CAPSS: Computer-assisted Patient Scheduling System

Mike Hughes [mlwh@city.ac.uk]
Centre for Measurement and Information in Medicine,
City University
Northampton Square
London EC1V 0HB.

Submitted in partial fulfilment of a PhD in Measurement and Information in Medicine, 2003

Abstract

The healthcare industry, and particularly the publicly funded healthcare industry, faces many challenges over the foreseeable future. These challenges come from a variety of sources such as trends in public opinion, advances in medical research and the changing healthcare demands of the patient population.

The publicly-funded healthcare industry has not kept pace with these changes, as evidenced by large waiting lists for many surgical procedures. If current standards of quality of care are to be maintained, and increased wherever possible, and healthcare budgets not to spiral upwards, then the only solution to the waiting list problem is to increase the cost-effectiveness of healthcare provision.

It will be hypothesised that the cost-effectiveness of healthcare delivery can be improved through the development of a computer-assisted patient scheduling system (CAPSS). This hypothesis will be supported by showing that providing healthcare managers with more complete and accurate information about the projected availability and demand for healthcare resources improves the ability to control the operational performance of the healthcare system. Moreover, that the ability to deliver this necessary information to the control system in a timely and efficient manner is only realistically attainable through the computerisation of the patient scheduling system, and hence the deployment of CAPSS.

The demonstration of the viability of a computerised model of patient scheduling is performed using the empirical domain of the Royal Brompton and Harefield NHS Trust (RBH). Using this data various models are developed, ranging from a mathematical model demonstrating the relationship between the degree of control over patient scheduling attainable by healthcare managers and the optimal level of cost-effectiveness thereby achievable, through to design models of computer simulation programs that may be used as the basis of a decision-support system for enhancing the processes of patient scheduling and the corresponding allocation of healthcare resources.

Acknowledgements

The Author wishes to acknowledge the grateful contributions of the following in the preparation of this thesis and the research described therein:

- Paul Butler, IT Manager, Royal Brompton and Harefield NHS Trust
- Ewart Carson, Director, Centre for Measurement and Information in Medicine, City University and project supervisor
- Michael Hutt, Director, Clinical Support Services, Royal Brompton and Harefield NHS Trust
- Surjeet Kaur, Nursing Manager, Adult Intensive Care Unit, Royal Brompton and Harefield NHS Trust
- Mahmoud Makhlouf, Senior Engineer, MITRE Corporation and project technical supervisor
- Cliff Morgan, Clinical Director, Adult Intensive Care Unit, Royal Brompton and Harefield NHS Trust and project clinical supervisor
- Amy Page, Patient Admissions Manager, Department of Surgery, Royal Brompton and Harefield NHS Trust
- Margo Shaffer, Directorate Manager, Cardiac Services, Royal Brompton and Harefield NHS Trust
- Paul Silvester, Patient Admissions Manager, Department of Surgery, Royal Brompton and Harefield NHS Trust
- Ron Summers, Professor, Department of Information Systems and Library Studies, Loughborough University and former project supervisor

CONTENTS

Acknowledgements 3
CONTENTS 4
1. Introduction 8
1.1. Background 8
1.2. Objectives 10
1.2.1. Theoretical Analysis 10
1.2.2. Empirical Analysis 11
1.2.3. Modelling Approach and Formalism 13
1.2.4. Operational Modelling 14
1.2.5. Design Modelling 17
1.3. Clinical Setting 18
1.4. Summary 18
2. Healthcare Resource Allocation 21
2.1. Modelling Healthcare Resources 23
2.1.1. Proactive and Reactive Resource Allocation Decisions 23
2.1.2. The Bed-slot Assumption 25
2.2. Effectiveness of Resource Allocation 28
2.2.1. The parametric assumption 28
2.2.2. Resource allocation as a control process 30
2.2.3. Measuring the Effectiveness of Healthcare Resource Allocation 33
2.3. Cost-effectiveness of Healthcare Delivery 34
2.3.1. Measuring the Cost-effectiveness of Healthcare Delivery 34
2.3.2. Modelling Not Generic and Fixed Resources 37
2.4. The Real World 38
2.4.1. The Real Patient Population 39
2.4.2. The Real Healthcare Facility 42
2.5. Resource Allocation and Cost-effectiveness in a Progressive Care System 44
2.5.1. Classification of Control-limiting Factors 48
2.6. Summary 51
3. The Empirical Domain 53
3.1. Operational Overview 55
3.1.1. High-Dependency Care 55
3.2. RBH Overview 58
3.2.1. RBH High-Dependency Information Systems 60
3.2.2. Patient Management 62
3.2.3. Carevue 64
3.2.4. The RBH Operating Room Suite (OR) 68
3.2.5. The Recovery Room (RR) 75

3.2.6. The Intensive-Care Units (AICU/PICU) 80
3.2.7. The Intermediate-Care Units (HDUs) 89
3.3. Resource Allocation 91
4. Modelling Approach and Formalism 101
4.1. Introduction 101
4.2. Modelling Approach 102
4.3. Modelling Formalism 112
4.3.1. Object-Oriented Petri-Nets 120
4.3.2. Process Diagrams 121
4.3.3. Object Class Diagrams 126
4.3.4. Object Relationship Diagrams 127
4.3.5. System Metrics 129
4.4. Summary 130
5. Operational Models 132
5.1. RBH Current Operational Model (COP) 133
5.1.1. Introduction 133
5.1.2. Object Classes 133
5.1.3. Databases 135
5.1.4. Processes 137
5.1.5. Process Diagrams 141
5.2. RBH Proposed Operational Model (POP) 146
5.2.1. Introduction 146
5.2.2. Objects 146
5.2.3. Databases 149
5.2.4. Processes 150
5.2.5. Process Diagrams 154
5.3. Discussion 160
5.3.1. Model Comparison 160
5.3.2. The Design Models 163
6. Design Models 165
6.1. 2.07: predictResources 165
6.1.1. Literature Review 166
6.1.2. Univariate Analyses 173
6.1.3. Multivariate Analyses 177
6.1.4. Conclusions 180
6.1.5. Discussion 183
6.1.6. 2.08: evaluateSchedule 187
6.1.7. Literature Review 195
6.1.8. Petri Nets 196
6.1.9. Coloured-Timed Petri Nets (CTPNs) 197
6.1.10. The Model 199

6.1.11. Discussion 203
6.2. Summary 204
7. Concluding Remarks 206
7.1. Modelling Healthcare Resource Allocation 206
7.2. The Empirical Domain 209
7.3. Modelling Approach and Formalism 217
7.4. Operational Models 220
7.5. Design Models 222
7.6. Summary 226
7.6.1. Theoretical Analysis 226
7.6.2. Empirical Analysis 227
7.6.3. Modelling Approach and Formalism 228
7.6.4. Operational Modelling 229
7.6.5. Design Modelling 231
7.7. Contributions to Knowledge 232
7.8. Future Research 234
8. References 235
9. Appendix: Intensive-Care Units 240
9.1.1. The Adult Intensive Care Unit (AICU) 240
9.1.2. The Paediatric Intensive Care Unit (PICU) 245
10. Appendix: Evaluating Resource Targeting in an Adult Intensive Care Unit 247
10.1. Data 249
10.2. The Hypotheses 251
10.2.1. The Predictability Hypothesis 251
10.2.2. The –G&F Proportion Hypothesis 253
10.2.3. The Variance Hypothesis 257
10.2.4. The Difference Hypothesis 263
10.3. General Discussion 268
11. Appendix: Operating Theatre Utilisation Analysis 271
11.1. Introduction 271
11.2. Method 271
11.3. Data 278
11.4. Scheduling Status: Non-monotonic bed-slot allocation 281
11.5. Scheduling Status: Monotonic bed-slot allocation 283
11.6. Results 285
11.6.1. Scheduling Status: Non-monotonic bed-slot allocation 285
11.6.2. Scheduling Status: Monotonic bed-slot allocation 286
11.6.3. Scheduling Status: Monotonic and non-monotonic bed-slot allocation 286

11.6.4. Identification and Classification of Control-Limiting Factors 288
11.7. Discussion 290
12. Operational Model Data Attributes 292
12.1. Current Operational Model Data Attributes 292
12.2. Proposed Operational Model Data Attributes 298
13. Appendix: Related Publications 307

1. Introduction

1.1. Background

The publicly-funded healthcare industry faces many challenges over the coming years, not least of which is to serve an ever-increasing demand for its resources, while at the same time increasing efficiency, keeping costs under control, and reducing waiting lists.

In the United Kingdom, the government has made a commitment to reduce waiting lists, with a large proportion of the funding necessary to accomplish this goal being expected to come from efficiency savings as well as de novo budgetary increases [NHS98], [DOH01], [DOH02]. At the same time, there are demographic changes in the developed world which increase levels of per capita healthcare expenditure necessary to maintain current levels of healthcare provision. There are also escalating costs involved in the development of new therapies, particularly drug therapies, which are inevitably passed on to healthcare consumers through higher insurance premiums or greater tax burdens.

The problems facing the publicly-funded healthcare industry can be seen through the growing size of waiting lists for NHS treatment. Figure 1.01 below shows the growth in the number of people waiting for NHS treatment in England for the period 1996 to 1998 [DOH98]. In 1996 just over one million people were waiting for treatment, while in 1998 the figure grew to nearly 1.3 million[1].

If the twin objectives of reducing waiting lists and not allowing healthcare budgets to spiral out of control, the only solution is to increase the cost-effectiveness of healthcare provision. The term 'cost-effectiveness' may be summarised as being able to treat more patients for the same amount of resource inputs, while not decreasing the quality of healthcare delivered.

---

[1] As of the date of writing, the figures for the numbers waiting for consultation or treatment under the NHS remains stubbornly high, although there has been some reduction from the 1998 figures.

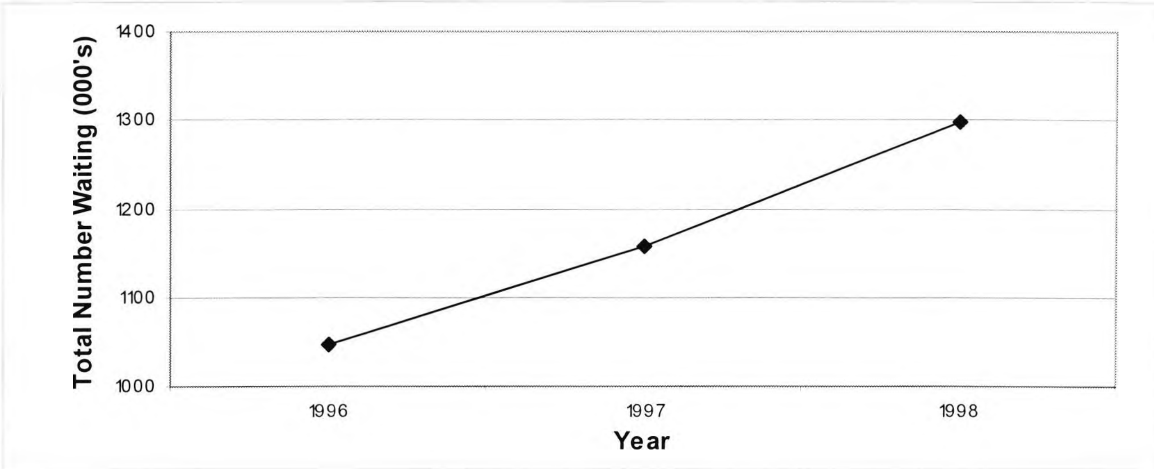

Figure 1.01. Total Number of Patients Waiting for NHS Treatment (England) 1996-1998 (Source: UK Department of Health)

A recent report [NHS98a] to The Department of Health identified three elements to the measurement of cost-effectiveness in healthcare. The first of these is the cost per unit of care/outcome. That is, how much it costs in financial terms to treat a typical case of, for example, cataract. The second measure is the productivity of 'capital estate'. This is a less encompassing measure than the first, although is nevertheless important as the purchase of new capital equipment and the depreciation of existing equipment make significant contributions to healthcare budgets. Thus, to increase the productivity of an important piece of capital equipment, such as a mechanical ventilator, cost-effectiveness can be increased by decreasing the number of mechanical ventilators necessary to treat the same amount of patients. The third and final measure proposed in the report is the productivity of labour. This is a more difficult measurement to make than the productivity of capital equipment, although is more important in terms of its share of the healthcare budget.

Both the productivity of labour and the productivity of capital equipment are less general measures of cost-effectiveness than the cost per unit of care/outcome insofar that increases in the productivity of labour or capital equipment should favourably impact on the cost per unit of care/outcome, whereas the converse is not necessarily the case. Cost per unit of care/outcome is also the only one of the three measures to explicitly measure cost-effectiveness relative to a standardised level of quality of care, expressed by the terms "unit of care/outcome".

There are various approaches which have been proposed in attempting to increase the productivity of resources such as capital equipment and labour. These approaches can broadly be categorised in terms of whether they involve the introduction of new technology or whether they involve making better use of existing technology, or a combination of making better use of existing technology through the introduction of new technology. The development of new information systems is typically a combination approach, where a new technology is introduced in the form of a new information system which, ideally, allows managers to make better use of resources and clinicians to provide better care to patients.

The UK government has identified improved information systems as a way to increase the cost-effectiveness of healthcare delivery. A recent report delivered to The Department of Health outlining a strategy for modernising the NHS identified the management of information as one way to improve the cost-effectiveness of healthcare delivery "by providing health planners and managers with the information they need" [NHS98]. An examination of the nature of this informational deficit as it occurs within the context of the operational management of hospital resources will be the main objective of this dissertation, along with the development of an information system aimed at reducing it.

It will be hypothesised that the cost-effectiveness of healthcare delivery can be improved through the development of a computer-assisted patient scheduling system (CAPSS). This hypothesis will be supported by making the claim that providing healthcare managers responsible for patient scheduling with more complete and accurate information about the projected availability and demand for healthcare resources improves the ability to control the performance of the healthcare system. Moreover, that the ability to deliver this necessary information in a timely and efficient manner is only realistically attainable through the computerisation of the patient scheduling system, and hence the deployment of CAPSS.

1.2. Objectives

To support the hypothesis that CAPSS may improve the operational cost-effectiveness of healthcare delivery the main body of this thesis will be divided into five chapters, each of which will satisfy a particular area in the modelling and development of CAPSS within an empirical domain.

1.2.1. Theoretical Analysis

Chapter 2 will provide a theoretical analysis of the process of patient scheduling. Chapter 2 seeks to satisfy the following objectives in supporting the hypothesis that CAPSS may improve the operational cost-effectiveness of healthcare delivery.

PATIENT SCHEDULING IS A CONTROL PROCESS

Chapter 2 will present a mathematical model that demonstrates the equivalence of patient scheduling as a control process. This will be the first logical step in supporting the hypothesis that CAPSS may improve operational cost-effectiveness by increasing the level of control that operational managers have over the patient scheduling process.

EFFECTIVENESS OF CONTROL MAY BE MEASURED BY THE CONSUMPTION DISTRIBUTION

It will be argued that the extent to which operational managers can control the patient scheduling process may be measured by the variance in the distribution of different consumption rates of healthcare facilities' resources. This represents an important starting point in developing a method for evaluating system performance.

COST-EFFECTIVENESS OF HEALTHCARE DELIVERY MAY BE DEFINED MATHEMATICALLY

An equation will be developed based on the distribution in consumption rates of healthcare resources that measures the cost-effectiveness of healthcare delivery. This equation provides the basis for a far more formal approach to healthcare facility performance evaluation than has been used thus far.

INCREASING THE EFFECTIVENESS OF CONTROL INCREASES THE COST-EFFECTIVENESS OF HEALTHCARE DELIVERY

Demonstrating a positive correlation between the effectiveness of control and the operational cost-effectiveness of a healthcare facility is necessary in supporting the hypothesis that CAPSS can improve cost-effectiveness with this improvement in cost-effectiveness being derived from an increase in the effectiveness of control over patient scheduling.

THE PROCESS OF RESOURCE ALLOCATION CAN BE CONSIDERED AS A MULTI-DIMENSIONAL TILING PROBLEM

The difficulty in optimising the clinical and economic performance of a healthcare facility is demonstrated through the equating of the patient scheduling process to that of a tiling problem. This equivalence is presented as a conceptual aid in the modelling of the patient scheduling process.

1.2.2. Empirical Analysis

Chapter 3 will provide an empirical analysis of the process of patient scheduling as it occurs in the Royal Brompton and Harefield NHS Trust (RBH) which will be the empirical domain used throughout this thesis. Chapter 3 seeks to satisfy the following objectives in supporting the hypothesis that CAPSS may improve the operational cost-effectiveness of healthcare delivery:

THE BED-SLOT ASSUMPTION IS VALIDATED FOR THE RBH HIGH-DEPENDENCY ENVIRONMENT

A fundamental assumption is introduced in the modelling of healthcare resource allocation stating that under certain conditions healthcare resources may be modelled as a single resource – the bed-slot. Chapter 3 summarises a study testing the bed-slot assumption as it applies to RBH, with the main results of the study included as an Appendix.

PATIENTS RESOURCE CONSUMPTION MAY BE PREDICTED WITH ACCEPTABLE ACCURACY USING A ROLLING PREDICTION MODEL

A requirement of any system with the objective of increasing the level of control that managers have over patient scheduling must be to be able to predict the amount of resources that each patient is likely to consume. This requirement is tested for RBH data in a study which is summarised in Chapter 3, with the main results included as an Appendix.

THE RESOURCE CONSUMPTION PROFILE OF RBH INTENSIVE-CARE PATIENTS IS TYPICAL

To demonstrate the applicability of the results of the theoretical analysis from Chapter 2 to the empirical domain represented by RBH, the resource consumption profile of RBH is presented and comparison made to the theoretical distribution presented in Chapter 2.

THERE IS A RELATIONSHIP BETWEEN PATIENTS RESOURCE CONSUMPTION AND MORTALITY, OPERATIVE CATEGORY AND AGE

The AICU Chronicity Study, included as an appendix and summarised in Chapter 3, examines the relationship between patients' length of stay in the RBH Adult Intensive Care Unit (AICU) and various patient variables, including mortality and patient's operative category and age. The results of this study indicate the viability of a prediction model for predicting patients' resource consumption.

A STATISTICAL METHOD MAY BE USED TO IDENTIFY AND QUANTIFY CONTROL-LIMITING FACTORS IN THE RBH PATIENT SCHEDULING SYSTEM

The success of CAPSS in increasing the cost-effectiveness of healthcare delivery at RBH is dependent upon the current process of patient scheduling being sub-optimal. Moreover, that any sub-optimality must be caused by control-limiting factors that result from a lack of knowledge regarding patients' projected resource consumption or the projected availability of healthcare resources if the introduction of CAPSS is to support the hypothesis that it is able to improve cost-effectiveness. This is opposed to control-limiting factors that are known by managers and are inherent in the system design. This distinction is defined as one between epistemological and non-epistemological control-limiting factors and a statistical method is developed that may identify these two types of control-limiting factor and quantify their effects on cost-effectiveness. The full method is included as an Appendix.

THERE ARE EPISTEMOLOGICAL CONTROL-LIMITING FACTORS PRESENT IN THE RBH PATIENT SCHEDULING PROCESS

Using the statistical method to identify and quantify control-limiting factors in the RBH patient scheduling system, a study is presented with the objective of demonstrating that there are epistemological control-limiting factors present. The presence of epistemological control-limiting factors would thereby indicate the sub-optimal performance of the system and moreover the capability of CAPSS being able to support the hypothesis that its introduction would improve cost-effectiveness.

1.2.3. Modelling Approach and Formalism

Chapter 4 proposes a modelling approach and formalism suitable for the development of operational models of RBH and the requirements of CAPSS. In the development of a suitable modelling approach and formalism Chapter 4 satisfies the following objectives:

BPR AND SOFTWARE ENGINEERING MODELLING APPROACHES ARE NOT INDIVIDUALLY ADEQUATE FOR MODELLING CAPSS

Chapter 4 evaluates the utility of modelling approaches and formalisms adopted in Business Process Re-engineering (BPR) and Software Engineering to be able to offer a suitable basis for the modelling of RBH and the requirements of CAPSS. It is concluded that neither formalism alone is adequate for the task. This conclusion is made on the basis of considering various properties necessary for the modelling of RBH and the requirements of CAPSS and whether or not those properties are included in conventional BPR and Software Engineering approaches or formalisms.

A HYBRID MODELLING APPROACH MAY SATISFY THE MODELLING NEEDS OF AN OPERATIONAL MODEL OF CAPSS

As a result of the evaluation of modelling approaches and formalisms deployed for the activities of software engineering and business process re-engineering an argument is made for a hybrid modelling approach that is particularly suited to the development of operational models of the RBH patient scheduling system and other similar systems.

THE REPRESENTATIONAL EFFICIENCY OF OBJECT-ORIENTED AND FUNCTION-ORIENTED MODELLING FORMALISMS MAY BE MEASURED

In evaluating the appropriate modelling formalism to adopt for the modelling of the RBH patient scheduling system the object-oriented modelling paradigm is compared against its function-oriented equivalent in terms of the efficiency with which each could represent and model systems. The comparison method used is based on information theory and the results indicated that the most efficient modelling paradigm to use was dependent on the degree of system complexity. In the case of modelling the RBH patient scheduling system, it will be concluded that due to the level of system complexity an object-oriented approach was appropriate.

AN OBJECT-ORIENTED VERSION OF PETRI-NETS PROVIDES AN ADEQUATE FORMALISM FOR THE MODELLING OF CAPSS

Although Petri-nets have been used extensively in system modelling, their use has been largely restricted to the design of system simulation models, rather than the specification of operational dynamics and software requirements as in the case of modelling RBH and the requirements of CAPSS. In Chapter 4 the properties of a basic definition of Petri-nets is combined with the static modelling properties of a generic object-oriented modelling formalism to create a comprehensive modelling formalism. It is argued that this formalism may be used in the modelling of RBH and the requirements of CAPSS.

SYSTEM METRICS MAY BE CALCULATED FOR OPERATIONAL MODELS THAT EVALUATE THE OPERATIONAL PERFORMANCE OF THE RESULTING SYSTEM

In order to provide a formal comparison of the current process of patient scheduling at RBH and the proposed process with the introduction of CAPSS, Chapter 4 includes the definition of a set of system metrics that may be used to evaluate the processing efficiency and degree of system integration in both the current and proposed operational models of the patient scheduling process.

1.2.4. Operational Modelling

Chapter 5 is the first modelling chapter whose purpose is the development of two models that represent the operational processes and data involved in the process of patient scheduling at RBH. The first model to be developed is the current operational model (COP) that represents the current patient scheduling process. The second model to be developed is the proposed operational model (POP) that represents the patient scheduling process with the inclusion of CAPSS. The main points of the two models are as follows:

THERE ARE THREE MAIN OBJECT CLASSES IN THE CURRENT AND PROPOSED OPERATIONAL MODELS OF THE RBH PATIENT SCHEDULING PROCESS

Three object-classes are used in both the current and proposed operational models:

1. Patient contains all of the data and processes defining patients within the RBH high-dependency environment
2. Bed-Slot contains all of the data and processes defining bed-slots within the RBH high-dependency environment
3. Unit contains all of the data and processes defining the component healthcare units of the RBH high-dependency environment

CAPSS INTRODUCES TWO NEW PROCESSES INTO THE RBH RESOURCE ALLOCAITON AND PATIENT SCHEDULING PROCESS

The main difference between the current and proposed operational models is the introduction of two new processes into the proposed operational model that are absent from the current operational model. Both of these new processes are represented as computerised processes and as such represent the main software components of CAPSS. The two new processes define the prediction of patients' resource consumption and the subsequent prediction of system performance.

In the proposed operational model predictions are made of each patient requiring admission to the RBH high-dependency environment. These predictions consist of the patient's projected bed-slot consumption and which of the component units of the RBH high-dependency environment where the bed-slots will be consumed and when. The manager responsible for scheduling patient admission may then propose a schedule of admissions. This proposed schedule of admissions is evaluated by the system by predicting the performance of each component unit of the RBH high-dependency environment that would result given each patient's predicted resource consumption and the availability of resources within the high-dependency environment. The operational manager is then able to modify the schedule of admissions and repeat the evaluation procedure, or to implement the schedule unmodified.

Because there is no closed-loop involved in the process of patient scheduling in so far that the decision-making is still undertaken by a human agent, CAPSS represents a decision-support system. Unlike other decision-support systems deployed in healthcare settings, however, CAPSS is not concerned with the diagnosis or treatment processes.

THE PROPOSED OPERATIONAL MODEL INCREASES THE DEGREE OF SYSTEM AUTOMATION

If the introduction of CAPSS is to improve control over patient scheduling at RBH then it must demonstrate measurable effects in control system performance. Even if this is achieved through CAPSS generating sufficiently accurate projections of patient resource consumption and predictions of system performance, this represents a necessary, but not a sufficient, condition for the confirmation of the hypothesis that such an improvement in control system performance results in an increase in the cost-effectiveness of healthcare delivery.

For CAPSS to improve control system performance it must not only result in sufficiently accurate predictions of patient resource consumption and system performance evaluation, but it must also accomplish these objectives in a cost-effective manner. For example, it must generate the information necessary to improve control system performance using as few resources as possible while not compromising on the accuracy of this information. In practical terms this implies a need for increased levels of system automation.

One of the system metrics developed in Chapter 4 is designed specifically to measure the degree of system automation and this metric is applied to both the current and proposed operational models of the RBH patient scheduling system to compare the degree of system automation implied by each model.

