., CAUSEMAECKER, P.D., PETROVIC, S. & VANDEN BERGHE, G. (2003). Variable neighbourhood search for nurse rostering problems. In M.G. Resende & J.P. de Sousa, eds., *Metaheuristics: Computer Decision Making*, chap. 7, 153–172, Kluwer. 101, 104, 107

- BURKE, E., CAUSEMAECKER, P.D. & VANDEN BERGHE, G. (2004a). Novel meta-heuristic approaches to nurse rostering problems in belgian hospitals. In J. Leung, ed., *Handbook of Scheduling: Algorithms, Models and Performance Analysis*, chap. 44, 44.1–44.18, CRC Press. 101, 106
- BURKE, E., CAUSEMAECKER, P.D. & VANDEN BERGHE, G. (2004b). The state of the art of nurse rostering. *Journal of Scheduling*, **7**, 7, 100, 136
- BURKE, E., LI, J. & QU, R. (2006). A hybrid model of integer programming and variable neighbourhood search for highly-constrained nurse rostering problems. technical report nottcs-tr-2007-2. Tech. rep., School of Computer Science and IT, University of Nottingham. 151
- BURKE, E., CURTOIS, T., POST, G., QU, R. & VELTMAN, B. (2008). A hybrid heuristic ordering and variable neighbourhood search for the nurse rostering problem. *European Journal of Operational Research*, **188**, 330–341. 100, 137, 140, 143, 151
- BURNS, R. & CARTER, M. (1985). Work force size and single shift schedules with variable demands. *Management Science*, **31**, 599–607. 41
- CANON, C. (2006). *Personnel scheduling in the Call Center industry*. Ph.D. thesis, Universite Francois-Rabelais de Tours, Laboratoire d’Informatique. 46
- CATTANI, K. & SCHMIDT, G. (2005). The pooling principle. *INFORMS Transactions on Education*, **5**, 3

REFERENCES

- CAUSEMAECKER, P.D. & VANDEN BERGHE, G. (2003). Relaxation of coverage constraints in hospital personnel rostering. *Lecture Notes in Computer Science*, **2740**, 129–147. 101, 105
- CAUSEMAECKER, P.D., DEMEESTER, P., VANDEN BERGHE, G. & VERBEKE, B. (2005). Evaluation of employee rosters with the extended linear numberings method. In *Proceedings of the 24th Workshop of the UK Planning and Scheduling group*, 125–133. 101, 107
- CLEVELAND, B. & MAYBEN, J. (1997). *Call Center Management on Fast Forward*. Call Center Press, Annapolis, MD. 16
- CMINSIGHT (2004). The truth about staff attrition and absenteeism in contact centres. 19
- CONTACTBABEL (2006a). North American contact centres in 2006: The state of the industry. 2
- CONTACTBABEL (2006b). UK contact centres in 2006: The state of the industry. 2
- DANTZIG, G. (1954). A comment on Edie's traffic delay at toll booths. *Operations Research*, **2**, 339–241. 7, 40
- DEERY, S. & KINNIE, N. (2002). Call centres and beyond: A thematic evaluation. *Human Resource Management Journal*, **39**, 471–492. 18
- DEERY, S., IVERSON, R. & WALSH, J. (2002). Work relationships in telephone call centres: Understanding emotional exhaustion and employee withdrawal. *Journal of Management Studies*, **12**, 1–13. 20

REFERENCES

- DOWSLAND, K. & THOMPSON, J. (2000). Solving a nurse scheduling problem with knapsacks, networks and tabu search. *Journal of the Operational Research Society*, **51**, 825–833. 93
- DTI (2004). The UK Contact Centre Industry: A Study. 15, 22
- ERLANG, A. (1948). *The life and works of A.K. Erlang*. The Copenhagen telephone company, Copenhagen. 6, 21
- ERNST, A., JIANG, H., KRISHNAMOORTHY, M. & SIER, D. (2004a). Staff scheduling and rostering: A review of applications, methods, and models. *European Journal of Operations Research*, **153**, 3–27. 5, 7, 39, 136
- ERNST, A., JIANG, H., OWENS, M.K.B. & SIER, D. (2004b). An annotated bibliography of personnel scheduling and rostering. *Annals of Operations Research*, **127**, 21–144. 136
- FEINBERG, R., KIM, I.S. & HOKAMA, L. (2000). Operational determinants of caller satisfaction in the call center. *International Journal of Service Industry Management*, **11**, 131–141. 22
- FOLKARD, S. (1992). Is there a ‘best compromise’ shift system? *Ergonomics*, **35**, 1453–1463. 31, 34
- FOLKARD, S. & TUCKER, P. (2003). Shift work, safety and productivity. *Occupational Medicine - Oxford*, **53**, 95–101. 30
- GANS, N., KOOLE, G. & MANDELBAUM, A. (2003). Telephone call centres: Tutorial, review and research prospects. *Manufacturing and Service Operations Management*, **5**, 79–141. 19, 20