THE PROPOSED OPERATIONAL MODEL INCREASES THE DEGREE OF SYSTEM INTEGRATION

While system automation measures the extent to which system processes are performed by non-human agents, system integration measures the extent to which system processes are distributed between a number of different processing agents – human or non-human.

System integration is an important consideration when developing and subsequently maintaining the system – the more that system processes are distributed amongst a proliferation of different processors, then the more expensive it will be to develop and maintain. For example, in the case of integration of different databases, supporting data conversion and communication software needs to be developed.

One of the system metrics developed in Chapter 5 is designed to measure an important aspect of system integration by measuring the degree of data distribution between different data stores. This metric is applied to both the current and proposed operational models of the RBH patient scheduling system to compare the degree of system integration implied by each model.

1.2.5. Design Modelling

Chapter 6 is the second modelling chapter whose purpose is the development of two models that are proposed as implementations of the two new processes introduced in the proposed operational model of the RBH patient scheduling process of Chapter 5. The first of these models is for the prediction of patient resource consumption requirements. The second of these models is for the evaluation of system performance based on the predictions generated by the first model. The main points of the two models are as follows.

PATIENTS RESOURCE CONSUMPTION MAY BE PREDICTED BY PREDICTION MODELS

The objective of the first model is to demonstrate that a suite of computerised prediction models is capable of predicting patients' resource consumption – within an acceptable degree of accuracy – for patients scheduled for admission to the RBH high-dependency environment.

The method used to accomplish the objective will be to evaluate and compare prediction models that have been developed and described in the literature through a systematic literature review. Both the methods used in the derivation of each model, as well as the accuracy of the predictions made by each model will be considered in evaluating and comparing the models.

If any of the models is capable of demonstrating a sufficient level of accuracy in its predictions, then it will be assumed that it has the capability of being implemented as a software routine in the context of CAPSS for the prediction of patients' resource consumption requirements and the integration with a model for the evaluation of system performance.

SIMULATION IS THE BEST APPROACH TO PREDICT THE OPERATIONAL PERFORMANCE OF HEALTHCARE SYSTEMS

The objective of the second model is the evaluation of the operational performance of healthcare systems such as the RBH high-dependency environment. This evaluation is based on the output of the first model for the prediction of patient resource consumption requirements.

Various methods will be compared for the design of the model based on the methods that have been successfully used for comparable models and have been described in the literature. Although none of the models that will be reviewed from the literature have precisely the same objectives as the model to be developed, a set of evaluation parameters may be extrapolated on the basis of the requirements of the model to be developed and applied to the models that have been presented in the literature. On the basis of this evaluation, it will be concluded that the model is best developed using a simulation of the RBH high-dependency environment.

COLOURED-TIMED PETRI-NETS MAY BE USED TO SIMULATE THE OPERATIONAL PERFORMANCE OF THE RBH HIGH-DEPENDENCY UNIT

The use of simulation as the basis for model development still leaves open the formalism to be used for the design of the simulation model. It will be argued that an abbreviation of the basic Petri-net formalism used in the development of the operational models whereby different resource and patient types are categorised according to colours is well-suited to the task. This argument will be based on the successful use of coloured Petri-nets in other comparable modelling domains

1.3. Clinical Setting

The healthcare facility in which this project is undertaken is the high-dependency environment of the Royal Brompton and Harefield NHS Trust, London (RBH). RBH provides an excellent setting for this project as it exemplifies many of the challenges facing publicly funded healthcare in their most extreme form. The drive for the cost-effective delivery of high-dependency healthcare is intensified both by the high-cost of providing such healthcare, as well as the increasing need for its provision, the complex difficulties involved in controlling the allocation of high-dependency healthcare resources, and it being the focal point for advances in healthcare and the introduction of new, and often expensive, technologies. All of these considerations are especially pertinent to RBH, given its position as a specialist teaching hospital and tertiary referral centre, and thus accepting a case-mix whose healthcare requirements pose a particular challenge to healthcare managers and clinicians alike.

1.4. Summary

The healthcare industry is undergoing a period of upheaval with increasing demands on healthcare resources against a backdrop of rapid technological development. The result is a need for increased cost-effectiveness of healthcare delivery. The aim of this thesis will be to show how CAPSS, a computer-assisted patient scheduling system, may improve the cost-effectiveness of healthcare delivery through enabling healthcare managers to exert a greater degree of control over the process of patient scheduling.

The demonstration of the viability of CAPSS will be performed using the empirical domain of the Royal Brompton and Harefield NHS Trust (RBH). The first step is to develop a mathematical model which highlights the positive relationship between the cost-effectiveness of healthcare delivery and the degree of control that managers are able to exert over patient scheduling.

A detailed examination of the current process of patient scheduling as it occurs at RBH is the next step, which also involves a series of studies that justify the deployment of a computer-assisted patient scheduling system at RBH.

Following the derivation of a modelling approach and formalism designed for the development of software requirements and process-re-engineering models, two complementary models of the patient scheduling process at RBH will be developed in Chapter 5. The first of these models used the modelling approach and formalism derived in the previous chapter to model the current process of patient scheduling; the second modelled the process as it would be with the introduction of two new processes implemented as software routines, in addition to supporting database and data collection systems.

Apart from the introduction of the two new processes into the system, using especially derived system metrics, it will be further shown that the proposed system of the RBH patient scheduling process involved a much larger degree of system automation and system integration, and a much lesser degree of fragmented data structures.

The first of the two new processes to be introduced is for the systematic and continual prediction of patients' resource consumption. This involved the use of computer models to predict patients' lengths of stay in each healthcare unit of RBH, as well as to which units they would require admission based on each patient's clinical and demographic characteristics. The output of this new process will be used as input into the second new process which used the resource consumption predictions in combination with information regarding the availability of healthcare resources to predict the projected performance of the component healthcare units of RBH given a proposed admissions schedule. This information could then be used by healthcare managers to modify the admission schedule of supply of healthcare resources to provide an optimal level of cost-effectiveness.

The objective of the second modelling chapter was to demonstrate the viability of implementing each of the two new processes whose requirements were defined in the previous chapter as software routines. To demonstrate the viability of the first process, a literature review was performed demonstrating that such models have already been proven viable in previous research – albeit for very different purposes. The viability of the second process was demonstrated through the design of a simulation model using coloured-timed Petri nets capable of satisfying the data and functional requirements of the process as defined in the previous chapter.

Given the results of the first chapter demonstrating the link between the cost-effectiveness of healthcare delivery and degree over control over patient scheduling and the demonstration of the viability of a computerised patient scheduling process to enable a greater degree of control presented in Chapter 5 and Chapter 6, the aim of this thesis will be supported.

2. Healthcare Resource Allocation

'Resource allocation' is here intended to describe the function of allocating resources between the different component processing units of a production process. As such, it is a complementary process to that of patient scheduling – the process of scheduling patients for admission to or discharge from healthcare units or services.

As a decision process it is worth considering resource allocation decisions in two distinct phases: the first phase being the making of the decision, the second phase being the implementation of the decision. The implementation phase is the allocation of a resource to a process, and as such may occur at any point in time before the actual consumption of the resource by the process. The resource remains allocated to the process throughout the production process. Thus, the fact that a resource has already been consumed will be indicated by it being allocated to a process at some time period in the past. Similarly, a resource will not have been consumed if its allocation to a process is for some future time period (or if it is yet to receive an allocation).

It will be assumed that, by definition, a resource may be consumed by a processing unit only if it has been allocated to that processing unit, and that the period of consumption will be the same as the period specified in the allocation. For those resources which are re-usable by the same process at different times, or which can be used in different processes either at the same or different times, multiple allocations are necessary. Thus, a resource may, for example, be allocated to process P_1 at time period t_0-t_1, and to P_2 at time period t_1-t_2.

Because the term 'resource' being used here refers not only to the raw material and labour used in a production process, but also, for example, the various machinery and tools, and because all resources need to be allocated to processes, the notion of a process itself becomes abstracted from its physical instantiation. Thus, resource allocation can be considered as the 'bringing together' of all the resources necessary to instantiate some production process at a particular point in time.

In the healthcare domain the resources are healthcare resources, such as nurses, drugs and the bricks and mortar comprising a healthcare facility. The production process is that of treating a patient or group of patients. Strictly speaking, patients should also be considered resources as, to use the analogy with the manufacturing industry, they are a raw material involved in the production process. As such, the terms 'resource allocation' and 'patient scheduling' are synonymous in the healthcare domain. Thus, healthcare resource allocation is the process of deciding which patients or groups of patients should be treated and when, given that a limited supply of resources means that all patients cannot be treated on demand. In this sense, it can be seen as the process of balancing supply and demand. However, it differs in two fundamental ways from the process of balancing supply and demand described in textbooks on economic theory. According to economic theory, supply and demand are determined by the consumers' willingness and ability to pay for the product. Thus, the product is sold only to those consumers who have decided that they can and will pay for it. While this is largely the case in the privately provided healthcare industry, in the publicly provided healthcare industry provision is based on perceived healthcare need rather than an individual's decision to pay for it. This means that a different decision process is needed to balance supply and demand than the 'invisible hand' of the free market. Moreover, the power to make the decision has to be transferred from the individual patient to the healthcare provider.

The power to make resource allocation decisions exists at two levels, determined by both the number of patients affected by the decision, and the time-span involved in making it. The policy level of decision-making considers entire groups of patients over a long-term perspective and is typically undertaken by central government or senior healthcare managers. The operational level of decision-making considers individual patients over a short-term perspective and is typically undertaken by more junior healthcare managers. It is the operational level of resource allocation which will be the main concern in this dissertation.

The objective of this chapter is to argue for the hypothesis that improving the process of healthcare resource allocation increases the cost-effectiveness of healthcare delivery, and that improving resource allocation can be facilitated by providing operational managers with more control over the patient scheduling process. To achieve this, a statistical model of healthcare resource allocation will be developed that is able to measure the outcome of resource allocation in terms of the cost-effectiveness of healthcare delivery. The relationship between cost-effectiveness of healthcare delivery and the effectiveness of resource allocation is examined by arguing that resource allocation may be seen as a control process, and that the effectiveness of resource allocation may therefore be measured as the degree to which it is able to control certain pertinent system variables.

2.1. Modelling Healthcare Resources

2.1.1. Proactive and Reactive Resource Allocation Decisions

For resource allocation to be a meaningful process, the resources to be allocated have to be, to some extent, generic in nature. That is, a resource has to be able to be involved in the treatment of more than one patient. A surgeon, for example, can operate equally well on a patient called Joe Bloggs as he can on a patient called John Smith. The objective of the resource allocation process is therefore to decide which, if either, of these two patients the surgeon will actually operate on, depending on their different healthcare needs. The same line of reasoning can be applied to all healthcare resources.

Although individual resources may be generic, a patient's overall healthcare needs tend to be more specific, if not unique. For example, although two patients may have the same diagnosis, they may have different bodyweights, and so require different amounts of anaesthesia, or one may require a longer stay in the intensive care unit, and so on. Thus, the generic nature of resources exists only in isolation. When a combination of resources is determined for the treatment of a particular patient, then that combination will often be specific to that patient with other patients requiring a different combination. Thus, the process of resource allocation involves many different decisions as to whether or not to allocate a resource, and if so in what quantity, for each individual resource for each individual patient.

Each resource allocation decision may be classified as being either reactive or proactive. Proactive resource allocation decisions are decisions which may be made well in advance of the consumption of the resource, based on predictions of the future healthcare needs of the patient. Reactive decisions, however, are made closer to the time of consumption and are made in response to unforeseen changes in the healthcare needs of the patient. The extent to which a patient's treatment is undertaken as a series of predominantly proactive decisions rather than predominantly reactive decisions depends on the state of medical knowledge and the skill of the clinicians concerned. For many common surgical procedures, for example, the treatment of the patient is predominantly proactive, since a mass of medical knowledge regarding the procedure will have been accumulated allowing the patient's healthcare needs to be more precisely determined. This line of reasoning lies at the basis of modern patient management paradigms, such as protocol management and the identification of critical pathways.

Because reactive resource allocation decisions are, by definition, made in response to unforeseen events, attempting to control when the resources will actually be consumed in an attempt to improve the cost-effectiveness of healthcare delivery is not possible. Proactive decisions, however, are made well in advance of consumption which leaves some scope for controlling when consumption occurs.

To see how this works, consider the process involved in making proactive resource allocation decisions. There are two components to resource allocation decisions, the 'if' component which determines if resources are to be allocated to a patient, and the 'when' component which determines when the resources are to be consumed. For example, in making resource allocation decisions for a surgical unit, the 'if' component is decided by the consultant responsible for the patient. If the consultant decides that the patient should undergo surgery, then the patient is placed on a surgical waiting list. A surgeon then decides when to perform the surgical procedure. The decision of when to operate will be based on the urgency of the patient's healthcare need and the availability of the resources necessary to perform the required procedure. If the patient's healthcare need is not urgent, then this leaves scope for controlling when the procedure is to be performed.

In surgical units, as with all healthcare units, the decision of when the patient begins consumption of the resources involved in treatment is, in effect, represented by the decision of when to admit the patient to the unit concerned. That is, once the decision to admit has been made and the resources have started being consumed, the resource allocation decision thereby becomes, in most cases, monotonic - it cannot be reversed. For example, when the patient is admitted to a surgical unit and the surgeon begins to operate, it becomes very difficult to reverse the decision to operate by ceasing the operation and sending the patient back to the ward. The situation may be seen as analogous to domino toppling; once the first domino is toppled, so the rest will be toppled in sequence. The decision to topple each domino, regardless of its position in the sequence, is made when one decides to topple the first domino. The difference between the decision to topple the first domino and the decision to topple the tenth domino in the sequence is therefore not when the decision is made, but when the decision takes effect. Thus, in the case of resource allocation, each proactive decision involved in a patient's treatment is made once the patient is admitted to the unit. It follows, therefore, that if the decision to admit a patient is in most cases monotonic when it takes effect, then any scope for controlling when the resources will be consumed exists in most cases only between the time that the 'if' component of the resource allocation decision is made and the 'when' component of the decision takes effect.

2.1.2. The Bed-slot Assumption

In modelling the resource allocation process a simplifying assumption is made with regards to the combination of resources which are thereby allocated to the patient. This assumption is that this package of resources is generic per unit of time for each patient admitted to the same unit. Thus, it is assumed that if Joe Bloggs and John Smith are admitted to a surgical unit, and both spend the same amount of time in that unit, then they both consume the same package of resources. If, however, only Joe Bloggs is admitted to the surgical unit, and John Smith is admitted to an intensive-care unit, then they will likely consume different resources, even if their duration in each respective unit is the same. This package of resources will be referred to as a bed-slot, and the assumption that the notion of a bed-slot may be effectively used in modelling the resource allocation process will be referred to as the bed-slot assumption.

Clearly, the bed-slot assumption is not an accurate reflection of reality. For example, a patient requiring a dental extraction and another patient requiring a heart transplant will require very different resources per unit time, even though they may both be admitted to the same surgical unit. The question to be addressed, however, is not whether the bed-slot assumption is an accurate reflection of reality, but whether or not it can be effectively used in modelling the resource allocation process. More specifically, for the purposes of this discussion, the question is whether or not the notion of a bed-slot can be effectively used in the development of a model of resource allocation which can be used as a tool in improving the cost-effectiveness of healthcare delivery.

As an initial attempt to answer this question it is necessary to first analyse whether the notion of a bed-slot can be applied to different types of resource considered in isolation. Second, it is necessary to determine, for each type of resource, whether or not its inclusion in a model of resource allocation has a significant effect on being able to use the model to improve the cost-effectiveness of healthcare delivery.

The classification of resource types needs to be informed by their particular characteristics which affect the different management techniques needed to optimise their productivity. The most important of these characteristics are as follows:

1. Whether or not the resource may be re-used or not;
2. the opportunity cost of using a resource;
3. the extent to which the consumption of the resource type is generic, and
4. the ease with which the supply of the resource can be modified to meet demand.

Table 2.01 below shows these characteristics for an initial classification of resource types defined along familiar lines as follows:

- **Space**. This refers to the physical space available within the healthcare facility.
- **Labour**. This refers to the human resources used either directly or indirectly in the delivery of healthcare. It thus includes not only clinical staff such as surgeons and nurses, but also administrative and other support staff.
- **Capital**. This refers to the physical equipment used either directly or indirectly in the delivery of healthcare. Thus, as with Labour, it includes not only equipment such as beds and mechanical ventilators, but also, for example, the office equipment used by administrative staff.
- **Consumables**. This refers to those non re-usable resources which are used either directly or indirectly in the delivery of healthcare, and includes item such as drugs, sterile gloves, disinfectant spray and disposable syringes.

Resource	Re-usable	Opportunity Cost High	Generic Consumption	Supply Fixed
Space	Yes	Yes	Yes	Yes
Labour	Yes	Yes	Yes	Yes/No
Capital	Yes	Yes	Yes/No	Yes
Consumables	No	No	No	No

Table 2.01. Characteristics of different types of resources.

In explanation of Table 2.01, Space, Labour and Capital are all re-usable, at least in the short-term, whereas Consumables are, by definition, not re-usable. The opportunity cost of Space, Labour and Capital is high, since in each case the cost of the resource must be met, regardless of whether or not the resource is consumed. A nurse, for example, must still be paid even though he/she is surplus to requirements in the short-term (although in the long-term, he/she may be made redundant). The opportunity cost of Consumables, however, is low, since these resources may only be used once. Space, Labour and Capital are all generically consumed to some extent. For example, a nurse or a bed may usually be used for any patient, regardless of the patient's diagnosis. For some Capital resources, however, and most Consumables resources, their consumption is non-generic. For example, not all patients admitted to an intensive-care unit will require a mechanical ventilator, or a particular drug. Although in the long-term the supply of all resources is variable, in the short-term Space, Capital and some Labour resources have a fixed supply insofar that any variation in supply has a very high opportunity cost attached to it. The supply of some Labour resources may be quite variable since it is possible to make overtime provisions or hire temporary staff when demand is high, and reduce overtime or temporary employment when demand is low. The supply of Consumables is not fixed, since they may be ordered on demand and/or kept in storage to be consumed as and when needed without incurring any extra cost.

These considerations suggest an alternative classification of resource types according to whether or not they are generic and whether or not they have a fixed supply, as follows:

- Generic and fixed: All Space resources, some Labour resources and some Capital resources.
- Generic and not fixed: Some Labour resources.
- Not generic and fixed: Some Capital resources.
- Not generic and not fixed: All Consumables.

This classification of resource types allows the question as to which resources the notion of a bed-slot may be meaningfully applied to be answered quite easily, since for a bed-slot to be able to be applied to a resource the resource must be generic. Therefore, the resource types Generic and fixed and Generic and not fixed from the above list may be modelled by the notion of a bed-slot; the resource types Not generic and fixed and Not generic and not fixed may not be modelled by the notion of a bed-slot.

The second question is to determine for each type of resource whether or not its inclusion in a model of resource allocation has a significant affect on being able to use the model to improve the cost-effectiveness of healthcare delivery. In this case, those resources which have a low opportunity cost if they are not used will have a lesser effect on being able to use a model to improve cost-effectiveness, since such resources involve a cost only when they are consumed and are not subsequently re-used. Therefore, reducing the cost of such resources could only be achieved through a corresponding reduction in effectiveness, which would have little or no consequent impact on overall cost-effectiveness. Thus, all Consumables resource types may be excluded from consideration in the development of a model aimed at improving cost-effectiveness. Conversely, all other resource types which have a high opportunity cost for not using those resources should not be excluded.

From this discussion, it can be concluded that Generic and fixed and Generic and not fixed resource types can both be modelled by the notion of a bed-slot; Not generic and not fixed resources cannot be modelled by the notion of a bed-slot, but the inclusion of these resources in a model aimed at improving cost-effectiveness would not have a significant impact on the model's results, and that Not generic and fixed resources cannot be modelled by the notion of a bed-slot, although excluding such resources in a model aimed at improving cost-effectiveness would have a significant impact on the model's results.

The extent to which these theoretical conclusions are validated by evidence will be analysed below and in Appendix 10. For the moment, it is worth noting that the bed-slot can be used for at least some resources. Further, as will be argued below, a model based on bed-slots can be used as a framework for complementary models of those resources whose patterns of consumption are not generic. Thus, on the principle that one should walk before one can run, attention will first be focussed on developing a model of resource allocation using the notion of a bed-slot as representing resources, before demonstrating how the same kind of model may be used to model those resources not able to be modelled by the bed-slot.

2.2. Effectiveness of Resource Allocation

2.2.1. The parametric assumption

The parametric assumption is that the distribution in the demand for resources may be approximated as a normal distribution. That is, that the analysis of the demand for resources can be performed using parametric statistics.

Consider a hypothetical healthcare facility with the following characteristics:

1. That each admitted patient has the same demand for resources equivalent to one bed-slot per patient per day, and that each patient consumes only one bed-slot
2. That all patients are admitted to the healthcare facility at exactly the same time each day (and therefore, since each patient stays for exactly one day, that all patients will be discharged from the healthcare facility at the same time each day).

3. That the average (mode) demand for bed-slots per day over a sample 1,000 day period is 15

The Parametric Assumption states that the distribution in demand for the bed-slots in this healthcare facility will approximate to a normal-shaped frequency distribution with mean level of demand being approximately equal to the mode level of demand, d^m, (15) and truncated such that $0 \geq d_i \leq N$, where N is the total size of the population. For the purposes of this discussion, this distribution will be truncated at 0 and 80[2] bed-slots per day. An instance of the resulting probability distribution is as shown in Figure 2.01 below (for bed-slot demand < 31).

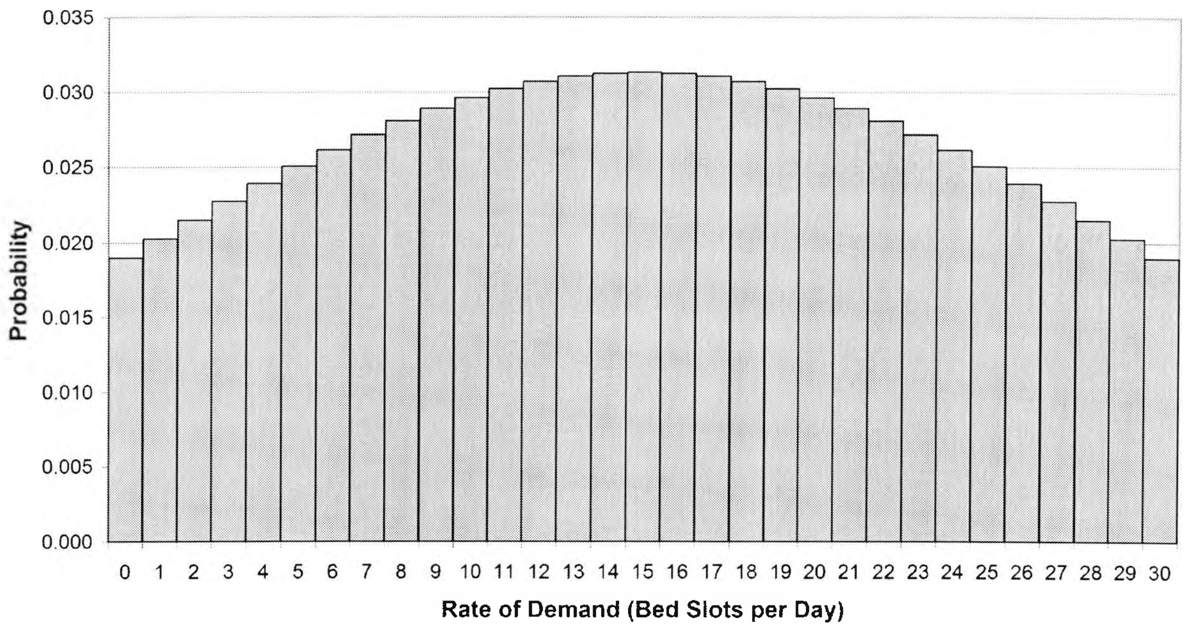

Figure 2.01. Distribution of bed-slot demand with no control over resource allocation and unlimited bed-slots.

Justification for the parametric assumption may be made by considering the generation of need for healthcare services as being a stochastic process where the distribution of bed-slot demand may be modelled as a binomial distribution. Let us say that the hypothetical healthcare facility is a surgical unit which serves a population of N = 250,000 people. If an estimate of the average demand per person for the bed-slots within the surgical unit is once in their lifetime, and the average lifetime is measured as 30,000 days, then a binomial distribution for the resulting bed-slot demand per day may be defined where the average demand is the number of trials, n, and the bed-slot demand is the probability of success, P. According to the rule of thumb that a binomial distribution may be

---

[2] Note that for the purposes of this discussion, the probability density has been normalised to sum to 1 over the interval 0 to 80 bed-slots per day (normalisation factor = 1.1774). Henceforward, all distributions will be thus normalised.

approximated by a normal distribution if $np > 5$ and $n(1 - p) > 5$, then the above distribution clearly satisfies both these conditions and may thus be approximated as a normal distribution.

Before proceeding, some simplifying assumptions will be made as follow:

1. In those cases where the distribution is truncated, the probability density distribution will be normalised over the truncated range, and
2. the mean that will be used to generate the distribution, d^0, will be assumed to be equal to the mean of the truncated distribution, which will be assumed to be equal to the mode level of demand, d^m, as above.

It should be further noted that the assumption of a homogenous patient intake is not necessary for the distribution of demand for bed-slots to be normally distributed, but is instead adopted here only to facilitate the discussion which follows in the next section.

2.2.2. Resource allocation as a control process

The main hypothesis of this chapter is that resource allocation is a control process, and that the cost-effectiveness of the underlying production process can be improved by improving the effectiveness of the control process.

To be classified as a control process, resource allocation needs to be defined as a process which maintains the system or some system variable at some ideal value or state. It is therefore necessary to identify what would constitute the ideal value or state of the system, the features of resource allocation which allow it to maintain that ideal value or state, and subsequently to develop a means of comparing the current value or state with the ideal value or state.

Consider the scenario where there is no control over the resource allocation within a healthcare facility serving the demand for healthcare services described above. That is, where the demand for the bed-slots of the healthcare facility is satisfied as and when it occurs. If the supply of bed-slots is unlimited, then the distribution of bed-slot consumption will be the same distribution as the demand for those bed-slots. For example, consider a unit, U, where the number of bed-slots, B=N. Assume that over a sampling period of 1,000 days, that the resulting distribution in demand for bed-slots may be approximated to a normal distribution shifted to the right such that the mode level of demand, $\bar{d} = 15$ and truncated such that $0 \leq d \leq 80$ as above. Figure 2.02 below shows the resulting probability distribution for U for $0 \le d \le 30$. As would be expected by the satisfaction of demand as and when it arises, the resulting distribution of consumption of healthcare resources is the same as the distribution of demand for healthcare resources shown in Figure 2.01.

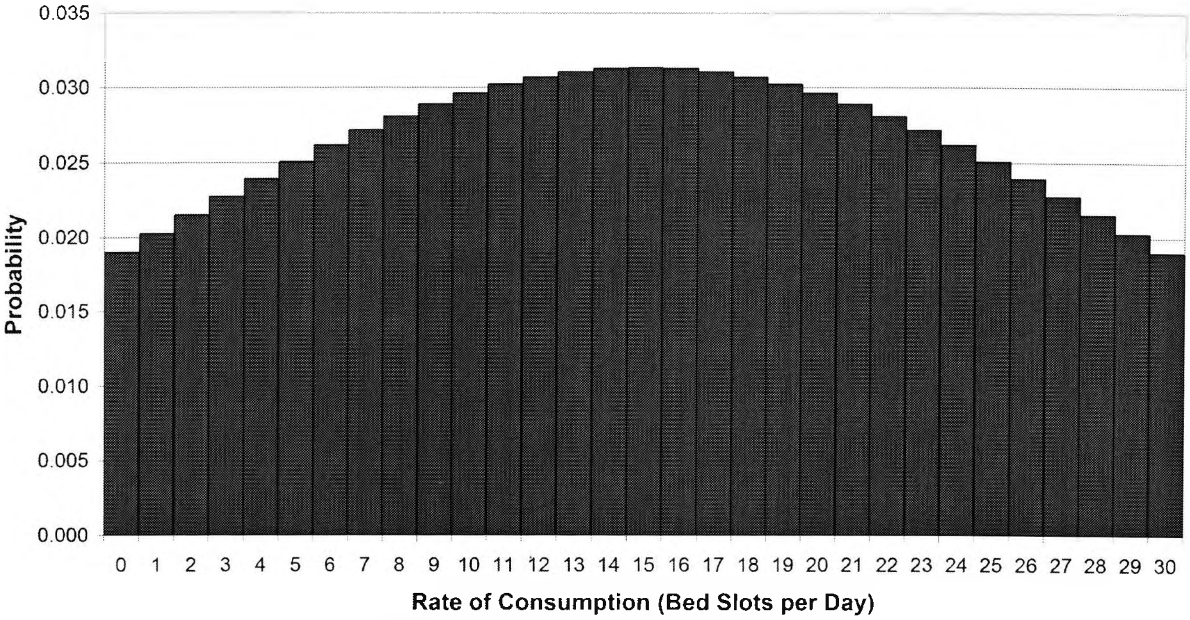

Figure 2.02. Consumption of bed-slot distribution with no control over resource allocation and unlimited bed-slots.

Naturally, however, the supply of bed-slots, B, is not the same as the size of the population, N..

Therefore, unless demand is always less than or equal to supply, it is not possible to satisfy all of the demand for bed-slots as and when that demand occurs. In the hypothetical healthcare facility depicted in Figure 2.02, for example, if it has a maximum supply of 20 bed-slots at any one time, the consumption – measured as the total amount of satisfied demand of bed-slot days - of bed-slots will be as shown in Figure 2.03 below.

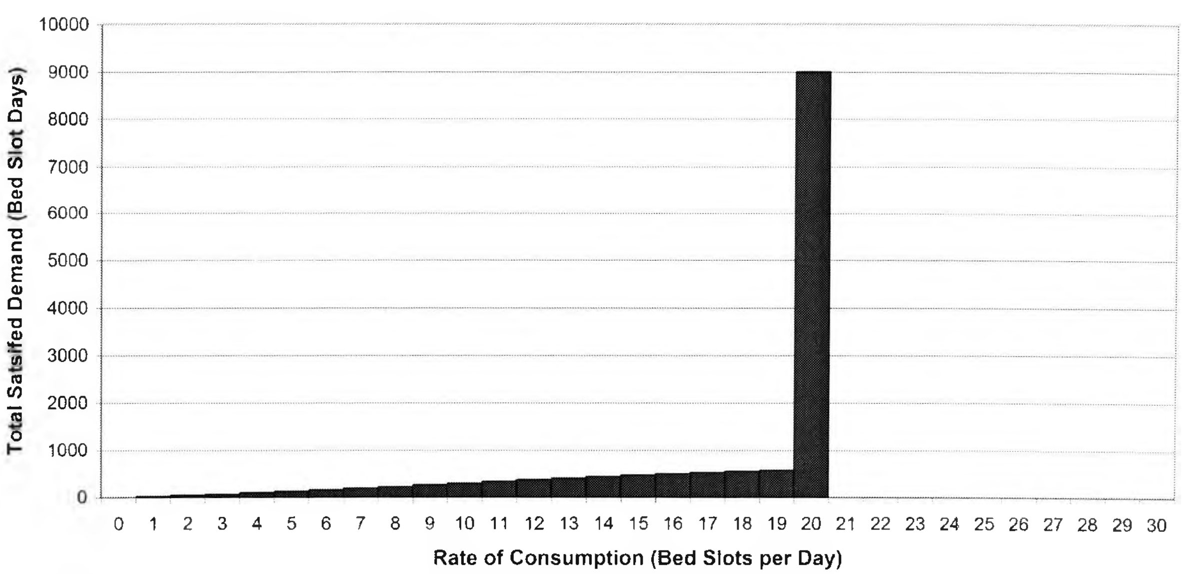

Figure 2.03. Distribution of satisfied bed-slot demand with no control over resource allocation and B = 20.

Note that the shape of the distribution is no longer normal, since the total bed-slot demand for those days when the level of demand generated by the population is greater than 20 bed-slots, the first 20 bed-slots demanded in each of those days may still be satisfied and the resulting consumption is therefore added on to the total consumption of bed-slots. Thus, the level of unsatisfied demand, D_{unmet} is calculated by Equation 2.01 below, where:

- N = The maximum amount of bed-slots that could be demanded per day = the total number of potential patients in the healthcare unit's service area;
- B = The maximum amount of bed-slots that could be consumed per day in the healthcare unit U;
- d = The level of bed-slot demand measured as the number of bed-slots demanded, where $0 \leq d \leq N$
- n_d = the frequency of days within the sampling period, T, when the level of bed-slot demand is d.

Eq.2.01.

$$D_{\text{unmet}} = \left(\sum_{d=B+1}^{d=N} n_d d \right) - \left(\sum_{d=B+1}^{d=N} n_d B \right) = \left(\sum_{d=B+1}^{d=N} n_d (d-B) \right)$$

A more meaningful measure of unsatisfied demand is to measure it as a proportion of total demand, D'_{unmet} as calculated by Equation 2.02 below.

Eq.2.02.

$$D'_{\text{unmet}} = \frac{\left(\sum_{d=B+1}^{d=N} n_d (d-B) \right)}{\left(\sum_{d=0}^{d=N} n_d d \right)}$$

Another useful measure is the mean utilisation rate of bed-slot resources, u. That is, the mean proportion of bed-slots which are allocated to a patient at any one time within T. This may be calculated by the dividing the total satisfied demand by the total availability of supply, BT. The mean utilisation rate is shown in Equation 2.03 below.

Eq.2.03.
$$u = \frac{\left(\sum_{d=0}^{d=B} n_d d\right) + \left(\sum_{d=B+1}^{d=N} n_d B\right)}{BT}$$

2.2.3. Measuring the Effectiveness of Healthcare Resource Allocation

Both u and D'_{unmet} are important variables when evaluating resource allocation in a healthcare facility as they both give indications of the cost-effectiveness of healthcare delivery, as discussed below. They are not, however, the best candidates for control variables. This is because their values depend in part on variables which do not measure the effectiveness of resource allocation, namely, the mode level of demand, d^m, and the number of bed-slots, B. A variable is needed, therefore which measures the effectiveness of resource allocation qua control process, which is independent of extraneous features of the healthcare facility in which the resource allocation occurs. The variable proposed here is the variance in the distribution of demand, measured as the standard deviation, σ, of the normal distribution function, $f(d^m)$, that best approximates to the shape of the distribution in demand.

In the hypothetical healthcare facility discussed above where $B = N$, $D_{\text{unmet}} = 0$ because the demand distribution and consumption distribution of bed-slots were identical. Thus, whenever $B < N$, $D_{\text{unmet}} > 0$. Crucially, this is not a necessary condition, however, if one assumes that it is not necessary to satisfy demand for bed-slots as and when it occurs. That is, when demand is greater than some reference value, the demand for bed-slots which is in excess of that value can be left unsatisfied until periods when demand is less than the reference value. For example, if someone is diagnosed with cancer, it is not normally essential to achieve a good treatment outcome for them to be admitted for chemotherapy or surgery on the same day of diagnosis. By delaying the admission – in effect, creating some system 'slack' by creating an admissions queue – the distribution in demand can be re-shaped to optimise bed-slot utilisation rates.

Thus, resource allocation becomes a control process whereby the number of bed-slots allocated per day is maintained at some ideal value. The degree to which this is achieved may be considered as determining the effectiveness of resource allocation.

It will be argued here that in re-shaping the distribution in demand, the original normal distribution $f(d^m)$ may be modified to a new distribution, $f'(d^m)$, where the standard deviation of $f(d^m)$, σ, is greater than that of $f'(d^m)$, σ'.

Consider the case where $\sigma' \underset{\lim}{\alpha} 0$, in which case the probability of demand at any point in time being equal to d^m will likewise approximate to 1. In this case, the bed-slot utilisation rate, u, will be as follows:

Eq.2.04.

$$u \underset{\lim}{\alpha} \frac{d^m}{B}$$

In other words, in the above example, $u \underset{\lim}{\alpha} 0.75$. Moreover, $D_{\text{unmet}} \underset{\lim}{\alpha} 0$ since the likelihood of demand ever exceeding the number of bed-slots, B, will tend towards 0.

Clearly, the distribution in demand cannot be re-shaped to the extent that $\sigma' \underset{\lim}{\alpha} 0$, but nonetheless, the more frequently that demand can be kept less than or equal to the number of bed-slots, the lower will be the proportion of unsatisfied demand.

2.3. Cost-effectiveness of Healthcare Delivery

2.3.1. Measuring the Cost-effectiveness of Healthcare Delivery

In the previous section, the standard deviation of the demand distribution was proposed as measuring the effectiveness of the resource allocation process. However, this measure is unable by itself to evaluate the outcome of the resource allocation process within the context of a real life healthcare facility insofar that it assumes a number of bed-slots equal to N, the number of people in the overall population. Moreover, it measures only the effectiveness of the control over resource allocation, rather than the effectiveness of the resulting healthcare delivery or the efficiency with which it is delivered. Thus, the question to be addressed here is, How to evaluate the outcome of the resource allocation process in terms of the efficiency and effectiveness of healthcare delivery? The evaluation parameter which will be proposed here is cost-effectiveness, in which case the question becomes, How to measure cost-effectiveness?

Unlike the notion of efficiency, cost-effectiveness considers the effectiveness of a process, as well as its cost, where cost is inversely proportional to efficiency, i.e., the more efficient a production process is, the lower the cost of production. In this sense it is therefore a superior parameter. It is often a very simple matter to increase efficiency (i.e. decrease cost) just by placing less emphasis on the quality (i.e. effectiveness) of the product. For example, a software company may be able to halve its production costs by employing less programmers, but the outcome may well be that the software has many more bugs in its code or has less functionality or a worse user interface than the main competitor's product. In such cases, although efficiency has increased, effectiveness has decreased. Thus, in order to determine whether or not the increase in efficiency was worthwhile the notion of cost-effectiveness is used which takes into consideration the effects on both efficiency and effectiveness. Similarly, to determine whether or not an increase in the effectiveness of a product is justified, the effects on cost need to be considered. These considerations point to cost-effectiveness being calculated as shown in Equation 2.05 below.

Eq.2.05

$$\text{Cost - effectiveness} = \frac{\text{Effectiveness}}{\text{Cost}}$$

The problem remains, however, of how to measure cost-effectiveness. Clearly, it is calculated using measurements of cost and effectiveness individually. Cost is relatively easy to calculate, whether it be in dollars, pounds or yen. Effectiveness is more difficult, and will vary for each product. To measure the processing speed may be a good measure of effectiveness for computers, for example, but not for fruit cakes. Moreover, effectiveness will often need to consider many different dimensions of a product's quality. In computers, for example, effectiveness may be measured not only be processing speed, but also hard disk capacity, reliability and aesthetics.

In determining the effectiveness of healthcare resource allocation, a variable needs to be identified which measures the quality of the healthcare which is provided by a particular resource allocation process. If the aim of resource allocation is to satisfy as much of the demand for healthcare resources as possible with the resources that are available, then the effectiveness of a resource allocation process may be measured by the extent to which this aim is achieved. Thus, in the case of a particular resource, the effectiveness of its allocation to patients may be measured by the proportion of the total demand which is satisfied by such allocation. In the case of bed-slots, therefore, effectiveness is measured using the variable D_{unmet}, as calculated by Equation 2.06 below.

Eq. 2.06
$$\text{Effectiveness} = D_{\text{met}} = (1 - D_{\text{unmet}})$$

Where D_{unmet} is calculated by Equation 2.02 above.

The measurement of cost, as it occurs in the healthcare industry, can be calculated by the average cost of a bed-slot per patient. This will be inversely proportional to the average utilisation rate of bed-slots, U, as shown in Equation 2.07 below.

Eq.2.07
$$\text{Cost}_{\text{mean}} = \frac{\text{Cost}_{\text{total}}}{\text{Consumption}_{\text{total}}} = C \times U^{-1}$$

That is, the average cost of a bed-slot per patient is the total cost of all bed-slots over the number of patients whose demand for bed-slots is satisfied, which is inversely proportional to the mean bed-slot utilisation rate, U, as calculated by Equation 2.03 above.

To make the discussion as general as possible, cost will be measured in terms of efficiency by fixing the constant, C in Equation 2.07 to 1. Thus, cost-effectiveness may be calculated by Equation 2.08 below.

Eq.2.08
$$\text{Cost - effectiveness} = U(1 - D_{\text{unmet}})$$

Thus, cost-effectiveness may take on any value between 0 and 1. A value of 1 signifies the maximum possible cost-effectiveness, and requires that $U = 1$ and $D_{unmet} = 0$. A value of 0 signifies the least possible cost-effectiveness, and requires that $U = 0$ or $D_{unmet} = 1$. It is interesting to note that for cost-effectiveness to be 1, both effectiveness and cost have to take on certain values, whereas for cost-effectiveness to be 0, only one has to take on a particular value. This is because to satisfy all bed-slot demand does not imply a utilisation rate of 1 (in fact, if there is to be any spread at all in the distribution of demand, U will necessarily be less than 1), whereas if utilisation rate of bed-slots is 0, then necessarily no demand will have been met since no bed-slots will have been consumed.

Consider the cost-effectiveness of the distribution $f(d^m)$ described above, and the modified distribution $f'(d^m)$, where the standard deviation of $f(d^m)$, σ, is greater than that of $f'(d^m)$, σ'. It can be seen that, because the distribution $f'(d^m)$ has a lower value of D_{unmet} it will likewise have a higher level of cost-effectiveness. In fact, as $\sigma' \underset{\lim}{\alpha} 0$, then both the utilisation rate and the cost-effectiveness will tend towards $\frac{d^m}{B}$ which, in the above example is 0.75. Moreover, in this case, cost-effectiveness may be increased further by increasing the mode rate of admissions to d'^m, where $d'^m > d^m \leq B$.

2.3.2. Modelling Not Generic and Fixed Resources

So far the discussion has been devoted to the effectiveness of bed-slot allocation. However, as argued above, the notion of the bed-slot excludes certain types of resources which are nevertheless important elements to be included in a model of resource allocation which aims at increasing the cost-effectiveness of healthcare delivery. These resources, which were classified above as not generic and fixed, include expensive capital equipment such as mechanical ventilators and cardiopulmonary bypass machines.

To include these resources in a model of resource allocation is a relatively straightforward matter once a model of bed-slot resource allocation has been developed. Consider, for example, the case of allocating mechanical ventilators amongst the same patient population used in the above analysis. There are two alternative ways by which this may be modelled along the same lines as the allocation of bed-slots.

First, the patient population may be divided into two groups – those who require a bed-slot and a mechanical ventilator allocation, and those who only require a bed-slot allocation. Alternatively, those patients requiring mechanical ventilation may be considered as a subset of those patients requiring a bed-slot. In either case two distributions are formed whose characteristics may be used in deriving measures for the effectiveness of resource allocation and for the cost-effectiveness of healthcare delivery along the same lines as the derivation of these measures for the single distribution of bed-slot demand discussed above.

In deriving an overall measure of cost-effectiveness where there are multiple demand distributions, it is possible to take an appropriately weighted sum of the individual measures of cost-effectiveness for each distribution $f(d_i')$. Thus, in the case of n distributions, the overall cost-effectiveness of healthcare delivery may be calculated by Equation 2.09 below.

$$\text{Cost - effectiveness} = \sum_{i=1}^{n} U_i (1 - D_{\text{unmet}i}) p_i w_i$$

Eq.2.09

Where p_i is the proportion of the total patient population to whom the resource is allocated, and wi is an additional weighting according to the importance to which the management ascribes to optimising the cost-effectiveness of that particular resource. For example, in the case of mechanical ventilators, because they are a relatively expensive individual resource item, and because their demand tends to be of a more critical nature than the demand for other resources, the management will consider the optimisation of the delivery of this resource of a greater priority than that of other resource, even though those patients requiring mechanical ventilation may constitute only a relatively small proportion of the total number of patients.

It is worth noting that Equation 2.09 may be used to derive a measure of cost-effectiveness for whichever alternative approach one adopts to modelling the different demand distributions, the only difference being the value to which the individual p_i sum.

A further point to note is that, if it is assumed that those patients demanding, for example, mechanical ventilation, will also demand bed-slots, controlling one resource thereby controls, to some extent, the other. The reasoning for this conclusion is based on the notion of explained variance. That is, if the only factor which predicts the level of demand for mechanical ventilators is the level of demand for bed-slots, then all variance in the level of demand for mechanical ventilators will be explained by variance in the level of demand for bed-slots. Therefore, a reduction in the variance in the demand for bed-slots will reduce the variance in the demand for mechanical ventilators. Unfortunately, the level of demand for mechanical ventilators is predicted by more factors than just the level of demand for bed-slots, in particular the proportion of the demand for bed-slots which derives from particular types of patients, classified according to the severity of their illness or their diagnosis. Thus, while demand for bed-slots can explain at least some of the variance in the distribution for the demand of mechanical ventilators, unless the proportion of the bed-slot demand which comes from each type of patient is assumed invariant, it will not explain all variance in demand for mechanical ventilators.

2.4. The Real World

The discussion so far has been based on very simplified conceptions of both the patient population and the healthcare facility whose resources they demand. In the case of the patient population, it was assumed to be homogenous insofar that they each had the same healthcare needs which equated to a 24 hour bed-slot in for each patient, which moreover was allocated to the patient at the same time each day. In the case of the healthcare facility, it was implicitly assumed that the facility was a 'stand alone' one whereby it had no affect on, nor in turn was affected by, a wider healthcare environment. In this section this simplified conception will be developed into a more accurate reflection of the real world by abandoning all of the assumptions made above.

2.4.1. The Real Patient Population

Different patients will have different healthcare needs, and thus represent different demands for healthcare resources. This difference in demand has two dimensions for each resource demanded – when the resource will start being consumed, and the length of time of its consumption.

A patient population with heterogeneous healthcare needs does not in itself imply that the cost-effectiveness of healthcare delivery will necessarily be sub-optimal (i.e., less than 1). For example, if all of the patients healthcare needs are known, and there are no urgent cases, and the number of resource units available within the healthcare facility is greater than or equal to the mean level of demand for those units, then complete control can still be exerted over the allocation of each resource unit. Unfortunately, in almost every healthcare facility, there are urgent cases, or at least varying degrees of urgency, whether the urgency originates from a genuine healthcare need or a requirement to achieve objectives aimed at reducing waiting lists. Further, in many cases it is not possible to predict a patient's exact healthcare needs in terms of the resulting demand for particular resources. These two problems will be discussed in turn.

The greater the proportion of urgent cases within a patient population, then the less control may be exerted over resource allocation. From the discussion above, the effectiveness of control of resource allocation is determined by the degree to which it is possible to reduce the variance in the $f'(d)$ distribution. And the degree to which it is possible to reduce the variance in the $f'(d)$ distribution is determined by the extent to which it is possible to delay the consumption of resources to those patients whose perceived demand for those resources first occurs on those days when overall demand is in excess of some reference value. However, in the case of urgent cases, the extent to which it is possible to delay consumption is necessarily limited, and hence the effectiveness of resource allocation is reduced.

Just as in the case of there being more than one type of resource, it is possible to derive different distributions for different types of patient. If the perceived demand for bed-slots is classified as being urgent or non-urgent on the basis of whether or not the consumption of the bed-slot may be delayed for a significant period of time, then the two distributions $f'(d)_{\text{urgent}}$ and $f'(d)_{\text{elective}}$ may be derived for the urgent and elective (non-urgent) consumption of bed-slots, respectively.

Using the simulated data from the hypothetical healthcare facility discussed in the preceding sections, if it is assumed that one third of the perceived demand for bed-slots is classified as being urgent, then the two distributions for the consumption of bed-slots are as shown in Figure 2.10 and Figure 2.11 below.

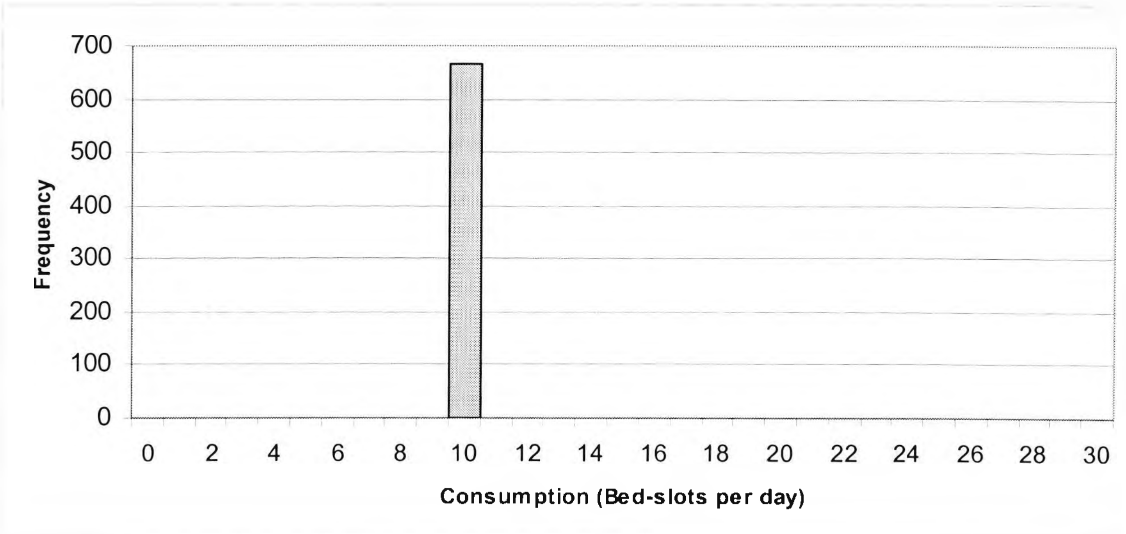

Figure 2.10. Distribution of non-urgent bed-slot consumption

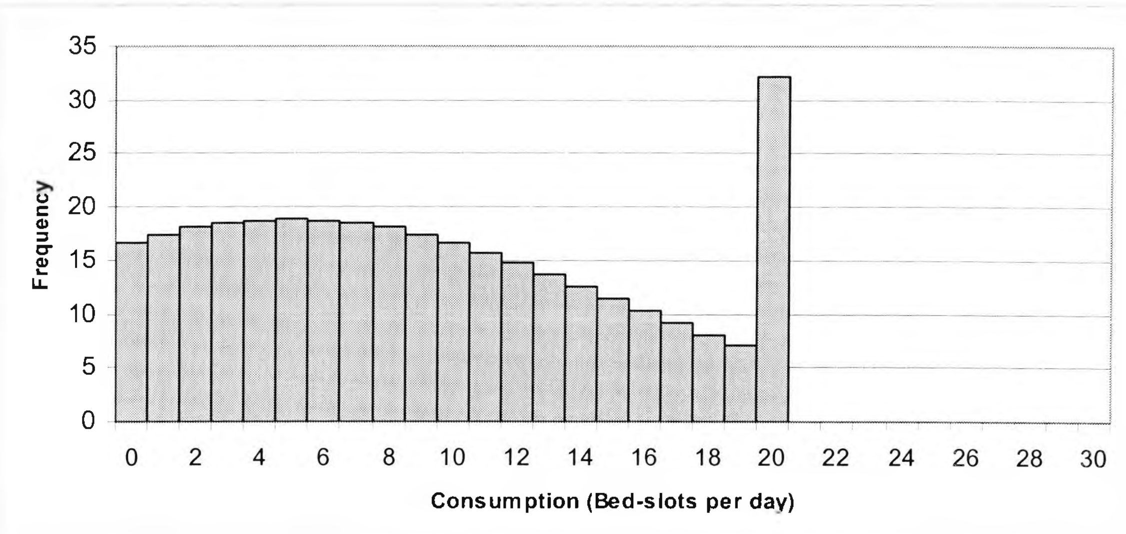

Figure 2.11. Distribution of urgent bed-slot consumption.