REFERENCES

- GLASS, C.A. & KNIGHT, R.A. (2008a). Call center tour scheduling with employee preferences, City University, Cass Business School *working paper*. 11, 12, 48, 138, 152, 154, 165
- GLASS, C.A. & KNIGHT, R.A. (2008b). The nurse rostering problem: A critical appraisal of the problem structure, City University, Cass Business School *Submitted to European Journal of Operational Research*. 11, 13, 135, 167
- GOODALE, J. & THOMPSON, G. (2004). A comparison of heuristics for assigning individual employees to labour tour schedules. *Annals of Operations Research*, **128**, 47–63. 7, 45
- GROSSMAN, T., OH, S., ROHLEDER, T. & SAMUELSON, D. (2001). *Encyclopedia of Operations Research and Management Science*, chap. Call Centers, 73–76. Kluwer Academic Publishers, Boston, MA, 2nd edn. 7, 39
- GUTHRIE, J. (2001). High involvement work practices: Turnover and productivity, evidence from New Zealand. *Academy of Management Journal*, **44**, 180–190. 18
- HOLMAN, D., BATT, R. & HOLTGREWE, U. (2007). The Global Call Centre Report: International Perspectives on Management and Employment. Tech. rep., Global Call Centre Network. 16
- HOULIHAN, M. (2002). Tensions and variations in call centre management strategies. *Human Resource Management Journal*, **12**, 67–85. 18
- HSE (2006). *Managing shift work: health and safety guidance*. HSE Books, ISBN no. 0717661970. 23, 26, 28, 55, 136

REFERENCES

- HUTCHINSON, S., PURCELL, J. & KINNIE, N. (2000). Evolving high commitment management and the experience of the RAC call centre. *Human Resource Management Journal*, **11**, 63–78. 18
- ISKEN, M. (2004). An implicit tour scheduling model with applications in health-care. *Annals of Operations Research*, **128**, 91–109. 43, 109, 115, 138
- JACOBS, L. & BRUSCO, M. (1996). Overlapping start-time bands in implicit tour scheduling. *Management Science*, **42**, 1247–1259. 42, 43, 52
- JARRAH, A., BARD, J. & DESILVA, A. (1994). Solving large-scale tour scheduling problems. *Management Science*, **40**. 42
- JAUMARD, B., SEMET, F. & VOVOR, T. (1998). A generalized linear programming model for nurse scheduling. *European Journal of Operational Research*, **107**, 1–18. 109, 113
- KINNIE, N., HUTCHINSON, S. & PURCELL, J. (2000). Fun and surveillance: the paradox of high commitment management in call centres. *International Journal of Human Resource Management*, **11**, 967–985. 23
- KNAUTH, P. (1993). The design of shift systems. *Ergonomics*, **36**, 15–28. 30, 31
- KNAUTH, P. (1996). Designing better shift systems. *Applied Ergonomics*, **1**, 39–44. 30, 31
- KOGI, K. (1996). Improving shift workers' health and tolerance to shiftwork: recent advances. *Applied Ergonomics*, **27**(1), 5–8. 34
- KOGI, K. & THURMAN, J. (1993). Trends in approaches to night and shiftwork and new international standards. *Ergonomics*, **36**, 3–13. 32