Because all of the non-urgent bed-slot consumption is able to be delayed, it is possible to exert complete control over the allocation of those bed-slots. Thus, the non-urgent consumption of bed-slots is maintained at the mean level of demand for those bed-slots of 10 per day as shown in Figure 2.10. Urgent bed-slot consumption, however, may not be delayed, nor therefore controlled, which results in the distribution shown in Figure 2.11.

Figure 2.12 below shows the resulting distribution for all bed-slot consumption which is derived through the summation of the frequencies of both urgent and non-urgent bed-slot consumption rates.

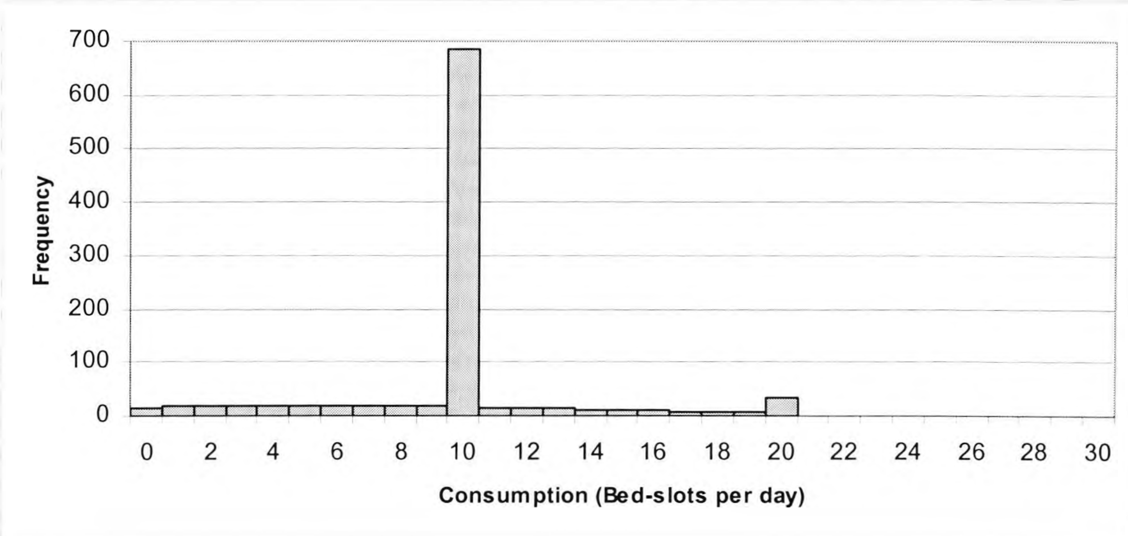

Figure 2.12. Distribution of urgent and non-urgent bed-slot consumption.

When measuring the effectiveness of resource allocation in the case of a mix between urgent and non-urgent bed-slot demand, the SD of the $f(d'_{\text{non-urgent}})$ should be considered, rather than the distribution of summed frequencies of both urgent and non-urgent bed-slot demand, since urgent bed-slot demand is beyond the control of resource allocation. Similarly, when calculating the cost-effectiveness of healthcare delivery, it is more meaningful to derive a value which is relative to the inherent limitations implied by urgent bed-slot demand. Such a value would measure relative cost-effectiveness as a proportion of the maximum value attainable given the urgent bed-slot demand. Thus, relative cost-effectiveness would be calculated by Equation 2.10 below

$$\text{Cost-effectiveness}_{\text{rel}} = \frac{U(1-D_{\text{unmet}})}{p_{1\text{-u}} + p_{\text{u}}[U_{\text{u}}(1-D_{\text{unmet,u}})]}$$

Eq.2.10

Where $p_{1\text{-u}}$ is the proportion of the total bed-slot demand which is non-urgent, p_{u} is the proportion which is urgent, and $U_{\text{u}}(1-D_{\text{unmet,u}})$ is the cost-effectiveness of the delivery of healthcare for urgent bed-slot demand (the cost-effectiveness of non-urgent bed-slot demand will be fixed at 1, the maximum value).

With regards to the problem of predicting patients' demand for resources, the degree to which it affects the effectiveness of resource allocation and cost-effectiveness depends on various factors. At prima facie level, being able to control resource allocation presumes the ability to determine which resources will be allocated to which patients and when. Therefore, if it is not possible to determine a patient's required resource allocation, it is not possible to control it. However, as already discussed, the main source of control over resource allocation is to determine when resource consumption starts. And in the case of non-urgent demand, the point at which consumption starts can be controlled. Thus, while it is true that resource allocation becomes a largely uncontrolled process once consumption begins, it is when it begins which is of most importance, and this may be controlled for non-urgent demand. To demonstrate this, consider the case where, a patient has an allocation of one day's bed-slot and whose consumption of that bed-slot has already begun. It then becomes clear that the patient will require an allocation of two day's bed-slot, and that this is beyond the control of resource allocation insofar that it is determined by the patient's disease process and therapeutic requirements rather than any cost-effectiveness targets. It remains possible, however, to control the allocation of resources to patients whose consumption of those resource is yet to begin by simply re-allocating the resources to one of these patients to be consumed one day later than originally intended. Thus, having imperfect knowledge over patient's required resource allocation does not necessarily result in sub-optimal cost-effectiveness of healthcare delivery if one retains control over when consumption of those resources begins.

There are two cases where control over when consumption begins is lost and imperfect knowledge over patients' resource needs can impact negatively on the cost-effectiveness of healthcare delivery. The first case is where there is a significant amount of urgent demand for bed-slots. If a patient consumes a bed-slot for longer than specified in the original allocation, then this could result in an urgent demand for a bed-slot being left unsatisfied. The second case is where the healthcare facility is part of a larger healthcare environment, as discussed below.

2.4.2. The Real Healthcare Facility

The discussion so far has considered a healthcare facility in isolation. However, most healthcare delivery is undertaken in the context of a system of interconnected facilities or units. For example, in the case of a surgical facility, a patient will first be admitted to a pre-operative ward before being admitted to the operating room itself. Then, being discharged from the operating room, the patient may subsequently be admitted to an intensive care unit or recovery room before finally being admitted to a post-operative ward.

Such progressive care systems are designed to reflect the progressive nature of a patient's therapy by making different units functionally distinct. This further allows for the specialisation of resources, and in particular labour resources and the more effective logistical management of those resources. It also makes resource allocation and the optimisation of the cost-effectiveness of healthcare delivery more difficult to achieve, at least so long as there is imperfect knowledge of patients' resource needs.

The control over resource allocation in the case of an isolated healthcare facility is gained by being able to determine when the resources will be consumed. In a progressive care system, however, such as the example of the surgical facility mentioned above, while it is possible to exert the same level of control over when the resources will be consumed at a global (i.e., system) level, it is not possible at a local (i.e. unit) level. Moreover, this localised loss of control is not dependent on imperfect knowledge of patients' resource needs.

In a progressive care system, which consists of three units A, B and C, patients are allocated bed-slots in each of these units so that they are first admitted to unit A, then B and finally C. Because a patient is admitted to unit B from unit A and to unit C from unit B, allocating a bed-slot in unit A requires that contiguous bed-slots be allocated in units B and C. Therefore, allocating a bed-slot in unit A thereby allocates contiguous bed-slots in B and C. In this sense, resource allocation within a progressive care system can only occur at a global level. This is especially the case in critical care environments such as the example of the surgical facility used above. In such environments, it is not usually possible for patients to queue between the discharge from one unit and the admission to the next unit because of the danger of delaying admission in critically ill patients. Moreover, in those cases where it is necessary for a patient to queue for admission to a subsequent unit, because the patient is typically too critically ill to, for example, wait in a waiting room, queuing must occur in the discharging unit, which is not only very expensive, but because of the global allocation of resources, also disruptive to the entire resource allocation process. For example, if a patient is to be discharged from an operating room to an intensive care unit, but the intensive care unit is full, the patient must queue for admission to the intensive care unit in the operating room where an equivalent level of care is able to be provided. This is very expensive because of the very high fixed costs of the operating room, and disruptive to resource allocation because any patient who was allocated operating room resources while it is being used as a queuing area for the intensive care unit must themselves queue for the operating room, and so on.

2.5. Resource Allocation and Cost-effectiveness in a Progressive Care System

The process of resource allocation in a progressive care system may be conceptualised as being analogous to the problem in mathematics of tiling the plane. For a two dimensional area and a set of floor tiles, the problem of tiling the plane is how to use the tiles to cover as much of the area as possible while not allowing for any overlap.

The ease with which it is possible to tile a given area (measured as, for example, the length of time it takes an AI search algorithm to reach an optimum arrangement of tiles) is dependent on the shape, size and variety of the tiles. In the simplest case where all of the tiles are the same shape and size and square, the area may be completely covered. In the most difficult case where none of the tiles are the same size or shape, then it is likely that a complete covering of the area will be impossible.

The analogy between tiling the plane and resource allocation may be constructed by the following identifications:

- White space (i.e. uncovered areas) = unallocated resources;
- Covered space = allocated resources;
- Floor tiles = patients.

In the case of bed-slots and a homogenous patient population, the situation is depicted in Figure 2.13 below.

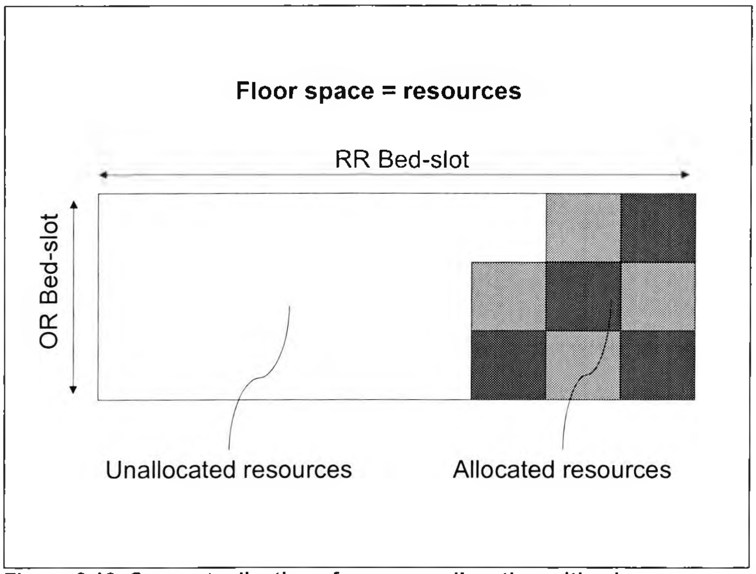

Figure 2.13. Conceptualisation of resource allocation with a homogenous patient population.

In Figure 2.13, the different dimensions correspond to the length of time for which the bed-slot is allocated to a patient. In this case a very simple progressive care system is represented consisting of an operating room (OR) and a post-operative recovery room (RR). More complex progressive care units will be represented by increasing the number of dimensions. Thus, a system consisting of three healthcare units may be represented by a cube, and systems with more than three units may be represented by hypercubes. In each case, however, the same principle remains of covering white space or filling an empty volume.

Because the patients are assumed to be homogenous, the floor tiles which represent them are all the same shape, and thus it becomes possible to completely cover the area. When a heterogeneous patient population is assumed, however, the possibility of being able to fully allocate all bed-slots becomes remote, as shown in Figure 2.14 below.

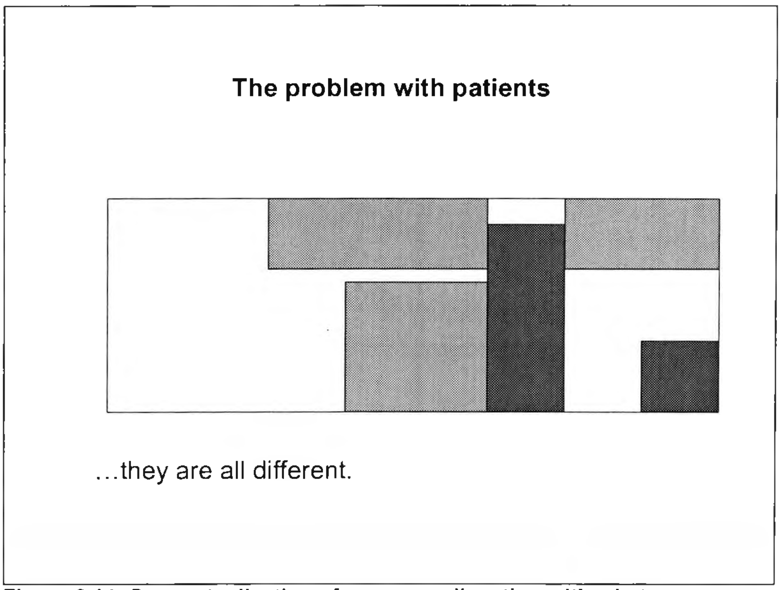

Figure 2.14. Conceptualisation of resource allocation with a heterogeneous patient population.

Figure 2.14 demonstrates how the allocation of a bed-slot in one unit of a progressive care system constrains the choice over the allocation of bed-slots in another unit. In the case of having complete control over resource allocation, however, the allocation of all bed-slots in all of the units all of the time remains a possibility despite a heterogeneous bed-slot demand. But this is only possible if the demand is so heterogeneous as to cover every possible combination of bed-slot requirements in each unit, so that any size or shape of any remaining white space may be allocated. This is an unrealistic situation since the possibility assumed is mathematical rather than clinical. That is, the degree of heterogeneity achievable within the patient population is constrained by the healthcare needs of the population rather than the ability of combinatoric mathematics to generate infinite variety. The complete control of resource allocation within a progressive care system with heterogeneous demand is therefore effectively impossible.

The possibility of complete control over resource allocation becomes even more remote when the problems of urgent bed-slot demand and imperfect prior knowledge of healthcare needs are considered. As mentioned in the case of an isolated healthcare facility, it is not possible to control the resource allocation for urgent bed-slot demand. This is equally true in the case of the progressive care system. Using the tiling analogy above, it is as if, for some tiles, the choice over where they are placed is removed, which therefore constrains the choice over where other tiles are placed.

The problem of imperfect information of patients' healthcare needs, as argued above, does not necessarily limit control over resource allocation, since non-monotonic resource allocation decisions may be modified to account for unforeseen resourcing requirements. In the case of a progressive care system, however, this flexibility may only exist at the local level. For example, in the case of a surgical facility consisting of an operating room and a post-operative recovery room, it may be possible to re-allocate an operating room bed-slot because a patient requires a longer than expected period of time in the operating room before being admitted to the recovery room without thereby reducing the cost-effectiveness of the delivery of operating room resources. Unfortunately, it will not be possible to re-allocate recovery room resources in light of the fact that the admission of a patient from the operating room will be delayed, thus leaving some resources unutilised. This is because all patients admitted to the recovery room will come from the operating room, and thus the allocation of recovery room resources is entirely subsequent to the allocation of operating room resources.

In terms of the tiling analogy above, the problem of imperfect knowledge over patients' healthcare needs is analogous to the situation of tiling a plane with a set of tiles whose dimensions are not known. One approach to this problem, short of possessing perfect knowledge of patients' healthcare needs, would be to increase the space between tiles so that some degree of variance in the dimensions may still be accommodated. For patients, this solution amounts to defining a confidence interval in, for example, the duration of bed-slot consumption and fix the bed-slot allocation to each patient at the upper confidence limit. The idea is also at the basis of the so-called block-scheduling of operating room resources, whereby the allocation of bed-slots are standardised to a block of time which is sufficiently large to accommodate most operative procedures without having to delay subsequent procedures.

There are two possible problems with adopting the confidence limit approach in units other than the operating room, however. The first is that the standardised allocation size will necessarily imply (for a normal distribution) that the cost of excluding only an acceptably small amount of excessive demand is to include a large number of cases where demand is less than the standard allocation size. This under utilisation will be lessened in the case of the distribution of required allocation sizes being skewed to the right. Unfortunately, the typical skewness of such distributions is actually to the left, and in many cases this degree of skew makes the notion that the distribution may be approximated as a normal distribution implausible.

Figure 2.15 below models a hypothetical distribution in the quantity of bed-slot resources consumed by each patient. The distribution used in this case is a gamma distribution, although a Weibull distribution may also be used.

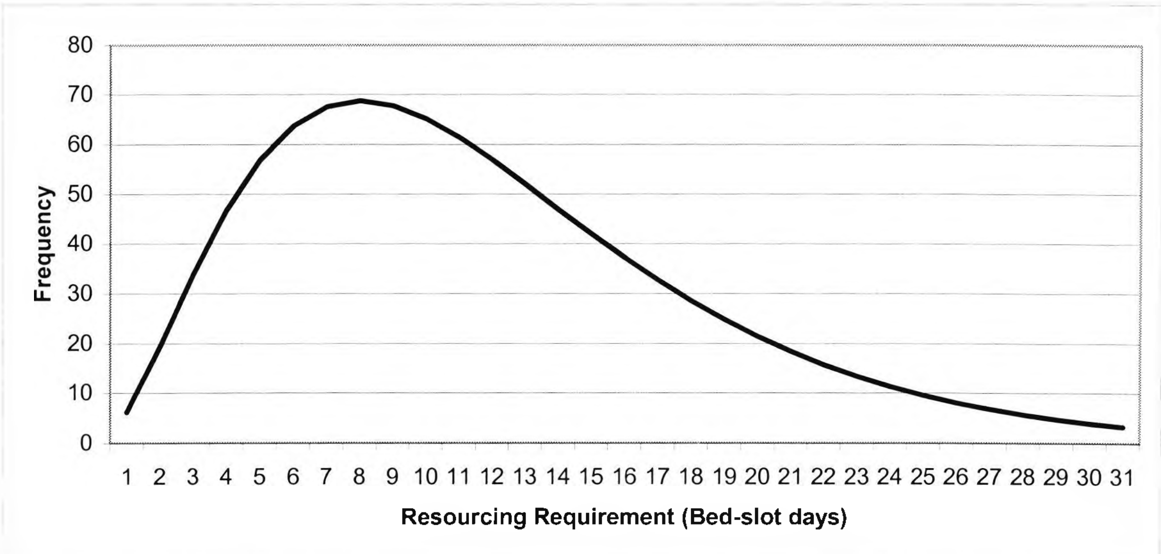

Figure 2.15. Distribution in resourcing requirements, modelled using a gamma function.

The second problem is that, although those cases whose demand is in excess of the standard allocation size are relatively small in proportion to those cases whose demand is not in excess, the proportion of the total demand represented by the excluded cases will be much greater, simply because their demand is greater. For example, in the above gamma distribution, if the standard allocation size is 20 bed-slot days, then approximately one-eighth of the patients will require a bed-slot allocation in excess of the standard. However, this one-eighth of patients corresponds to almost one quarter of total demand. Moreover, it is often found that those patients who have a larger overall demand for resources are also those patients whose resourcing requirements are the least predictable. Thus, while the use of confidence limits to determine a standard allocation size may be of use in the allocation of operating room resources, it is not necessarily applicable to all types of unit.

As with the case of the isolated healthcare facility, the inclusion of urgent cases amongst the patient population will accentuate any detrimental affects that imperfect knowledge over patients' healthcare needs will have over the cost-effectiveness of healthcare delivery because of the high degree of interdependency between the different units.

2.5.1. Classification of Control-limiting Factors

So far various factors have been discussed which restrict the possibility of a healthcare manager attaining complete control over the allocation of resources. All of these control-limiting factors derive from properties of the patient population. That is, their healthcare needs and the degree to which those needs may be accurately assessed. However, just as the properties of the patient population can limit the level of attainable control, so too can the properties of the healthcare resources themselves.

Complete control over patients consists in being able to determine when they will consume healthcare resources. Similarly, complete control over healthcare resources consists in being able to determine when they will be consumed by patients.

The most obvious example of a control-limiting factor which derives from healthcare resources is the restrictions on the allocation of those resources which are implied by staff scheduling. For example, if a nurse decides that he will take some annual leave, then this implies a restriction of choice over when he may be allocated to care for a patient. Similarly, if a surgeon is only able to perform procedures in the morning because of other commitments, he may not be allocated to perform procedures on a patient during the afternoons. Nonhuman resources may also restrict control through, for example, maintenance requirements when they are taken out of production for a period of time.

In a progressive care system a particular control-limiting factor arises when one of the component units operates for a different period of time than other units. For example, in the case of the surgical facility described above, the operating room may only operate for non urgent cases during normal working hours, whereas the recovery room may operate for 24 hours each day. To optimise cost-effectiveness of recovery room operation, therefore, those patients who require bed-slots of longer duration must be allocated bed-slots which begin later in the day so that recovery room resource utilisation is maintained at a high level during those hours when the operating room is closed.

There are various ways that the control-limiting factors have been discussed other than in terms of whether they derive from patients or not. For a healthcare manager whose aim is to increase the effectiveness of resource allocation, it is important to classify control limiting factors in terms of whether or not they can be changed or not. That is, whether they should be considered as inherent to the process of healthcare delivery or not. In this respect, another more fundamental system of classification is also helpful; whether or not the control limiting factor is epistemological in nature. That is, whether control is limited by a lack of knowledge of the controlled system, or because of the system itself.

The set of control limiting factors which are identified for a healthcare facility and the extent to which they limit control will vary between different healthcare facilities and between different patient populations. A method for the identification and classification of control-limiting factors in a healthcare environment is presented in Appendix 11, along with the application of the method to the RBH High-dependency Environment.

In this thesis the type of healthcare facility which will be considered is a surgical progressive care system. The aim of this thesis is to develop a model of an information system which is capable of reducing the extent to which epistemological factors limit the effectiveness of resource allocation in this type of system. The main epistemological control-limiting factors which will be the focus of this dissertation are summarised below.

DETERMINATION OF URGENCY

Being able to more precisely determine the degree of urgency associated with the demand for healthcare resources prior to any allocation of those resources enhances control since it decreases the likelihood of having to allocate resources for a previously unrecognised urgent demand for those resources at the last minute.

DETERMINATION OF RESOURCING REQUIREMENTS

Being able to more precisely determine the resourcing requirements associated with a patient is important since it allows the allocation of resources to be made with a greater degree of confidence that the allocation will not have to be modified once the consumption of those resource have already begun. In a progressive care system, determination of resourcing requirements has two components, as follows:

DETERMINATION OF DURATION

That is, determining the length of time a resource needs to be allocated to a patient. With the bed-slot assumption, this amounts to determining the length of stay a patient will stay in each healthcare unit.

DETERMINATION OF ORDER

That is, determining the order in which the resources are to be allocated to a patient. With the bed-slot assumption, this amounts to determining the order in which the patient will require allocation of bed-slots in each healthcare unit.

According to the classification of healthcare resource-allocation decisions made above between proactive and reactive decisions, the aim of this thesis may be recast as the development of a model of an information system which increases the degree to which patient management is guided by proactive, rather than reactive, resource-allocation decisions.

In later chapters it will be argued that the proposed information system is able to increase the degree to which patient management is undertaken proactively, not necessarily through any improvement in the accuracy of the information involved, but rather through the way in which that information is processed, disseminated and utilised. That is, that the epistemological control-limiting factors are as much down to the information being managed ineffectively as it is to the accuracy of the information itself.

2.6. Summary

The objective of this chapter has been to represent the process of healthcare resource allocation as a control process whereby the cost-effectiveness of healthcare delivery may be increased through increasing the effectiveness of control over the allocation of healthcare resources. In order to achieve this, a mathematical model capable of demonstrating the relationship between the effectiveness of control over the allocation of resources and the cost-effectiveness of healthcare delivery was developed. This model was used to depict various scenarios according to the level of available resources, the level of demand for those resources, and the effectiveness of the control over the allocation of those resources.

Two basic assumptions were used in the support of the mathematical model. The first of these was the parametric assumption, which states that the distribution of demand for healthcare resources may be approximated as normal. The second assumption was the bed-slot assumption which states that resources may be modelled as a package of co-present resources. The bed-slot assumption was proposed as a simplifying assumption which required evidential support before it could be successfully utilised in a model of healthcare resource allocation. It was noted, however, that for those resources which could not be successfully integrated within the notion of a bed-slot, the demand for those resources could still be modelled using the same kind of model which was utilised in the modelling of bed-slots.

To demonstrate the relationship between the effectiveness of control over healthcare resource allocation and the cost-effectiveness of healthcare delivery, it was necessary to define how those two variables may be measured. The measure which was proposed for the effectiveness of control over healthcare resource allocation was the standard deviation of the distribution of the demand for healthcare resources which would be satisfied assuming an effectively unlimited supply of those resources. This distribution was referred to as the $f'(x)$ distribution. This distribution was also used in the derivation of the cost-effectiveness of healthcare delivery, which was composed of a measure for the efficiency of healthcare delivery and the effectiveness of healthcare delivery. The efficiency of healthcare delivery was defined as the resource utilisation rate; the effectiveness of healthcare delivery was defined as the proportion of demand for healthcare resources which was satisfied.

Having demonstrated the relationship between the effectiveness of control and the cost-effectiveness of healthcare delivery, those factors which limit the control over healthcare resource allocation were discussed. It was noted that before any attempt at increasing the effectiveness of control could be made, it is necessary to determine which control-limiting factors were able to be modified or not; that is, whether or not the control-limiting factor was an inherent feature of the demand for healthcare resources or the way in which that demand is satisfied. In this regard it was noted that an important distinction to be made was whether the control-limiting factors were epistemological in nature or not. Epistemological control-limiting factors are those whose origin lies in a lack of knowledge of the underlying system, rather than the underlying system itself. Two epistemological control-limiting factors, labelled the determination of urgency and the determination of resourcing requirements, were identified as being particularly important in the control over healthcare resource allocation in a progressive care system.

In subsequent chapters, a model of an information system will be developed which aims to increase the control over resource allocation by addressing the problem of epistemological control-limiting factors as they occur in a progressive care system. The progressive care system which will be used as the problem domain in the development of the information system is the surgical high-dependency environment of the Royal Brompton and Harefield NHS Trust in West London. The in-depth description and analysis of the operational aspects of this system will be the topic of the next chapter.

3. The Empirical Domain

The empirical domain for this study is the high-dependency progressive care system of Royal Brompton and Harefield NHS Trust (RBH) situated in London. The role of RBH for this study is to provide the empirical environment for the development of the current and proposed operational models in Chapter 5.

This chapter is divided into three main sections. The first section provides an examination of the core principles behind high-dependency medicine and its role in modern healthcare systems. The objective of this section is to provide the context in which the RBH High-Dependency Environment operates, which will be the subject of the second main section where a narrative operational model will be detailed which describes the current operations of RBH in terms of its organisational components and structure. The approach for achieving this will be to first discuss some of the fundamental principles of high-dependency care and the systems in place to deliver it. Following on this discussion is a description of each of the component healthcare units of the RBH high-dependency environment. Each unit will be described in terms of its typical level of resourcing, its patient intake and the relation with other units within the progressive-care system.

To supplement the description of each unit, summaries of various empirical studies which have been performed using RBH data will be presented, the main body of each study being included as appendices 8, 9 and 10 of this thesis.

The first of these studies is an examination of the relationship between patients' clinical and demographic characteristics and the required bed-slot allocation size in the adult intensive-care unit. The purpose of this study is to highlight some of the difficulties in allocating high-dependency resources that arise from the typical intensive-care patient.

The second empirical study will also be based on data collected from the adult intensive-care unit. The objective of the study is to empirically evaluate the bed-slot assumption that was defined in chapter 2. The database that will be used for this analysis contains records for the different interventions for each patient for each day. By classifying the different interventions according to the types of resources which are required to make the intervention, the bed-slot assumption may be tested by determining the degree to which different types of patients consume different quantities of resources. In this study a distinction is made between being justified to use the bed-slot assumption as a descriptive shorthand for a package of typically co-present resources, and being justified to use the bed-slot assumption as the basis of a system of resource allocation. It will be concluded in this study that the use of the bed-slot assumption as a descriptive shorthand is justified, but that the use of it as the basis of resource allocation should be considered as remaining an open question. It is shown, however, that it is reasonable – in the absence of any further evidence to the contrary – to adopt the bed-slot assumption in this manner as a working hypothesis.

The third and final empirical study is aimed at evaluating the current system of resource allocation in place at RBH. This study is divided into three separate analyses. The first analysis uses a data set of so-called monotonically allocated bed-slots in the RBH high-dependency environment. That is, each record corresponds to an operating room bed-slot which was both allocated to and consumed by a patient. The second analysis contrasts with the first in that it uses data of non-monotonically allocated operating room bed-slots. That is, each record corresponds to a high-dependency environment bed-slot which was allocated to a patient, though not necessarily consumed by that patient since the allocation could subsequently have been cancelled before consumption began. In this second set of data, in those cases where a bed-slot was allocated but not consumed, a control-limiting factor may be assumed to be the cause of the discrepancy. The most likely control-limiting factor in these cases is the cancellation of the existing allocation and the subsequent re-allocation (which would not be recorded in this database) of the bed-slot to a patient requiring urgent admission to the operating room. The two databases could therefore be thought of as representing the two contrasting scenarios of a situation with intentional resource allocation and a situation with actual resource allocation. Thus, a comparison of the two databases would be able to identify those control-limiting factors which inhibit the intentional resource allocation becoming realised. This comparison of the two databases will be the subject of the third analysis of Section 2. The purpose of this study is to identify control-limiting factors in resource allocation and on that basis assess the current system of resource allocation at RBH in terms of the degree to which those control-limiting factors which have been identified are in-principle surmountable through the re-engineering of the resource allocation process.

The final section of this chapter will be a narrative descriptive model of the current system of resource allocation in place at RBH. This model will be the main source of material for the subsequent development of the current operational model to be presented in Chapter 5.

3.1. Operational Overview

3.1.1. High-Dependency Care

A high-dependency environment comprises a number of healthcare units or wards within a hospital that are sufficiently well equipped to admit patients who are in need of a level of care that is in excess of that available in other units or wards within the hospital. The term 'dependency' refers to patients' dependency on particular types of resources; 'high-dependency' refers to an increased level of such dependency. For example, patients in regular hospital wards can be said to be dependent on a certain level of nursing care. Patients in high-dependency units also are dependent on a certain level of nursing care, but whereas a patient in a regular ward may have access to a fraction of a nurse's time, a patient in a high-dependency unit may have access to all of a nurse's time.

Monitoring resources are an interesting type of resource when comparing the situation in higher and lower dependency units. Patient monitoring plays a far more significant role in the high-dependency environment than elsewhere. There are various reasons why this should be so:

- **Physiological states tend to change more rapidly in high-dependency patients; and**
- **When they do, the change is more likely to require a faster response; and**
- **The resource inputs associated with the kinds of physiological states found in the high-dependency environment are expensive, in which case there are good economic reasons for knowing more precisely when and when not to apply such resources.**

These factors all result in a proliferation of information systems capital - including monitoring capital - in high-dependency environments that is not found in other areas of the hospital. In these terms, higher-dependency patients will tend to have a greater dependence on systems capital, in which case the boundaries between different levels of dependency may best be defined in terms of quantity of data generated per unit of time. One reason why this should be so is the case of so-called high-risk monitoring patients. For these patients, although in physiological terms their condition may not be serious at a particular moment it could nevertheless, and at some unpredictable time, undergo a rapid deterioration. Such patients have an obvious need for frequent and extensive monitoring. The same may be said of some pre-operative patients, whose condition needs to be monitored immediately preceding surgery to ensure the detection of any potential per-operative complications in sufficient time to plan for them in advance.

Every high-dependency environment will be different, and this difference will be a reflection of differences in case-mix, funding priorities, and clinical approaches to critical care medicine. There are some common factors however. Firstly, each operative high-dependency environment will be oriented around the operating theatres both clinically, financially, and often geographically also. Second, one would expect to find in each a recovery area and an intensive-care area, although the two roles may be combined within one unit.

Finally, there will be similarities in the flows of high-dependency patients around the system. Normally, for operative patients, the first high-dependency unit the patient enters is the operating room suite, followed by either the recovery room or the intensive-care unit, often followed by a lower-dependency care unit before eventual discharge to a regular ward. Such a typical flow round the system for an operative patient corresponds to changing levels of dependency. That is, there is initially an increase in dependency during the operative period. This decreases on completion of the operative procedure and on admission to either the recovery room or the intensive-care unit before final admission to a so-called intermediate-care unit or regular ward.

The functions of the recovery room and intensive-care units often overlap to a certain extent. In each case, on admission of a patient from theatre, the function is to recover the patient from the effects of the anaesthesia, stabilise the patient's physiology within an acceptable range, and to bring the level of pain down to an acceptable level. This, however, is the only function of the recovery room, and if it is unable to perform this role for a particular patient, then the patient is transferred to the intensive-care unit for more intensive long-term therapy. In some instances, there is no overlap in function between the recovery room and the intensive-care unit. In these cases all operating room discharges are admitted to the recovery room, where the patients are assessed for subsequent admission to either the intensive-care unit or some other lower-dependency unit.

The typical intensive-care unit itself serves a number of functions. The first is that already mentioned in relation to the recovery room. Secondly, in those cases where the patient's condition is too critical and/or the post-operative recovery period is expected to be too long-term for discharge via the recovery room, the intensive-care unit is able to provide the resources for such longer-term and higher-dependency care. Finally, intensive-care units often admit (sometimes exclusively admit) non-operative patients. That is, patients whose treatment is not expected to involve any (non-exploratory) surgical interventions. Non-operative patients' conditions often require extensive monitoring, and are often relatively long-term. They are usually admitted into the intensive-care unit from lower-dependency wards as the patient's condition deteriorates, or from other intensive-care units if the admitting hospital has a specialisation in the area of medicine that the discharging hospital lacks. In some cases a non-operative patient will eventually require some form of surgical intervention, in which case they will often be readmitted to the intensive-care unit as a post-operative patient. In other cases, the therapy provided may be purely palliative in nature with the focus being on the management of pain.

In many hospitals there are also intermediate-care units, also known as step-down units, which act as an interface between the higher dependency operating room suite, intensive-care unit and recovery room, and the lower dependency regular wards. The interface of the intermediate-care units may be for post-operative patients coming from the intensive-care unit or recovery room, or it may be for pre-operative or non-operative patients going into the operating room suite or the intensive-care unit. Whatever the case, the intermediate-care unit is able to offer a higher level of care than regular wards, though not as high a level found in the operating room suite, recovery room, or the intensive-care unit.

The overall design of the high-dependency environment will be primarily dependent on the number and case-mix of patients admitted. In general, it could be said that the number of patients affects the number of bed-slots and the case-mix affects how those bed-slots are distributed around the units comprising the high-dependency environment. For example, if the case-mix is comprised of a large proportion of non-operative patients, then the intensive-care unit would be expected to be relatively large, and the operating room suite to be relatively small. If the case-mix is comprised largely of operative cardiac patients, then one would expect both the operating room suite and the intensive-care unit to be relatively large, since cardiac procedures take a relatively long time, as does the post-operative recovery process. Conversely, if the operative workload was comprised largely of orthopaedic patients, then the intensive-care unit would be quite small, although the operating room suite would be of reasonable size.

A second factor prominent in the distribution of resources around high-dependency units is the clinical approach that has been adopted to certain aspects of the high-dependency care process. For example, as mentioned above, in some hospitals all post-operative patients are admitted to the recovery room. This would require a large recovery room, and also one that is well resourced in terms of quality of capital equipment and staff due to the increased severity of cases being admitted. Also, there are hospitals that have adopted a procedure of pre-operative monitoring within the intensive-care unit. It would be reasonable to infer that this would require an expanded intensive-care unit, although the reverse might actually be the case if unanticipated peri- and post-operative complications, which would otherwise have occurred due to a lack of information regarding aspects of the patient's physiological condition, could be anticipated.

Another obvious factor in the design of high-dependency environments is funding. If the money is not available to pay for a large operating room suite, for example, then a small one has to suffice. In general, increased levels of funding would tend to increase the size of the intensive-care unit, since this tends to be the most bottomless of bottomless pits within the healthcare sector. It is in the intensive-care unit where the changes in rationing of healthcare resources involve the largest quantities of money.

In the following sections there will be a discussion of how all of these factors combine to form the design of the empirical domain. The structure of the discussion will be to consider in turn each of the units comprising the RBH high-dependency environment, and noting in each case the relevant aspects of the case-mix necessary to consider in a model of resource allocation. Such relevant aspects will be, for example, the lengths of stays for different types of patient (i.e. the size of the required bed-slot allocation), the origin of admitted patients, and the destinations of discharged patients. After a brief description of RBH as a whole, the first unit to be considered is the operating room suite, followed by the recovery room, intensive-care units, and intermediate-care units respectively. In the final section of this chapter the problem of allocating resources for improving cost-effectiveness will be discussed in relation to the particular set-up at RBH.

3.2. RBH Overview

RBH is a publicly funded National Health Service (NHS) trust. The trust has two main sites, one located in South Kensington in Central London, the other located at Harefield, located just outside of London. Academically, RBH is associated with the medical school of Imperial College, University of London.

RBH is a specialist provider of cardiothoracic medicine (i.e., medicine relating to the organs of the upper chest cavity – heart and lungs). It has both secondary and tertiary referrals. Secondary referrals are those referrals made to RBH from a referring general practitioner (GP). Tertiary referrals are those referrals made to RBH from another, usually non-specialist, healthcare facility. Secondary referrals are typically for the kind of healthcare which could equally be provided by a non-teaching and non-specialist healthcare facility such as a district general hospital. In this regard, RBH is a provider of cardiothoracic medicine for two main populations of patients under the purchaser-provider contracts of the reformed NHS. The first population is that of the local regions of London in which RBH is located. The second population is farther afield in Bath in the west of England. Tertiary referrals to RBH could, in theory, come from any region within the UK or farther afield. Tertiary referrals differ from secondary referrals in the degree to which the specialist nature of RBH is more equipped in terms of resources and experience to deal with the particular healthcare needs of the patient than would be a general hospital. Thus, tertiary referrals typically involve relatively uncommon operative procedures such as heart or lung transplants or the redoing of procedures undertaken elsewhere.

Although RBH is an NHS Trust, and therefore is funded through the national healthcare budget, a small but significant proportion of its patients are self-funding. Self-funding patients are usually treated in the same units and by the same staff and other resources as NHS-funded patients. The exception to this is with regards to the provision of intermediate care which is often provided within a unit specifically intended for self-funding patients. Also, to prevent the provision of healthcare to self-funding patients interfering with the contractual obligations towards the NHS-funded patients, the operating rooms are set aside for self-funded patients on most Saturdays, with the rest of the week devoted to NHS-funded patients.

RBH is a healthcare provider of both adult and paediatric cardiothoracic medicine. Correspondingly, its Central London site has one intensive-care unit for adult admissions and one for paediatric admissions. The majority of patients are adults, although the paediatric admissions tend to require a disproportionately higher level of resourcing, and so is nevertheless an important consideration in modelling resource allocation.

The empirical domain of this study is the high-dependency environment of RBH's Central London site. This high-dependency environment is an example of a progressive-care system and consists of an operating room suite, two intensive-care units (one paediatric, one adult), a recovery room and various intermediate care units. Each of these units will be discussed in turn in later sections. Before this, however, there will be a discussion of the system of information management in RBH. This discussion will be divided into three sections: i) an overview of the main properties of high-dependency information systems; ii) an examination of the role that information systems play in the process of patient management; and iii) an overview of one of the most important information systems at RBH with respect to the objectives of this study, which is the Carevue system implemented in the operating room suite, the recovery room and both intensive-care units.

3.2.1. RBH High-Dependency Information Systems

Very broadly, there are two types of information involved in health-care: the clinical and the administrative. This distinction is workable to an extent, although becomes rather fuzzy when, for example, a clinical audit department will audit the functioning of an intensive-care unit in terms of patient outcome and resource utilisation, the results of which may then be used to improve the quality of care.

Similarly, only a broadly defined distinction may be made between human information systems (HISs) and computerised information systems (CISs). More often than not the inherent weaknesses involved in one type of system will be complemented by inherent strengths that are found in the other, in which case the ideal system involves both a human and a computer implementation. A good example of this is in clinical decision-making, where the human's capacity for pattern recognition tends to be less prone to make serious blunders and is generally superior, given the current programming techniques currently in fashion, than computerised diagnosis systems. Yet the task of collecting the data necessary for diagnosis is often the fortè of computer-controlled systems. This is especially true in the high-dependency environment, as will be discussed below.

Despite the vagueness in the categorisation, the distinction between HISs and CISs on the one hand, and clinical and administrative systems on the other is generally workable. The term "information system" itself is quite easy to define: "a system of communicating components each of whose function is to transfer and/or manipulate and/or extract and/or display data/information for a specified domain and purpose".

In the general healthcare environment there are a number of different information systems, ranging from the very localised and clearly defined systems that are in place, for example, in the scheduling of nurses in a particular unit, to hospital-wide systems such as the patient administrative system (PAS) that is found in all hospitals and are typically computerised to varying degrees.

In the high-dependency environment there is a much greater need for clinical information, both in terms of breadth (i.e., a greater amount of variables are recorded) as well as depth (i.e., an increased recording rate). There is also a need for increased precision and accuracy of measurement [AMB92], [BUT89], [MET95]. This increased need for both the quality and quantity of information in the high-dependency environment is a reflection of the typically faster rates of change in underlying physiological states, as well as the increased significance of such changes, both in terms of future resourcing requirements as well as healthcare outcome.

With regards to administrative information, the need is much the same in other, lower-dependency, units of the hospital. For example, the nurse scheduling system in the adult intensive-care unit is much the same as in a regular ward in terms of shift durations and experience-mix, although unlike in regular wards the system is computer-based as a decision-support system, implemented using a commercially available generic spreadsheet program.

One area in which the requirements for administrative information differ from regular wards is in the allocation of bed-slots. Operating room bed-slot allocation can only be an effective process if complemented by a corresponding system of allocation for bed-slots in interdependent units. Thus, the confirmation of an operating room bed-slot will be accompanied by a confirmation of a subsequent bed-slot in at least one of the post-operative units. For the transfer of patients between these post-operative units and subsequent units, a member of staff in each of those subsequent units is designated to co-ordinate the transfer. The co-ordination of operating room bed-slots and other bed-slots is undertaken by an operational manager called the master scheduler, in co-operation with the surgeons, the ward consultants, anaesthetists, and nursing staff. The co-ordination of patient transfers between other units is usually procedural - post-operative patients discharged from the adult intensive-care unit, for example, will normally be readmitted to the intermediate-care area within the ward that originally admitted them. For those patients for whom this procedure is not applicable, either because of a localised bed-slot shortage or an unexpected change in the patient's physiological condition, the process is co-ordinated at a daily meeting between the nursing representatives of the relevant units (the wards with intermediate-care areas, the adult intensive-care unit and the recovery room) and the master scheduler. A similar meeting also occurs daily for paediatric cases.

The resource allocation system is thus very much a human information system. The use of computers usually arises as representational tools used in the preparation of the resource allocation schedule which details which patients are allocated which bed-slots over the following working week.

A more detailed discussion of the resource allocation system in place at RBH will be provided in a later section. In the next section the principles of patient management are described, particularly as they relate to the information systems which support the patient management process.

3.2.2. Patient Management

As mentioned earlier, there is a greater need in the high-dependency environment for effective information systems. In terms of patient management, this implies various systems. Firstly, good communications are a prerequisite between the members of staff within a unit - in particular, between staff members on different shifts and between clinicians and nurses - to ensure consistency in patient care. Equally, there is a need for good communication between different units so as to provide consistency in the management of the patient despite the transfer between different units, through an extensive and precise communication of the patient's past, present, and expected physiological states. Both inter- and intra-unit communication is achieved through a mixture of human and computerised information systems.

In patient management at RBH, there are four identifiable components to the overall process that are normally identified, each component being undertaken by either human or non-human agents as follows:

DATA COLLECTION. The collection of raw (i.e. un-processed) physiological patient data by either machine sensors or human agents.

DATA INTERPRETATION. The analysis of the data in terms of the development of a diagnosis (analysis) and/or a plan for consequent therapy (interpretation). Machine data analysis/interpretation is able to interpret data only in some circumstances. The usual procedure is for the machine to provide a primary analysis and pass the result on for human data analysis/interpretation.

TREATMENT PLANNING.

TREATMENT IMPLEMENTATION. The implementation of therapy as determined by the data analysis/interpretation process, either by machine effectors or human agents.

Each of the above sub-processes are ordered chronologically, with the whole process being cyclical in nature ([CRA95]) (see Figure 3.01 below).

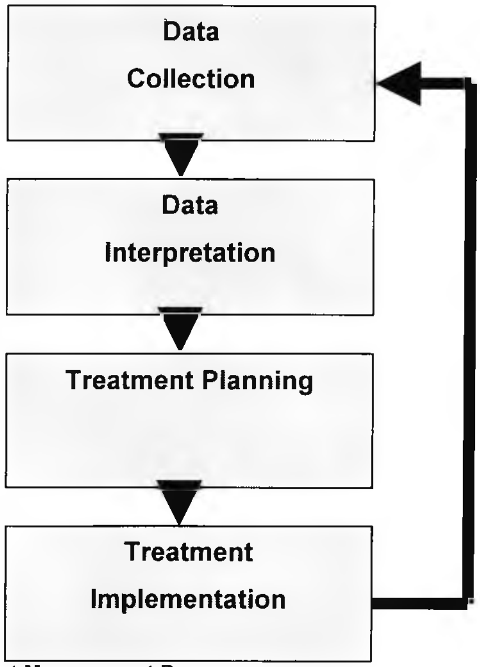

Figure 3.01 The Patient Management Process

The cyclical process of patient management is found in both lower and high-dependency health-care environments. However, in the high-dependency environment, because the underlying physiological state is both more severe and more dynamic, the cyclical process of patient management needs to reflect this by being undertaken at a much faster rate than would by the case in a lower-dependency environment. In other words, in the high-dependency environment, there is an increased need for an information system that can support such a rapid rate of processing, and further, whilst also continuing to be effective - in terms of cost, precision and reliability - at the processing itself.

Traditionally, the human element in the patient-management information system dominated every one of the above four sub-processes. However, for various reasons that will be discussed below, machines have begun to dominate the first two sub-processes, and are slowly making inroads in to the last two. This situation is reflected at RBH, where, in terms of the proportion of raw, uninterpreted data, the first two sub-processes are almost completely machine-based.

Although the process of data collection is dominated by machines there remains important elements of data which remain collected by humans. An example of such data might be the Glasgow Coma Score[3], which requires some neuropsychological tests unable, as yet, to be performed by machine. In the case of the analysis and interpretation of data, the situation will be reversed. In fact, although machines at RBH do a significant amount of analysis of data, it is only of a very routine arithmetical nature, and will usually only interpret data as a support tool, leaving the final interpretation to human agents. Those instances where machines make treatment planning decisions are rare and usually exclusive to certain aspects of routine and continuous processes that involve a degree of responsiveness to physiological condition, such as machine dialysis.

In the case of the storage and display of physiological data, it is difficult to apportion the workload accurately between either humans or machines, since the human's method of storage and display is necessarily a subjective phenomenon (excluding the use of paper-based records), and may well vary between individuals. It is probably true, however, that the type of physiological information stored and displayed by humans will be of a different emphasis and structure than that stored and displayed by machines. Specifically, the information stored and displayed by humans will often be either in the form of experience and skill that cannot easily be given a procedural linguistic encoding, or of a set of very general heuristics or laws involved in patient management.

At RBH, the machine-based storage and display of information is split between various monitors and a computer system known as Carevue[4] [CAR95]. Carevue also undertakes the bulk of machine-based data analysis, although it plays little role in the interpretation of data. In the next subsection, the Carevue system will be discussed in more detail, along with the rationale behind its introduction and the impact it has on patient management.

3.2.3. Carevue

Traditionally, information technology resources in the high-dependency environment served only as monitoring tools. However, recently there has been an increase in clinical information systems that also analyse the incoming data to varying degrees; thus, the extension from patient monitoring systems, to patient data management systems, or PDMSs. The Carevue system is an example of a

---

[3] A score which measures the degree of comatization

[4] Carevue 9000, Hewlett-Packard Company Ltd.

PDMS. It does not itself collect physiological data, nor does it continuously monitor physiological variables. Its purpose is to store, display and analyse physiological and other administrative data.

The Carevue system is a computer system operating from a Unix base, with various workstations situated throughout the high-dependency environment at RBH (see Table 3.01 below). Each workstation consists of a 17-inch screen with keyboard and mouse, along with a purpose-built trolley. A data-warehousing system is operated with inactive data being uploaded to an Oracle database, of which queries may be made for the purposes of clinical investigations.

Location	**Number**
Adult Intensive-Care Unit	20
Paediatric Intensive-Care Unit	7
Recovery	5
York HDU	2*
Theatres	5
Review Sites	1
Total	40

*Shared between two beds

Table 3.01. CareVue Terminal Locations

The Carevue system is integrated with various patient-monitoring systems, and will periodically collect data from such sources. It also allows data input direct from a keyboard/mouse by nursing or clinical staff. It has links to other hospital information systems - the patient administrative system (PAS), the pathology information system and the pharmacy system. These links enable, for example, orders to be sent to the pharmacy via the pharmacy system, pathology results to be accessed via the pathology system or patient administrative details to be accessed via PAS. The Carevue architecture is shown in Figure 3.02 below.