- KOOLE, G. (2005). Call centre mathematics: A scientific method for understanding and improving contact centres, dept. of Mathematics, Vrije University, Amsterdam. 3
- KUHN, H.W. (1955). The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, **2**, 7, 12, 78
- KUNDI, M. (2003). Ergonomic criteria for the evaluation of shift schedules. *Theoretical Issues in Ergonomic Science*, **4**, 302–318. 30, 136
- MEHROTRA, V. (1997). Ringing up big business. *OR MS Today*, **24**, 18–25. 17
- MILLAR, H. & KIRAGU, M. (1998). Cyclic and non-cyclic scheduling of 12 h shift nurses by network programming. *European Journal of Operational Research*, **104**, 582–592. 110, 116
- MILLER, H., PIERSKALLA, W. & RATH, G. (1976). Nurse scheduling using mathematical programming. *Operations Research*, **24**, 857–870. 109, 111, 113
- MIOZZO, M. & RAMIREZ, M. (2003). Service innovation and the transformation of work: the case of UK telecommunications. *New Technology, Work and Employment*, **18**, 62–79. 15, 17
- MORRIS, J. & SHOWALTER, M. (1983). Simple approaches to shift, days-off, and tour scheduling. *Management Science*, **29**, 942–950. 40
- NACHREINER, F., QIN, L., GRZECH, H. & HEDDEN, I. (1993). Computer aided design of shift schedules. *Ergonomics*, **36**, 77–83. 30
- OZKARAHAN, I. (1991). A disaggregation model of a flexible nurse scheduling support system. *Socio-Economic Planning Sciences*, **25**, 9–26. 115

REFERENCES

- OZKARAHAN, I. & BAILEY, J. (1988). Goal programming model subsystem of a flexible nurse scheduling support system. *IIE Transactions*, **20**, 306–316. 109, 115
- PINE, B. (1999). *Mass Customization: The new frontier in business competition*. Harvard Business School Press, 1st edn. 17
- SCHOENFELDER, E. & KNAUTH, P. (1993). A procedure to asses shift systems based on ergonomic criteria. *Ergonomics*, **1**, 65–76. 30, 31
- SILVESTRO, R. & SILVESTRO, C. (2008). Towards a model of strategic roster planning and control: an empirical study of nurse rostering practices in the UK National Health Service. *Health Services Management Research*, **21**, 93–105. 91, 160
- TAYLOR, P. & BAIN, P. (1999). An assembly line in the head: work and employee relations in the call centre. *Industrial Relations Journal*, **30**, 101–117. 2, 14
- THOMPSON, G. (1995a). Improved implicit modelling of the labor shift scheduling problem. *Management Science*, **41**, 595–607. 42
- THOMPSON, G. (1995b). Labor scheduling using npv estimates of the marginal benefit of additional labor capacity. *Journal of Operations Management*, **13**, 67–86. 7, 39
- THOMPSON, G. (1997a). Assigning telephone operators to shifts at New Brunswick telephone company. *Interfaces*, **27**, 1–11. 7, 45, 47

REFERENCES

- THOMPSON, G. (1997b). Labor staffing and scheduling models for controlling service levels. *Naval Research Logistics*, **44**, 719–740. 46
- TOPALOGLU, S. & OZKARAHAN, I. (2004). An implicit goal programming model for the tour scheduling problem considering the employee work preferences. *Annals of Operations Research*, **128**, 135–158. 109, 114
- TOTTERDELL, P. (2005). Work schedules. In J. Barling, E. Kelloway & M. Frome, eds., *Handbook of Work Stress*, chap. 3, 35–62, Sage Publications. 22, 75, 136
- WALLACE, C. & EAGLESON, G. (2000). The sacrificial HR strategy in call centers. *International Journal of Service Industry Management*, **11**, 174–184. 16
- WARNER, D. (1976). Scheduling nursing personnel according to nurse preference: A mathematical programming approach. *Operations Research*, **24**, 842–856. 109, 110, 113
- WARNER, D., KELLER, B. & MARTEL, S. (1991). Automated nurse scheduling. *Journal of the society for health systems*, **2**, 66–80. 112
- WEDDERBURN, A. (1981). Is there a pattern in the value of time off work? In A. Reinberg, N. Vieux & P. Andlauer, eds., *Night and Shiftwork: Biological and Social Aspects*, 495–504, Oxford: Pergamon Press. 30