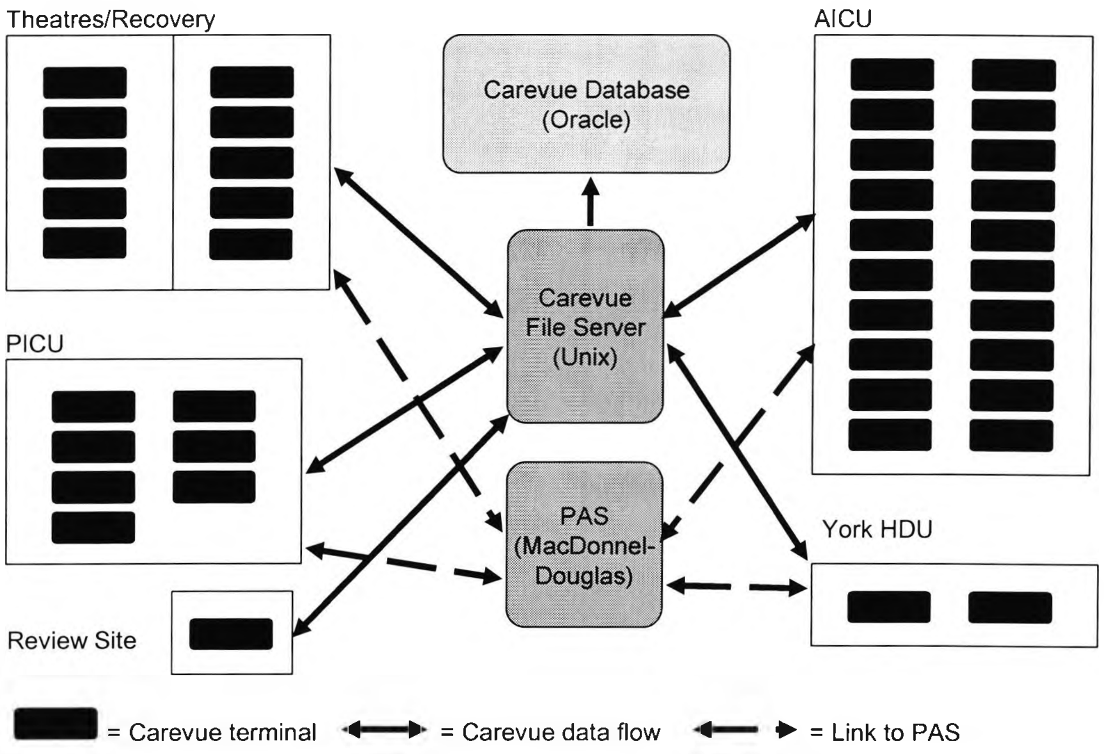

Figure 3.02. Carevue system architecture.

The level and content of inputted information in to a Carevue terminal depends very much on where the terminal is located. In common with all terminal sites, however, will be the recording of vital physiological statistics such as invasive blood pressure, oxygen saturation, etc. In OR, data will be collected regarding, for example, the nature and duration of the surgery from incision to surgery, the anaesthesia used, a free-text record of any complications in the surgery, the operating surgeon, etc.

Because OR and RR are on the same Carevue network, the data inputted in OR will also be available to RR who will continue to input information. RR will record statistics such as the level and need of mechanical ventilation and other statistics that are relevant in making the decision to discharge the patient to an HDU or intensive-care unit. RR staff are also able to use Carevue to estimate the time of admission of patients from OR by looking at the data for the patients currently undergoing surgery, which has obvious advantages in predicting the need for future bed-space.

In the intensive-care units the need for the management of data is the greatest since they generate the greatest amount of data. The quantity of data generated by a patient in intensive care tends to be greater than that generated in other units since there is both a relatively high rate of data collection, as well as a relatively large period of collection. This is in comparison to OR which, while collecting more data per unit time, will have a period of collection of only a few hours, compared to possibly days or weeks in an intensive-care unit.

Most of the data collected in AICU, apart from the physiological monitoring data, are represented within dynamic knowledge structures called intervention scores. Intervention scores are designed primarily as indicators of resource utilisation ([KEE83], [LEM94b], [TER94b], [CUL74], [REI97], [MOR97], [LIT88], [HAS93], [TER94a], [MIR91], [MAL92], [HOL93]). The basic structure is to list a selection of statistically significant clinical interventions and to weight each one in terms of the impact on resource utilisation. Each patient generates a scoring by counting each intervention that was implemented, multiplying by the weight of the intervention and summing to derive the total score.

In AICU, the scoring system used is Therapeutic Intervention Scoring System (TISS) ([KEE83], [CUL74]). TISS, as used at RBH, consists of 64 interventions ([SQU92a], [SQU92b]) and typically ranges from around 40-80 for patients in AICU. Apart from giving an accurate indication of resource utilisation, TISS also provides an indication of severity in terms of probability of mortality. In this case, although not the purpose for which it was primarily designed, TISS may also be used as a severity score. The underlying principles and utility of severity scores have been much discussed in the literature (for example [BAR96a], [KLE90], [KNA85a], [LEM94b], [TER94b], [TER94c], [WAT94a]), although as a guide to clinical practice their use has, in the author's opinion, still to be justified. Considered as a severity score, TISS generally performs quite badly in comparison with more developed purpose-built scoring systems such as the Acute Physiology and Chronic Health Evaluation II (APACHE II) ([KNA81], [KNA85a], [KNA91], [ROW94], [ROG94]), although there remains a close correlation between them. The interventions and weightings which go to make up TISS are given in Appendix 10.

To give a rough idea of the data generated by a typical AICU patient by TISS, consider the situation of the patient staying for 2 whole days in AICU. In that time the patient's TISS score will be calculated 6 times over (once every 8 hours). This generates 6 x 65 (6 scores multiplied by 64 intervention fields plus the total amount field) = 390 different data entries. Further, if we also consider the data from the monitors, which are recorded minute-to-minute, with - let's say - just three vital signs being measured, gives a total of 2880 x 3 (number of minutes in 2 days multiplied by number of vital signs) = 8640 more data entries.

Although before Carevue was introduced, the sampling rate for the recording of vital sign information was not as frequent as each minute, there was nevertheless a very large amount of data which had to be collected and recorded, all of which was done by pen and paper [BUT89]. The introduction of Carevue reduced this burden significantly, although most components of TISS remain unautomated, which gave nurses more time to spend caring for the patient.

In PICU a scoring system is also implemented in the Carevue data sheet to TISS in AICU. The scoring system used is Paediatric Risk of Mortality score (PRISM) ([POL88]) which, as the name suggests, is designed as a severity score rather than an intervention score. Nevertheless, it serves much the same purpose as TISS in AICU.

The data collected in York HDU is often simply the data extracted from the monitors. There is no scoring system used in the HDU (although a version of TISS specifically designed for the intermediate-care unit has been proposed ([CUL94]), which is a reflection of the lower level of dependency - and hence lower risk of mortality - associated with intermediate-care units.

3.2.4. The RBH Operating Room Suite (OR)

There are five fully equipped operating rooms at RBH, plus three catheter labs where minor cardiac operations are undertaken. All operating rooms are on one floor, along with the adult intensive-care unit and the recovery room (see Figure 3.03 below).

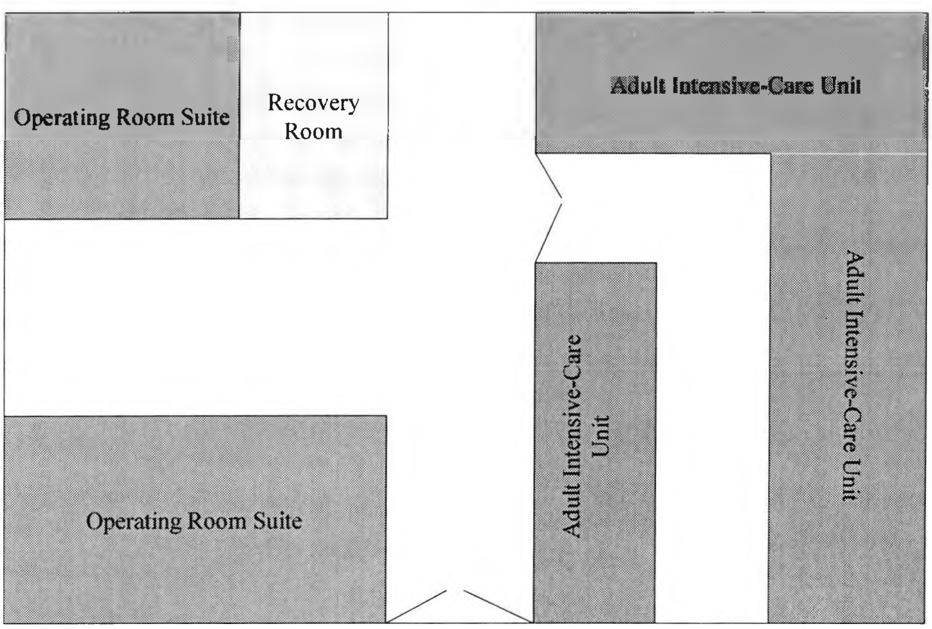

Figure 3.03. Schematic layout of the empirical domain.

With regards to the catheter labs, although they should be considered as high-dependency units, they will not be considered in this discussion since they operate almost autonomously from the rest of the empirical domain. There are, however, three areas of interaction that are worth mentioning at this point. First, some post-operative catheter lab patients are recovered in the recovery room. Second, if a problem occurs during the intraoperative period in the catheter lab, the patient may require emergency surgery in one of the fully equipped theatres. Finally, there are some shared resources between the two types of theatre. In particular, theatre technicians and anaesthetists are shared. Surgeons, however, are not shared. Consultant cardiologists do procedures undertaken in the catheter labs. Cardiac surgeons do procedures undertaken in the fully equipped theatres.

Of the five fully equipped operating rooms, four are fully operational, the fifth is used for one and a half days per week for elective (i.e., pre-planned) procedures, and the rest of the week is used only for emergency patients. Table 3.02 and Figure 3.04 below gives a breakdown of the total operative workload for the period 01/08/98 to 31/03/99 in terms of the number of procedures undertaken in each operating room.

OR Number	Number of Procedures	Percent	Valid Percent	Cumulative Percent
1	397	21.9	22.2	22.2
2	561	30.9	31.3	53.5
3	309	17	17.3	70.7
4	290	16	16.2	86.9
5	234	12.9	13.1	100
Total	1791	98.8	100	
Missing	22	1.2		
Total	1813	100		

Table 3.02. Number of procedures by operating room for the period 01/08/98 – 31/03/99

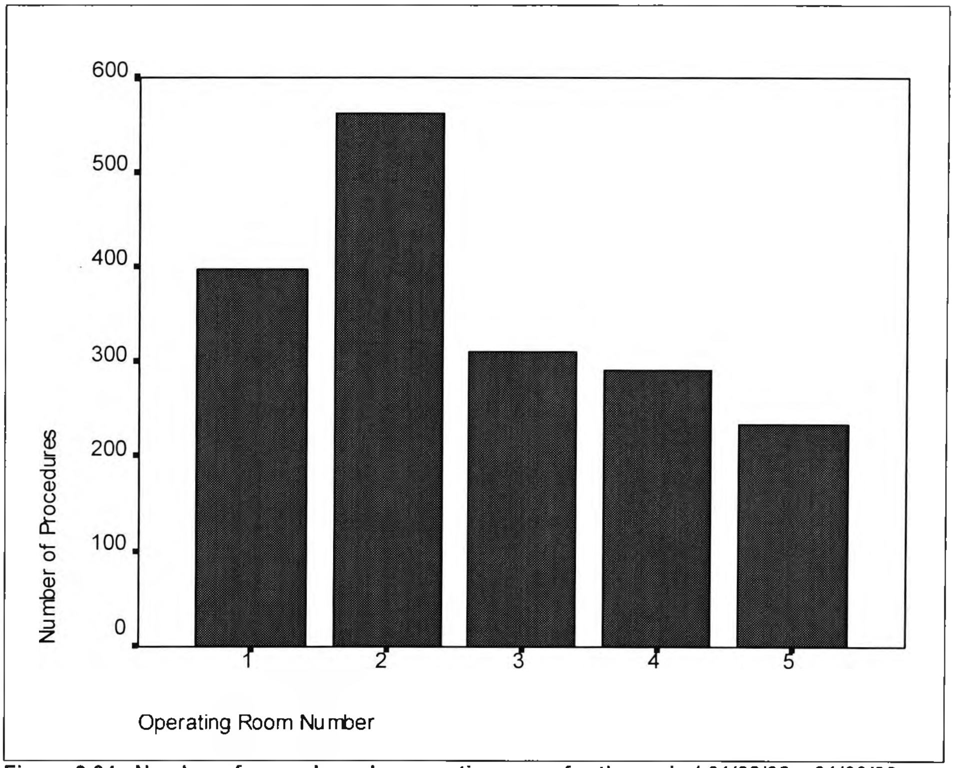

Figure 3.04. Number of procedures by operating room for the period 01/08/98 – 31/03/99

As may be seen from Table 3.02 and Figure 3.04, operating room 2 had the lion's share of procedures with 31.3% of the total; 5 had the least with 13.1%. The data shown in Table 3.02 and Figure 3.04 is only a very rough approximation to the actual workload, however. For example, a major cardiac procedure usually takes approximately twice as long to perform than does the typical thoracic procedure. This is an important consideration because, as can be seen from Table 3.02 and Figure 3.04 below, different operating rooms have different proportions of cardiac, thoracic and other procedures performed in them. This is a direct consequence of the block scheduling procedure of resource allocation employed at RBH, which will be discussed below.

OR Number	**PROCEDURE TYPE**			**Total**
	Cardiac	**Thoracic**	**Other**	
1	372	18	7	397
2	24	532	5	561
3	299	6	4	309
4	283	4	3	290
5	135	86	13	234
Total N	1113	646	32	1791
Total %	65.40%	32.40%	2.20%	100%

Table 3.03 Number of procedures of each type by operating room for the period 01/08/98 – 31/03/99

A way of more accurately measuring the operative workload for each operating room is therefore to weight each procedure type in terms of the quantity of operating room resources which need to be allocated for each type of procedure. This is relatively easy to accomplish given the bed-slot assumption and the block scheduling procedure employed for the allocation of operating room resources. Operating room time is divided into two blocks of time each day, one in the morning session and one in the afternoon, each one having a duration of approximately 4 hours. Given the high proportion of operating room resources being of the generic-and-fixed or generic-and-not-fixed types, it is reasonable to model each block of operating room time as representing a definite number of bed-slots. According to the current system of resource allocation in place at RBH, each block of operating room time is initially allocated either one or two procedures, depending on the type of procedure being performed. The appropriate number of bed-slots to associate with each time block is therefore two. If this schema is adopted, the time blocks and bed-slots allocated to each procedure type is as shown in Table 3.04 below.

Procedure Type	Typical OR Time Block Allocation	Corresponding OR Bed-slot Allocation
Cardiac:		
Major Cardiac	1	2
Minor Cardiac	1	1
Thoracic	1	1
Other	1	1

Table 3.04 Typical OR time block and corresponding OR Bed-slot allocations for each procedure type

Using the figures shown in Table 3.04 to measure OR workload, Table 3.05 and Figure 3.06 below show the total OR workload and mean OR workload per procedure (WPP), measured as the total number of bed-slots and the mean amount of bed-slots consumed for each procedure which is performed for each operating room, respectively.

OR Number	Number of Procedures	Total OR Workload	Mean OR WPP	Std. Deviation
1	397	684.9838	1.7254	0.4469
2	560	568.008	1.0143	0.1188
3	309	518.0076	1.6764	0.4686
4	290	544.997	1.8793	0.3263
5	233	330.9998	1.4206	0.4947
Total	1789	2647.0044	1.4796	0.4997

Table 3.05 Total OR workload and mean OR workload per procedure for each operating room

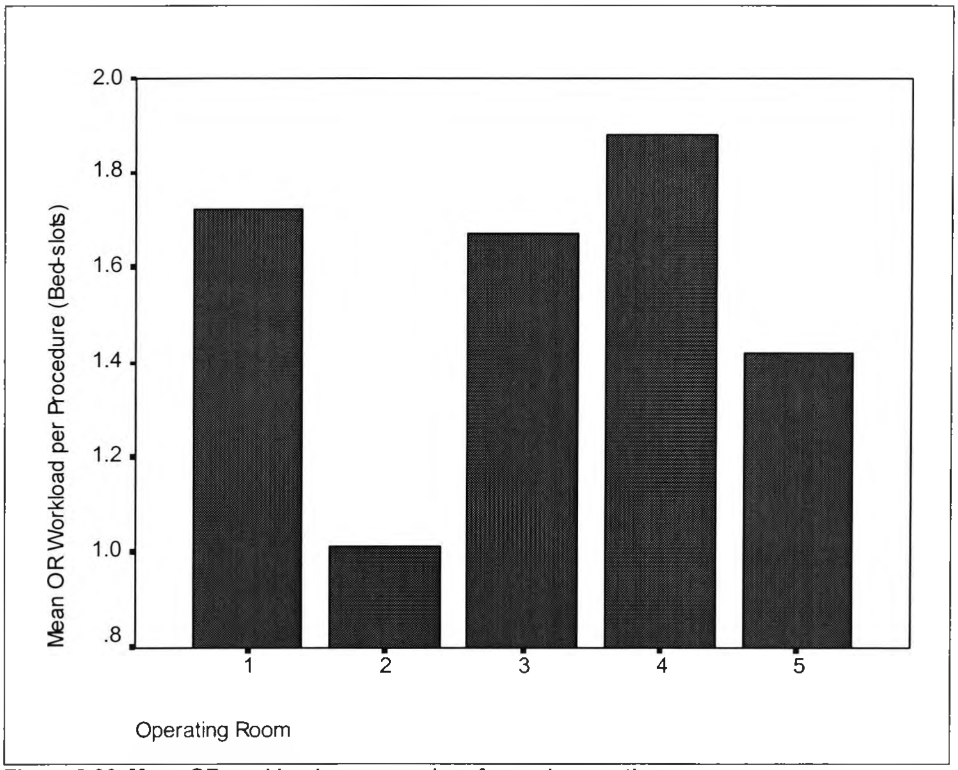

Figure 3.06. Mean OR workload per procedure for each operating room

Table 3.05 and Figure 3.06 shows that, although OR 2 has the largest workload, when measured as total number of procedures performed, the mean workload per procedure is relatively low when measured as number of bed-slots allocated to each procedure. This is as to be expected, since OR 2 is usually allocated thoracic procedures which typically consume only one bed-slot per procedure. This is contrasted with other operating rooms which are used predominantly for performing cardiac, and particularly major cardiac, procedures which typically require a double bed-slot allocation per procedure.

The heterogeneous makeup of the total operative workload is also reflected in the type of patient, as well as the type of procedure. Again, this is a consequence of the block scheduling system of OR resource allocation, with each block of OR time being allocated to surgeons whose specialisation is expressed both in terms of procedure type (i.e., cardiac or thoracic) and patient type (i.e., adult or paediatric). Table 3.06 and Figure 3.07 below demonstrates how this affects the makeup of the operative workload for each operating room.

OR Number	Number of Procedures	Proportion Paediatric	Std. Deviation
1	397	0.1209	0.3264
2	561	9.45E-02	0.2927
3	309	0.5146	0.5006
4	290	9.66E-02	0.2959

5	234	0.3974	0.4904
Total	1791	0.2127	0.4094

Table 3.06 Makeup of OR workload by proportion of procedures for paediatric patients

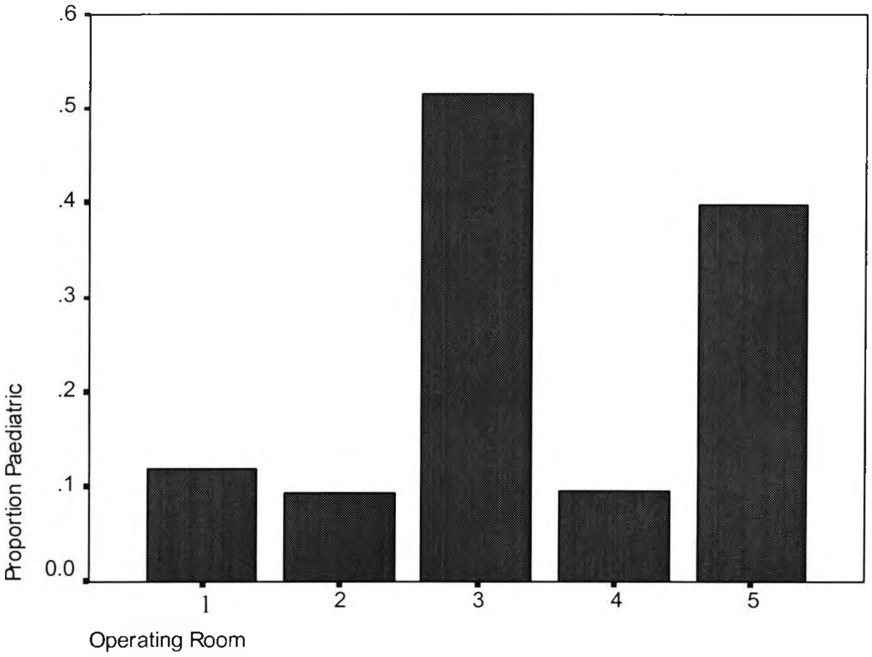

Figure 3.07. Makeup of OR workload by proportion of procedures for paediatric patients

With regards to the proportion of procedures which are performed as a matter of urgency as compared with those which are performed on a normal elective basis, there should be no reason for the proportion of urgent cases to be different between different operating rooms, except for a possible elevation in the proportion of urgent procedures undertaken in OR 5, due to the increased likelihood of it being unused at any given time and therefore more appropriate for use in performing urgent procedures. This hypothesis is not, however, supported by the data. The proportion of procedures for each operating room that were classified as urgent is shown Figure 3.08 below.

To support the hypothesis that there are significant between-operating room differences in the proportion of urgent procedures, a one-way ANOVA test was performed on the data and was confirmed ($p<0.05$). The reason why this should be so is unclear, although one possibility is that, because different surgeons are responsible for the treatment of slightly different types of patients, and because surgeons' time blocks tend to be within the same operating room, and that wherever possible urgent procedures are allocated within the time block of the surgeon responsible for the patient, that some surgeons may be responsible for types of patients who have an elevated risk of requiring an urgent allocation of operating room resources.

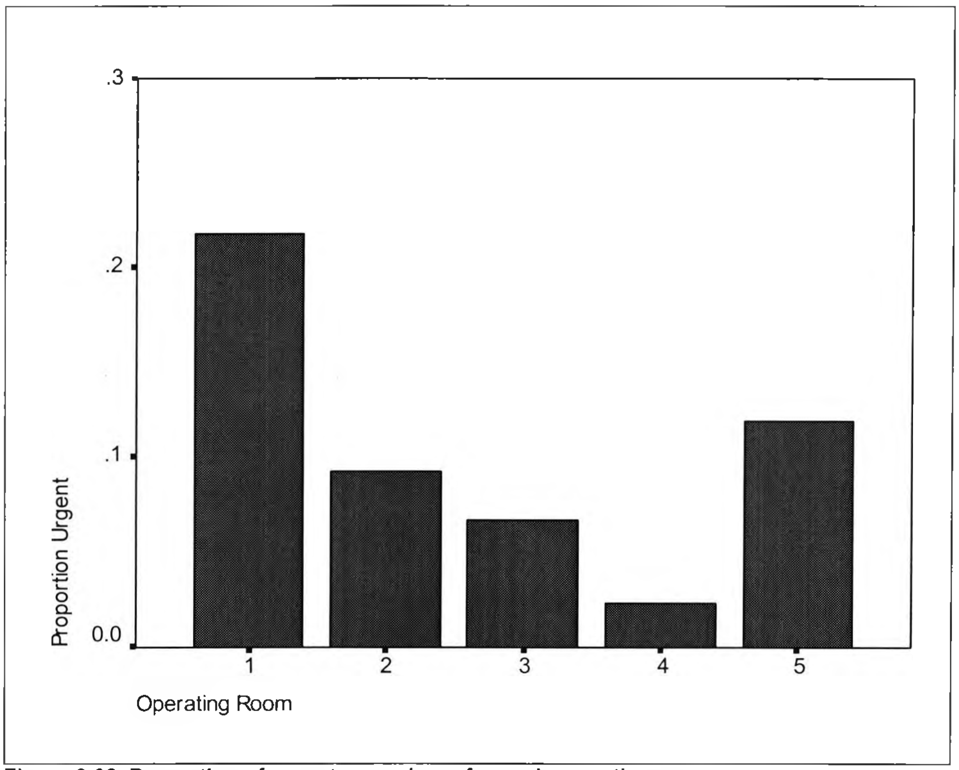

Figure 3.08. Proportion of urgent procedures for each operating room

With regards to the pre-operative locations that patients allocated operating room bed-slots come from, most patients came from an adult ward within the hospital with 74.1% of the total. Only 5 patients (0.3%) came from the recovery room Figure 3.09 below show the proportion of patients allocated operating room bed-slots coming from each location.

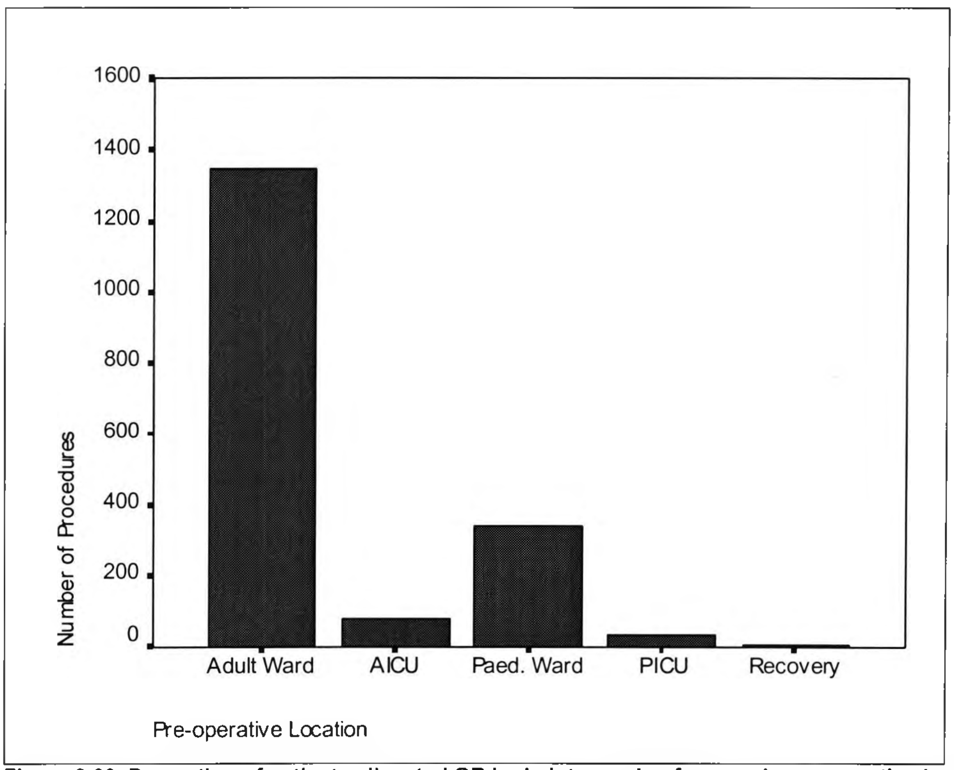

Figure 3.09. Proportion of patients allocated OR bed-slots coming from each pre-operative location.

With regards to the post-operative locations of patients coming from the operating room suite, the lion's share went to the adult intensive-care unit (AICU) with 35.0% of the total. Adult wards were close behind with 32.8% of the total. However, 11.4% of cases were recorded as having the recovery room as a post-operative location, most of which would finally end up in an adult ward. Figure 3.10 below shows the distribution of procedures between the different post-operative locations.

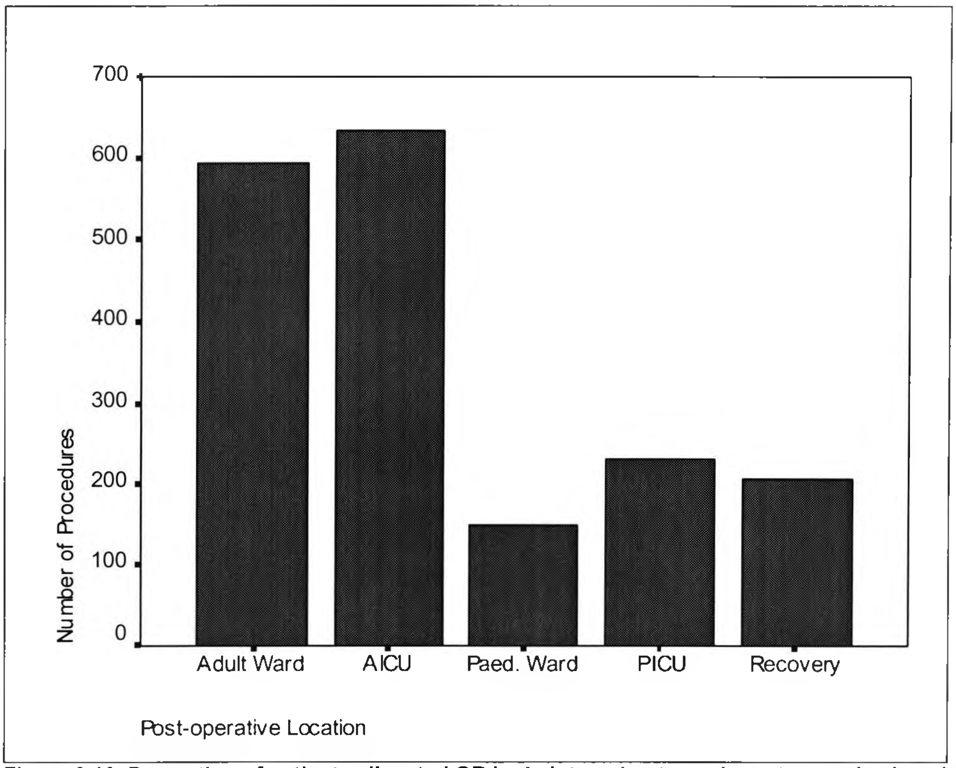

Figure 3.10. Proportion of patients allocated OR bed-slots going to each post-operative location.

3.2.5. The Recovery Room (RR)

The RBH recovery room (RR) has a maximum of 5 bed-sots at any one time, although if more are required it is possible to make use of an area within AICU which is situated close by (see Figure 3.03 above).

The function of RR is to stabilise post-operative patients before admission to the wards. This involves ensuring that the patient experiences an acceptable level of pain, is able to breathe without mechanical assistance, and more generally, whose condition will not deteriorate once transferred to either a lower dependence area of the hospital. If any of these criteria are not satisfied, then the patient is transferred to an intensive-care unit. Transfer from RR to an intensive-care unit is not a common outcome for most classes of patients admitted to RR. The exception to this is Cardiac Major patients that pass through RR, more of which will be said later. The following extract summarises the function and philosophy of RR:

The aim of the RR unit is to offer a safe environment in which to recover patients following surgery and procedures involving an anaesthetic maximising all available resources.

Prior to leaving the area, patients are safely recovered with our aim to achieve a pain level which is acceptable to the individual.

Care is provided by a professional, proficient practitioner .. and is delivered in an intuitive, reflective and research based manner using the named nurse approach

RR is staffed at all times that theatres are open. Additionally, an overnight service is provided Monday to Thursday, which caters for a class of Cardiac Major patients classified as overnight recovery patients (see below).

In terms of patient-mix, RR deals with a large class of operative patients, including both thoracic and cardiac, paediatric and adult. As a very general rule, all patients discharged from OR to RR are those that have undergone thoracic operative procedures or minor cardiac procedures. An exception to this general rule are so-called fast-track and overnight recovery patients which together form a significant proportion of all cardiac major patients. All fast-track and overnight recovery patients undergo the same surgical procedure[5], known as Coronary Artery Bypass Graft (CABG)[6]. CABG is a very common and standardised procedure, thus allowing surgeons to safely send their patients via RR rather than the adult intensive-care unit because of the experience that they, both as a community and individually, have acquired in dealing with such cases, and despite the procedure being major cardiac surgery.

With the increasing use of RR as a post-operative destination for CABG patients, it now deals with a larger proportion of such patients than AICU and represents a large proportion of its total workload. CABG patients which are admitted to AICU are either failed fast-track or overnight recovery patients, or those for whom the surgeon feels the level of dependency or length of stay would be too great for RR. The decision as to whether or not a patient requiring CABG is able to be classified as fast-track/overnight recovery or requiring post-operative admission to AICU is taken by the surgeon

[5] Although there are occasionally cases of aortic valve replacement (AVR) which are fast tracked. There has also been one case of mitral valve replacement (MVR) fast tracked, although this was an exceptional case.

[6] This excludes redo CABG procedures.

responsible according to a list of criteria which the patient must satisfy. This list is shown here as Figure 3.11 below.

Criteria for Cardiac Post-Operative Patients Admission to Recovery Room

Pre-Operative Criteria

1. Aged between 12 months and 70 years.
2. Left ventricular ejection fraction should exceed 0.3.
3. No serious pre-existing lung disease.
4. Normal liver function.
5. Normal renal function.
6. Normal coagulation.
7. No previous cerebral vascular event.
8. If there has been a myocardial infarction within the previous month the LVEF should exceed 0.5.
9. No recent alcohol or drug abuse.
10. Obesity excludes (i.e. >20% ideal body weight).
11. No insulin dependent diabetics.
12. Systemic hypertension must be controlled pre-operatively.
13. Electrophysiological surgery excludes.

Peri-Operative Criteria

1. Total anaesthetic and operation time must be less than 5 hours.
2. Cross clamp time of less than 75 minutes and bypass time of less than 120 minutes.

Figure 3.11. List of criteria for classification as fast-track/overnight recovery.

Of particular importance in the list of criteria shown in Figure 3.11 above is the inclusion of peri-operative criteria. This represents difficulties from a resource allocation perspective as it implies that a decision to allocate a CABG patient a RR bed-slot may subsequently be rescinded during the peri-operative period and instead an AICU bed-slot.

Table 3.08 below shows the breakdown of all admissions to RR in the period 1/8/94-31/5/97. As can be seen, Cardiac Major is the largest category of patients.

Operative Procedure	Total Number	% of Total	Mean Number per Day
Angio Procedure	1135	24.50	1.60
Cardioversion	1113	24.02	1.57
Fast Track Total	781	16.86	1.10
Other Major	490	10.58	0.69
Other Minor	325	7.01	0.46
Over Night Recovery	223	4.81	0.31
Portacatheter Procedure	105	2.27	0.15
Thoracic Major	104	2.24	0.15
Thoracic Minor	98	2.12	0.14
Thoracic Paediatric	97	2.09	0.14

Vascular Major	78	1.68	0.11
Vascular Minor	57	1.23	0.08
Wounds	27	0.58	0.04
Total	**4633**	**100.00**	

Table 3.08. Breakdown of all admissions to RR in the period 1/8/94-31/5/97

As mentioned before, not all of the patients scheduled to be fast-track or overnight recovery cases are actually able to go through RR and need instead an AICU bed-slot allocation. If this reclassification occurs when the patient has already been admitted to RR, it may be because RR needs the bed-slot for the next day's patients, because RR staff have been unable to 'wean' the patient from mechanical ventilation, or because the patient's condition has generally deteriorated/failed to stabilise in some other respect. A breakdown of Cardiac Major is given in Table 3.09 below, including the proportion (11.49%) of fast-tracks cases which are re-directed into AICU. Because overnight recovery cases only came into existence on 9 January 1996 (before then RR closed each night), the figure for these cases is disproportionately low. The figure in brackets gives the pro rata quantity.

Fast Track Success Status	**Total**	**% of Total**	**Mean Number per Day**
Unsuccessful Fast Track Total	130	11.49	0.18
Successful Fast Track Total	1001	88.51	1.41
Fast Track Total	1131	100.00	1.60

Table 3.09. Breakdown of Cardiac Major category admissions to RR

Of all the categories of patients in Table 3.08, all will be admitted to RR direct from OR, with the exception of most cardiac minor cases. These cases will come from one of the catheter labs, having undergone procedures such as angioplasty or other minor operations involving catheters. Many of the other cases undertaken in the catheter labs are treated as day cases and involve exploratory procedures to determine the need for major cardiac surgery. For the majority of cases admitted to the catheter labs, there will be no need for subsequent admission to RR. Table 3.10 below shows a breakdown of total RR workload according to whether the workload originates from patients being discharged from either OR or the catheter labs (CL). Workload is measured by scoring each patient admitted to RR according to the rule of scoring every patient as representing 1 unit of workload with the exception of cardiac major patients which have a score of 4 which represents the approximate ratio of recovery times for the two types of patient.

Originating Workload	Total	% of Total	Mean Workload per Day	Std. Deviation
CL-originating Recovery Workload	644	7.88	0.91	0.93
OR-originating Recovery Workload	7529	92.12	10.62	4.46
Total Recovery Workload	8173	100.00	11.53	4.57

Figure 3.10. Breakdown of RR workload according to origin.

With regards to the destinations of patients discharged from RR, and excluding the case of failed fast-tracks, the normal route would be to one of the post-operative wards, and in the more immediate post-operative period, to one of the intermediate-care units situated within the post-operative wards.

There are four wards designated as catering for post-operative patients, whether being admitted from RR or an intensive-care unit. Three of these wards are for adult cases (Elizabeth, Alexandra and Reginald Wilson wards), the other specifically for paediatric cases (Rose ward). Elizabeth is intended to be the destination ward for NHS adult thoracic patients, Alexandra ward is for NHS adult cardiac patients, and Reginald Wilson ward is for private adult patients, irrespective of surgery. All of the adult wards have an intermediate-care unit. There is also a paediatric intermediate-care unit that has recently been established, although it is situated within PICU rather than Rose ward. The intermediate-care units will be the topic of discussion later in this report.

Unfortunately, the data needed to establish the exact proportion of patients for each category going to each of the destination wards after discharge from RR is currently not available. However, in the majority of cases, the destination ward will be the one which is intended to cater for that type of patient. This will be particularly the case for paediatric discharges from RR - the vast majority will go either to the newly established paediatric intermediate-care unit, or to Rose ward. In the adult cases, the situation is slightly less clear-cut. For example, if there is no bed available in either Elizabeth or Alexandra, a patient might go to Reginald Wilson, or may be one of the medical wards - York (which also has an intermediate-care unit) or Paul Wood (which doesn't).

The lengths of stay for patients in RR vary considerably for each type of case. At one extreme are overnight RR cases, which could in theory stay for a maximum of around 21 hours, assuming that the procedure was admitted to theatre at the earliest time, and the duration of the procedure from anaesthesia to discharge was 3 hours. Second are (successful) fast-track cases, which typically have a recovery time of around 8 hours. All other cases will have recovery times substantially less, typically between 1-2 hours. Unfortunately, the exact lengths of stay are not available on a database, and in the case of overnight recovery cases, the measure has little significance since duration will often be dependent on the time of day at which the patient was admitted.

3.2.6. The Intensive-Care Units (AICU/PICU)

The function of the intensive-care units (AICU and PICU) at RBH is to be providers of critical care medicine. They each serve both operative patients discharged from, or to be admitted to, OR, and also non-operative patients admitted either as tertiary referrals from other hospitals, or admitted from within an internal hospital ward. In either case, patients need the level of monitoring and therapy that, outside of OR, may only be effectively provided within intensive-care units. The critical care medicine itself, which is provided in intensive-care units, has been defined by the following statement[7]:

A multidisciplinary and multi professional medical/nursing field concerned with patients who have sustained or are at risk of sustaining acutely life-threatening disease or injury. These conditions necessitate prolonged minute-to-minute therapy or observation in an intensive care unit (ICU) which is capable of providing a high level of intensive therapy in terms of quality and immediacy.

This definition is broadly correct, although in recent years there has been a tendency to move patients whose need for critical-care medicine is defined by a dependency on monitoring rather than therapy out of the intensive-care unit and into less expensive intermediate-care or coronary-care units. Such patients, referred to by Knaus [KNA93], as "low-risk monitoring patients", may then be moved into the intensive-care unit if their condition worsens, requiring more intensive therapeutic intervention. Alternatively, they may be moved to a lower-dependency (in terms of monitoring at least) area of the hospital if their condition improves.

The two intensive-care units at RBH both serve a mixture – operative and non-operative - of patients, although in each case the majority will be post-operative patients discharged from OR. In terms of age range, there is very little overlap between the paediatric and adult intensive-care units. There is also little overlap in terms of medical diagnosis between the two case-mixes. Paediatric cases will be almost exclusively congenital conditions, adult cases almost exclusively non congenital. This impacts

---

[7] National Institute of Health Consensus Development Conference Summary, 1983. In Parillo, J E, (ed.), Critical Decisions: key issues in the recovery of the critically ill. Toronto, PA: BC Decker, 1988:125. Quoted from Mallick, R et al, *The Intensive Care Unit Medical Director as Manager*, Medical Care, 33:6 (1995) p.611.

on both length of stay as well as predictability of length of stay in intensive care; congenital cases tend to require longer periods in intensive care, and are also less predictable before admission[8].

The AICU has a maximum capacity of 20 bed-slots at any one time, although only 15 of these are financed by the hospital trust. Thus, although bed-occupancy may exceed 15 at a particular time in the AICU, a level of occupancy less than 15 must compensate this at some other time in order to balance the books. A similar situation is also true for PICU - the maximum capacity is 11 beds, only 7 of which are financed. The bed-slots financed by the hospital trust in each unit are designated only for NHS/GP Fundholder patients. Private patients will finance their own bed-slots in either unit, in which case the financed level of occupancy may be in excess of the level financed by the hospital trust alone.

Apart from the differences mentioned between AICU and PICU, in terms of resource allocation they face very similar problems. Thus, to avoid unnecessary repetition of many points, attention will be focussed on the AICU rather than the PICU. In what follows 3 studies of the AICU is presented by way.

AICU CHRONICITY STUDY

The purpose of this study is to examine the relationship between the consumption of AICU bed-slots and the characteristics of the patient population which consumes them. This is an important relationship in resource allocation because of the skewed nature of the distribution for bed-slot allocation sizes and the difficulty in being able to predict the bed-slot allocation requirements of those patients who require chronic intensive-care against those who do not. In this respect, a measure of chronicity is introduced into the analysis according to the following rule: If AICU bed-slot allocation is greater than 2, it is classified as chronic, otherwise non-chronic.

There were three independent variables considered in the study: Sex, Operative Category and Previous Bed-Slot Allocation. AICU Outcome was also considered retrospectively in analyzing the relationship between outcome (dead or alive) and chronicity.

There were significant relationships between Operative Category and Previous Bed-Slot Allocation and AICU chronicity, but not between Sex and AICU chronicity. As expected, there was also a strong

---

[8] The lack of predictability is also true for intraoperative duration, with some paediatric cases consuming a whole day of theatre-time.

relationship between AICU chronicity and AICU outcome, with 21.8% of patients discharged alive from AICU have a chronic stay, compared with 68.2% of patients discharged dead having a chronic stay.

The full results and discussion of the AICU Chronicity Study is included in Appendix 9.

AICU Bed-slot Assumption Validation Study

The objective of this study is to examine the extent to which the Bed-slot Assumption referred to in the previous chapter is validated by the available data. The Bed-slot Assumption states that the adoption of the notion of a bed-slot is a valid simplification for use in resource allocation. As it stands, the notion of validity being used here is in need of further clarification. In order to determine the extent to which the assumption is validated by the data it needs to be re-cast in terms of testable hypotheses. Four such hypotheses may be identified, as follows:

1. **The Predictability Hypothesis**. Individual non-generic resources or different types of bed-slot whose individual allocation sizes are not predictable in advance of consumption within an acceptable degree of accuracy for different types of patient should be modelled as generic resources, either individually or as components of a bed-slot type in the case of individual non-generic resources, or in the case of types of bed-slot as instances of a more generic type of bed-slot; and

2. **The –G&F Proportion Hypothesis**. The –G&F resources whose individual allocation sizes are predictable in advance of consumption within an acceptable degree of accuracy do not constitute an excessively large proportion of the total resource consumption; and

3. **The Variance Hypothesis**. The overall level at which those resources which are modelled as components of the bed-slot are consumed at the population level is within an acceptable degree of variance within different categories of patients, or that such variance may not be reduced through the categorisation of patients into different categories where the categorisation of patients into those categories may be made in advance of consumption within an acceptable degree of accuracy; and

4. **The Difference Hypothesis**. There is no significant difference in the overall level of consumption of those resources which are modelled as components of the bed-slot between different categories of patients, or that where there is a difference between different

categories, the categorisation of patients into those categories may be made in advance of consumption within an acceptable degree of accuracy.

The first two of these hypotheses refer to ability to identify the existence of a bed-slot type. That is, for any given healthcare facility, if both of these hypotheses may be shown to be supported, then there is a type of bed-slot that can be validated by the data. It does not show, however, that this bed-slot type should be used as the basis of resource allocation.

The last two of these hypotheses refer to a specific application of the notion of the bed-slot. That is, for any given healthcare facility, if both of these hypotheses may be shown to be supported by the available data, then any bed-slot(s) which have already been identified may not effectively be categorised further into subtypes of bed-slot. Conversely, if the hypotheses are not supported by the available data, then there are subtypes of bed-slots which can be identified and consequently result in improved healthcare resource allocation.

The first hypothesis was confirmed analytically by arguing for the conclusion that a resource should be considered generic only if the knowledge regarding its consumption is generic, rather than the consumption itself; a resource is non-generic only if the knowledge regarding its consumption is non-generic. Thus, for example, the consumption of a mechanical ventilator may be non-generic insofar that some patients will consume the resource and others not consume it. However, for each patient, the knowledge of whether or not that patient will consume a mechanical ventilator may be generic insofar that for each patient it will not be known if the patient will consume a mechanical ventilator.

The second hypothesis was supported by categorising each component of the Therapeutic Intervention Scoring System (TISS) in terms of whether or not it implies consumption of generic or non-generic healthcare resources. More specifically, each resource was classified according to whether or not it implied the consumption of high-cost capital equipment, and also whether its consumption rate per patient was between the range of 0.05 and 0.80.

From an analysis of the TISS data it was seen that there are 8 –G&F TISS components that a) implied the consumption of high-cost capital equipment and b) had a consumption rate between 0.05 and 0.80 per patient. This represented a percentage of 11% of the total number of TISS components (74), which is less than the threshold of 25% necessary for the data to be considered inconsistent with the –G&F Proportion Hypothesis.

The third and fourth hypotheses were confirmed through a statistical analysis of the mean and variance of the incidences of the TISS components. In each of these cases three aggregate measures of TISS were derived. The first of these was the TISS Score as defined by Lemeshow et al ([LEM94b]); the second was a count of the total number of TISS interventions; the third the ratio between the two scores, effectively measuring the mean TISS weighting per TISS intervention.

For both the mean and variance of these aggregate measures and determining whether or not an alternative sub classification of bed-slots could be derived which was predictable in advance of consumption, a cluster analysis was performed on the data to derive two clusters of patients – those with high TISS aggregate scores and those with low TISS aggregate scores.

Although the last two hypotheses were not confirmed conclusively, it was argued in the study that a rolling system of prediction of consumption could be utilised which made daily predictions of resource consumption requirements which would improve the accuracy of prediction. The use of the TISS clusters to argue for a further sub-classification of bed-slots failed on the basis that the cluster could not be predicted prospectively, and thus the hypotheses was supported in the case of the RBH AICU.

The full results and discussion of the study is included as Appendix 9.

Resource Allocation Evaluation Study

The purpose of this study was to develop and use a method for objectively and statistically evaluating the patient scheduling process currently in place at RBH.

The method was based on two fundamental distinctions. The first is between what shall be called a system artefact, and what shall be called a system constraint. Both of these types of properties are control-limiting insofar that they represent limitations in the degree to which the process of patient scheduling may be controlled, and hence that system performance is optimised relative to one or more performance variables. However, system artefacts differ from system constraints in that the latter are necessary and permanent features of the system; system artefacts are contingent or transient artefacts in the operation of the system. Of course, what one person considers to be a necessary feature of a system may be different to what another considers necessary, and so this distinction is to some extent subjective, although for the purposes of this paper it nonetheless represents a workable and meaningful distinction.

A second distinction is made between two different types of system artefact. This distinction is formed as a cognitive distinction, and depends on whether or not those artefacts arise from a cognitive deficit on behalf of the control system that may be identified by the absence of any foreknowledge that such an artefact would occur.

On the basis of this distinction, an experimental method was developed that both identified each type of system artefact, and subsequently attempted to quantify their impact on overall system performance. The method was tested using data from the RBH master scheduler for the allocation of bed-slots in the operating room suite, recovery room and intensive-care units.

In the study the following performance variables were used:

Variable	Type	Range
AICU Bed-slot Allocation Status	Categoric	(Allocated, Not Allocated)
Fast-track Status	Categoric	(Fast-track, Non Fast-track)
Number AICU Bed-slots Allocated per Day	Numeric	N
Number Fast-tracks per Day	Numeric	N
Number Paediatrics per Day	Numeric	N
Number PICU Bed-slots Allocated per Day	Numeric	N
Number Urgent Bed-slots Allocated per Day*	Numeric	N
Operating Room	Categoric	(1, 2, 3, 4, 5)
Operative Category	Categoric	(Cardiac, Thoracic, Other)
OR Workload per Day	Numeric	R
OR Workload per Procedure	Numeric	R
Patient Type	Categoric	(Adult, Paediatric)
PICU Bed-slot Allocation Status	Categoric	(Allocated, Not Allocated)
Post-operative Location	Categoric	(Adult Ward, Paediatric Ward, RR, AICU, PICU)
Pre-operative Location	Categoric	(Adult Ward, Paediatric Ward, RR, AICU, PICU)
RR Workload per Day	Numeric	R
RR Workload per Procedure	Numeric	R
Urgency*	Categoric	(Urgent, Non-Urgent)

* Variable not included in non-monotonic dataset.

Table 3.11. Variables included in the Resource Allocation Evaluation Study

For each performance variable in the Table 3.11, the hypothesis was tested as to whether or not there was a significant difference in the value of the variable between a) the different months of the study period and b) between the days of the week. This hypothesis was tested for mean values using ANOVA and for frequencies using Chi-squared statistical tests for a dataset comprising monotonic bed-slot allocations, as well as another dataset comprising non-monotonic bed-slot allocations. The exception was those variables which occurred only in the monotonic dataset.

A comparison of means or variable present in both datasets between the two datasets was also made using ANOVA and a comparison of frequencies made using Chi-squared.

The following summary of results is broken down into those results for a) non-monotonic bed-slot allocations; b) monotonic bed-slot allocations, and c) a comparison of monotonic and non-monotonic bed-slot allocations.

Scheduling Status: Non-monotonic bed-slot allocation

The summary of results for each variable tested from the data set of non-monotonic bed-slot allocations is shown in Table 3.12 below. The first column lists the dependent variables which were tested; the second and third columns list the level of significance for whether or not variation in the corresponding independent variable could have occurred by chance or not over the months of the study period or amongst the different days of the week.

Dependent Variable	p for Month	p for Day of the Week
AICU Bed-slot Allocation Status	p<0.05	p<0.01
Fast-track Status	NS	p<0.01
Number AICU Bed-slots Allocated per Day	p<0.01	p<0.01
Number Fast-tracks per Day	NS	p<0.01
Number Paediatrics per Day	NS	p<0.01
Number PICU Bed-slots Allocated per Day	NS	p<0.01
Operative Category	NS	p<0.01
OR Workload per Day	NS	p<0.01
OR Workload per Procedure	NS	p<0.01
Patient Type	NS*	p<0.01
PICU Bed-slot Allocation Status	NS	NS
RR Workload per Day	NS*	p<0.01
RR Workload per Procedure	NS*	p<0.01

'NS' = not significant; *trend towards significance (0.1 > p >0.05).

Table 3.12. Summary of results from the non-monotonic bed-slot allocations dataset

Scheduling Status: Monotonic bed-slot allocation

The summary of results for each variable tested from the data set of monotonic bed-slot allocations is shown in Table 3.13 below. The first column lists the dependent variables which were tested; the second and third columns list the level of significance for whether or not variation in the corresponding independent variable could have occurred by chance or not over the months of the study period or amongst the different days of the week.

Dependent Variable	p for Month	P for Day of the Week
AICU Bed-slot Allocation Status	p<0.01	NS
Fast-track Status	p<0.01	P<0.01
Number AICU Bed-slots Allocated per Day	p<0.01	NS*

Number Fast-tracks per Day	p<0.01	P<0.01
Number Paediatrics per Day	NS	NS*
Number PICU Bed-slots Allocated per Day	NS	P<0.01
Operative Category	NS	P<0.05
OR Workload per Day	p<0.05	P<0.01
OR Workload per Procedure	NS	P<0.01
Patient Type	NS	P<0.05
PICU Bed-slot Allocation Status	NS	P<0.01
RR Workload per Day	p<0.01	P<0.01
RR Workload per Procedure	p<0.01	P<0.01
Number Urgent Bed-slots Allocated per Day	NS	NS
Urgent Bed-slot Allocation Status	NS	NS

'NS' = not significant; *trend towards significance (0.1 > p >0.05).

Table 3.13. Summary of results from the monotonic bed-slot allocations dataset

Scheduling Status: Monotonic and non-monotonic bed-slot allocation

The summary of results for each variable tested from the data sets of both monotonic and non-monotonic bed-slot allocations is shown in Table 3.14 below. The first column lists the dependent variables which were tested; the second column lists the level of significance for each dependent variable against the independent variable Scheduling Status.

Dependent Variable	**p for Scheduling Status**
AICU Bed-slot Allocation Status	p<0.01
Fast-track Status	p<0.05
Number AICU Bed-slots Allocated per Day	p<0.01*
Number Fast-tracks per Day	p<0.05*
Number Paediatrics per Day	NS*
Number PICU Bed-slots Allocated per Day	NS*
Operative Category	NS
OR Workload per Day	NS*
OR Workload per Procedure	p<0.01*
Patient Type	NS
PICU Bed-slot Allocation Status	NS
RR Workload per Day	NS*
RR Workload per Procedure	NS*
Month	NS
Day of the Week	NS

'NS' = not significant; *Equal variances not assumed.

Table 3. 14 Summary of results for comparison of non-monotonic and monotonic datasets

Table 3.15 below shows the differences in mean workload measure for each (numeric) dependent variable listed in the first column between non-monotonic and monotonic bed-slot allocations. As can be seen from the table, the hypothesis that there is a difference in the mean workload between monotonic bed-slot allocations and non-monotonic bed-slot allocations is confirmed for the dependent variables Number Fast-tracked per Day, Number AICU Bed-slots Allocated per Day and OR

Workload per procedure. Each of these workload measures showed a significant reduction in workload from the non-monotonic to the monotonic bed-slot allocations.

Dependent Variable	**Non-Monotonic**	**Monotonic**	**Difference**	**p* (2-tailed)**
Number Paediatrics per Day	1.981	2.024	-0.043	NS
OR Workload per Day	15.745	14.994	0.751	NS
RR Workload per Day	9.340	8.542	0.798	NS
Number Fast-tracked per Day	1.314	1.012	0.302	p<0.05
Number AICU Bed-slots Allocated per Day	3.475	2.892	0.583	p<0.01
Number PICU Bed-slots Allocated per Day	1.181	1.024	0.157	NS
OR Workload per Procedure	1.563	1.505	0.058	p<0.01
RR Workload per Procedure	0.897	0.856	0.041	NS

*p values between the two cases of equal variances being assumed and equal variances not being assumed in no case affected the overall judgements of significance. 'NS' = not significant.

Table 3.15. Comparison of means for numeric variables between the non-monotonic and monotonic datasets

Table 3.16 below shows the F values and corresponding level of significance for the hypothesis that there is no difference in the variance between the workload distributions of the monotonic and non-monotonic bed-slot allocations for each numeric measure of workload. As can be seen from Table 3.16, there is no significant difference in the variance in workload distributions for the dependent variables OR Workload per Day and OR Workload per Procedure using Levene's Test for the Homogeneity of Variance.

Dependent Variable	**F**	***St.Dev.*** [a]	***St.Dev.*** [b]	**Sig.**
Number AICU Bed-slots Allocated per Day	2.476	1.31	1.45	NS
Number Fast-tracked per Day	0.672	3.95	5.18	NS
Number of Paediatric Bed-slot Allocations per Day	3.215	4.35	4.92	NS*
Number PICU Bed-slots Allocated per Day	1.634	1.01	1.11	NS
OR Workload per Day	6.016	1.30	1.51	p<0.05
RR Workload per Day	2.785	0.83	0.94	NS*
OR Workload per Procedure	26.094	0.50	0.50	p<0.01
RR Workload per Procedure	9.493	1.27	1.16	p<0.05

'NS' = not significant; * trend towards significance; a) variance for non-monotonic bed-slot allocations workload distribution; b) variance for monotonic bed-slot allocations workload distribution.

Table 3.16 Standard deviation comparisons for numeric variables between the non-monotonic and monotonic datasets

Conclusions

The conclusions of the Resource Allocation Evaluation Study are that there are various weaknesses in the process of patient scheduling currently in place at RBH. In particular, there were epistemological control-limiting artefacts for the variables Number Fast-Tracks per Day and Fast Track Status and OR Workload per Day when considered across the different months of the study period. There were also epistemological control-limiting factors for PICU Bed-Slot Allocation Status considered across the different days of the week.

In comparing the means of performance variables between monotonic and non-monotonic bed-slot allocation datasets there was a significant difference in the mean values of Number Fast Tracks per Day, Number AICU Bed-Slots Allocated per Day and OR Workload per Procedure. In each of these variables, there was a reduction in the mean value of the variable as it occurred in the monotonic dataset from the non-monotonic dataset. As with the testing of significant variances in values across the different months of the study period or days of the week, this result again indicates the presence of epistemological control-limiting factors in the patient scheduling process currently in place at RBH.

The full results and discussion of the Resource Allocation Evaluation Study is included as Appendix 10.

3.2.7. The Intermediate-Care Units (HDUs)

In all, there are five HDUs, four of which are adult, one of which being designated primarily for privately funded patients, rather than patients of a particular operative category. Each adult HDU is situated within a larger hospital ward and, although HDU beds are separated geographically from regular ward beds, there are many resource inputs that are shared. In particular, the nursing and clinical staff are not specific to either HDU or non-HDU bed, although the nursing duties are organised hierarchically, with more experiences nursing staff being responsible for the HDU beds. The paediatric HDU is situated within PICU rather than a regular ward.

Each HDU has 4 bed spaces with usually 1 nurse responsible for all HDU patients. This level of nursing is much less than that found in either intensive-care unit or RR, although is significantly higher than in regular ward areas. This is a reflection of the HDU as a half-way house in terms of therapeutic dependency, i.e., a dependency on therapeutic interventions. Historically, however, the HDU is not necessarily a half-way house in terms of observational dependency, i.e., a dependency on physiological monitoring. Indeed, one of the main reasons for the establishment of HDU beds was to distinguish between the need for monitoring and the need for therapeutic intervention - a patient may require the former, though not necessarily the latter, during particular phases of a disease process. Thus, instead of placing such patients in an expensive intensive-care bed, they may instead be put in an HDU bed which is usually a much less expensive option. Nevertheless, the level of observational dependency catered for in an HDU bed will typically still be less than that in an intensive-care bed.

A second reason for the establishment of the intermediate-care unit was as a means of reducing the difference in dependency between the intensive-care unit and the regular wards. This has the effect of making the flow of the patient around the high-dependency environment be more a reflection of the underlying physiological condition. Without the HDU, a patient with a level of dependency somewhere below that catered for in an intensive-care bed, but above that catered for in a regular ward bed, would have to stay in an intensive-care bed.

The following table summarises the details of each of the adult HDUs. Although in each case the exact figures are currently not available, it is nevertheless possible to quantify to the extent of saying, for example, that the majority of admissions are in category X, from the data available from other high-dependency units, along with the experience of the nursing staff involved.

HDU	**Typical Patient Type**	**Typical origin(s) of admissions**
Alexandra	Publicly-funded Cardiac Post-Operative	AICU/RR
Elizabeth	Publicly-funded Thoracic Post-Operative	AICU/RR
Reginald Wilson	Privately-funded post-operative	AICU/RR
York	Cardiac medical	Other Hospital/AICU/Internal Hospital Ward

Table 3.17. Summary details of the four adult HDUs at RBH.

In each case, the destination for discharges from an HDU will usually be the regular area of the particular ward in which the HDU is situated. In the case of York, however, a proportion of the patients will go to OR or AICU if their condition deteriorates and requires more intensive therapy or surgical intervention. York will also discharge a proportion of patients to other hospitals as part of its function as a local coronary care unit or CCU, which admits cardiac non-operative patients having a relatively high level of observational dependence, usually to determine any subsequent need for surgical intervention.

The HDUs can be an important consideration in developing a model of the high-dependency environment since they often represent the end of a patient's high-dependency care process, in which case they have the greatest potential to cause disruption upstream simply because there is more upstream for them to disrupt than there is for other units. In practice, however, the RBH HDUs do not cause disruption to upstream units. This is primarily because they are often not as adequately resourced to fulfil their role as an intermediate stage in patients' post-operative recovery process, in which case there is little functional discrimination between the HDU area of a ward and the regular ward area. However, it is felt by some at RBH that HDUs could further relieve the pressure on upstream units - in particular AICU and RR - if they were better resourced to take a more dependent case-mix. For this reason, in the development of the operational models in later chapters, the HDUs will not be considered as distinct from regular wards.

3.3. Resource Allocation

Resource allocation in the RBH high-dependency environment consists of allocating bed-slots to patients awaiting admission to the environment. Typically, a patient will first be admitted to the operating room from a pre-operative ward, followed by either an intensive-care unit (AICU or PICU) or the RR room, followed by an intermediate-care unit, as shown in Figure 3.20 below.

Although the patient flows depicted in Figure 3.13 below represent the resourcing requirements of most patients admitted to RBH, there are very many other possible patient flows. In particular, the first unit of the high-dependency environment to which the patient will require admission is not necessarily OR. RBH admits both operative and non-operative patients. Some operative patients and all non-operative patients will require primary admission to an intensive-care unit (AICU or PICU). These patients may subsequently require admission to OR and/or an intermediate-care unit. The effect this has on the process of resource allocation is to make it more of a distributed system, with the allocation of intensive-care unit bed-slots being managed both by the managers of the intensive-care units themselves in the case of those patients with a primary admission to AICU or PICU and by the managers of OR in the case of those patients with a primary admission to OR and subsequent admission to an intensive-care unit.

Because of the distributed nature of resource allocation at RBH, the co-ordination of the different processing streams becomes essential for the effective operation of the whole system. This is achieved by the role of Master Scheduler, who has the responsibility of both allocating resources for those patients with primary admission to OR as well as co-ordinating the allocation of resources within the intensive-care units for those patients with primary admission to AICU or PICU.

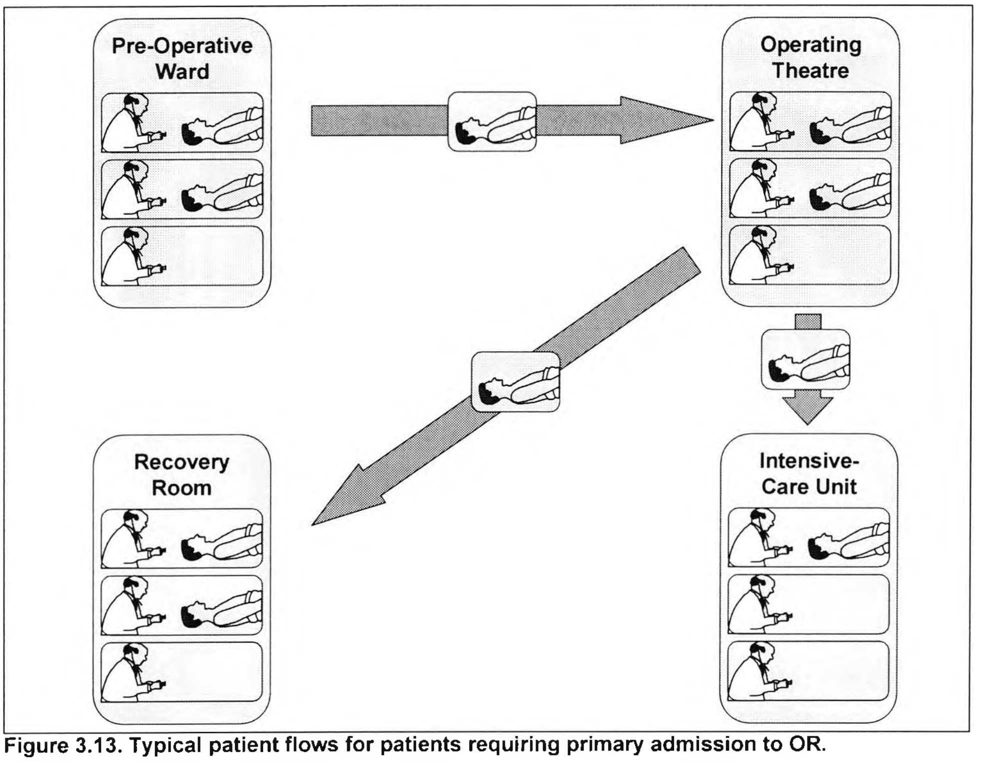

Figure 3.13. Typical patient flows for patients requiring primary admission to OR.

The role of Master Scheduler is depicted in the data flow diagram[9] of Figure 3.14 below. In this diagram, Master Scheduler receives admission requests from surgeons, usually via the surgical secretaries. The admission request contains information regarding the patient's identity, the surgeon making the request, the operative procedure involved, the degree of urgency, special resourcing requirements such as blood transfusion requirements and recommendations for post-operative care.

Each admission request is taken from patients on the surgeon's waiting list. Because of the way the NHS has evolved, the surgeons are normally individually responsible for patient care up to 48 hours after discharge from OR, at which point patient care becomes the responsibility of the director of either AICU or PICU for those patients in those units, or the consultant on duty in intermediate-care units or regular wards. Thus, during the period when resource allocation decision are made, each patient is the responsibility of a particular surgeon and appears on that surgeon's waiting list, and it is the surgeon's decision of when to request that a patient be admitted to OR.

Each OR bed-slot is reserved for the use of a particular surgical team, where a surgical team is identified by the combination of a leading surgeon and a consultant anaesthesiologist. The availability

[9] The formalism for data flow diagrams here is to represent processes by yellow boxes and data stores – which may be computerised databases or some other form of paper-based data storage – by green boxes. The arrows connecting processes and data stores denote flows of data from origin to source.

of surgeons and anaesthesiologists are determined by their respective rosters. Thus, when allocating an OR bed-slot to a patient, Master Scheduler must allocate a bed-slot which is reserved for the use of the surgeon responsible for the patient.

The allocation of OR bed-slots in this manner has become known as block-booking [MRG73]. There is now much evidence showing that the block-booking of OR bed-slots is more effective in maximising the utilisation rates of OR resources ([OZK95], [HAN92]) than the first-come-first-serve method where each OR bed-slot is allocated to a surgical team as and when it became available. Block-booking is more effective for two reasons. Firstly, each surgical team is able to start another patient straight after finishing a previous patient without having to move from one operating room to another or wait for another surgical team to complete their procedure. Secondly, because the surgeon has more control over his own operative workload, he will be more willing to perform procedures later in the day than otherwise since he can be more confident of starting on time.

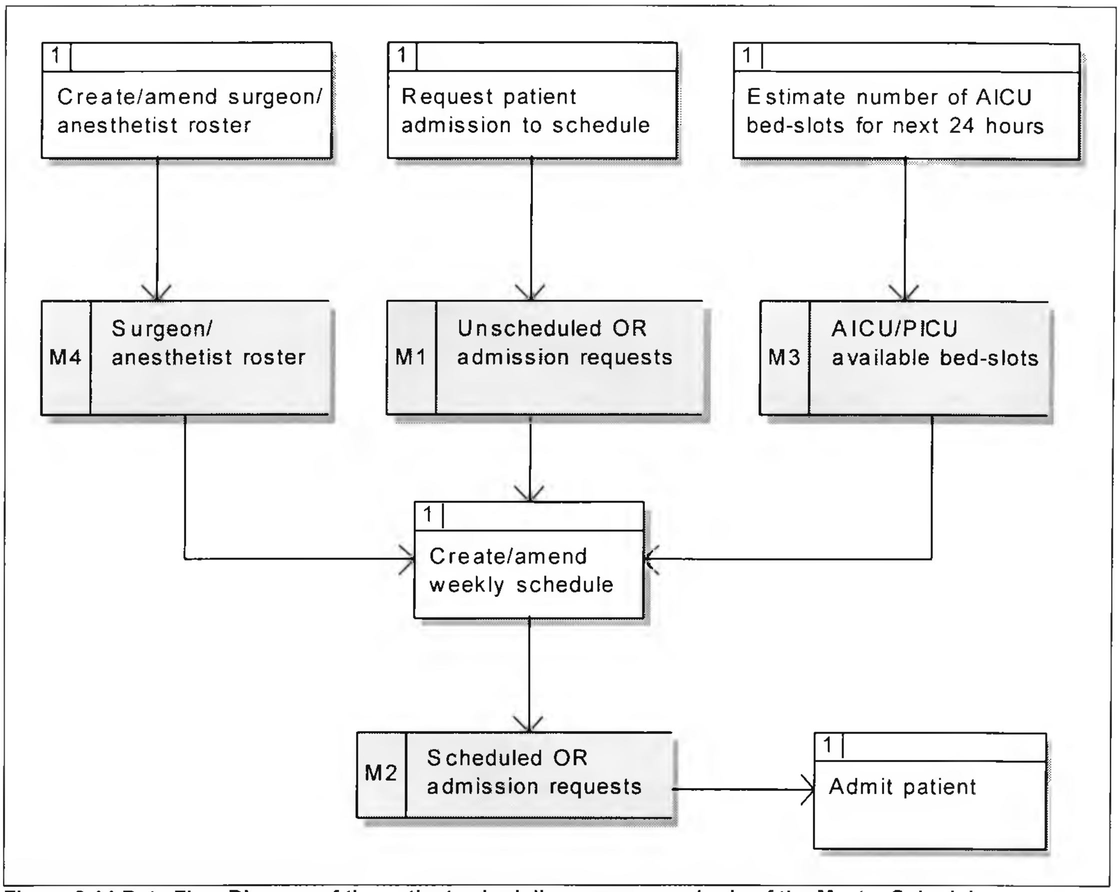

Figure 3.14 Data Flow Diagram of the patient scheduling process and role of the Master Scheduler

As well as the rosters for the surgeons and anaesthesiologists and the admission requests, Master Scheduler also requires information regarding the availability of AICU and PICU bed-slots. The availability of AICU and PICU bed-slots is estimated every morning by clinical staff within those units for the next 24 hours and this information is then passed on to Master Scheduler. The estimate of bed-slot availability is based on the physiological condition of patients already admitted to the intensive-care units, which is taken as an indicator of how much longer the patients already admitted will continue to consume intensive-care unit resources before they can be safely discharged to a lower-dependency unit. In addition, any patients whose primary admission will be to either intensive-care unit and is expected within the next 24 hours will also be taken into consideration.

With the estimates of intensive-care unit bed-slot availability, OR bed-slot availability – in the form of the surgeons' and anaesthesiologists' rosters – in hand, Master Scheduler is then in a position to allocate resources to patients appearing on the admission requests for primary admission to OR. The outcome of this resource allocation process is a weekly schedule. A weekly schedule allocates each patient requiring primary admission to OR within the next working week to at least one OR bed-slot and at least one bed-slot within either RR, AICU or PICU. Until a patient requiring primary admission to OR appears on a weekly schedule, the patient is classified as unscheduled; when a patient does appear on a weekly schedule, the patient is classified as scheduled.

Figure 3.15 below is an example template which is used by Master Scheduler in the construction of the weekly schedule. It can be seen that each operating room is divided into two blocks of time each day – one in the morning and one in the afternoon. (It is these blocks of time which were defined as individual OR bed-slots in the Resource Allocation Evaluation Study of the previous section). Each block of OR time is classified by three attributes as follows:

- Operative category of the procedure (i.e., cardiac or thoracic, designated by cell colour in the figure);
- Leading surgeon (designated by the surgeon's initials in the figure); and
- The post-operative bed-slot requirements of the patient (i.e., RR, AICU or PICU)

DAY	TIME	OR # 1	OR # 2	OR # 3	OR # 4	OR # 5
Monday	AM	JP - AICU	UP - RR	DS - PICU	TBA - F/T	
Monday	PM	JP - F/T	UP - RR	DS - PICU	TBA - F/T	
Tuesday	AM	MY - F/T	PG - RR	CL - PICU	NM - F/T	MY - F/T
Tuesday	PM	MY - AICU	PG - RR	CL - PICU	NM - AICU	CL - F/T
Wednesday	AM	JP - AICU	PG - RR	CL - PICU	DS - PICU	TBA - RR
Wednesday	PM	JP - F/T	PG - RR	CL - PICU	DS - AICU	TBA - RR
Thursday	AM	MY - F/T	UP - RR	DS - F/T	NM - AICU	
Thursday	PM	MY - PICU/AICU	UP - RR	DS - AICU	NM - F/T	
Friday	AM	JP - F/T	PG - RR	CL - F/T	NM - AICU	
Friday	PM	JP - AICU	PG - AICU		NM - AICU	

Key:
= OR bed-slot designated for cardiac procedures
= OR bed-slot designated for thoracic procedures
= Operating room closed for scheduled procedures
F/T = OR bed-slot designated for fast track OR overnight recovery procedures
RR = OR bed-slot designated for patients requiring admission to RR
AICU = OR bed-slot designated for patients requiring admission to AICU
PICU = OR bed-slot designated for patients requiring admission to PICU

Figure 3.15. Example template for the weekly schedule used in the patient scheduling process

The classification of OR bed-slots according to leading surgeon, operative category and post-operative bed-slot requirements has a degree of flexibility, especially with regard to the post-operative bed-slot requirements. Thus, for example, if AICU has few bed-slots estimated to be available on a particular day, OR bed-slots may be re-allocated for patients requiring instead post-operative bed-slots in RR or PICU.

As argued in Chapter 2, so long as a patient is scheduled but not yet admitted to OR, the allocation of any bed-slots to that patient is non-monotonic. Consequently, it is possible to amend the weekly schedule on a daily or even hourly basis by deleting or adding patients to the weekly schedule. Such additions or deletions are inevitable given the difficulty in estimating the availability of intensive-care resources for any period of time greater than 24 hours and the need to admit emergency patients to OR.

The allocation of OR bed-slots to emergency patients follows the same process as that depicted in Figure 3.14 above. In those cases where the patient needs to undergo surgery within a 48 hour time frame, (so-called urgent patients [FRO94]), it is usually possible to swap the patient with one already allocated OR bed-slots in the same time-frame and for whom the surgery may be safely delayed. In so-called emergent cases, where the patient must undergo surgery within a much smaller time-frame, 1-2 hours for example, it is necessary to allocate the next 'available' OR bed-slot to the patient by deleting the allocation for the patient who would have otherwise consumed that bed-slot in the case where the next bed-slot is already allocated. At RBH urgent cases are relatively common, although emergent cases are very rare and usually occur as a result of a problem arising in one of the catheter labs or in an intensive-care unit.

As argued in Chapter 2, the cost-effectiveness of the resource allocation process is determined by the extent to which it is able to optimise resource utilisation under the constraint of at least maintaining the quality of patient care. In the case of a progressive-care system this requirement translates into allocating resources in such a way that workload across all component units of the system is at a consistently high level, although not so high as to risk over-stretching resources or cancelling admissions. Thus, in allocating bed-slots to patients with primary admission to OR, Master Scheduler must aim to optimise not only resource utilisation of OR bed-slots, but also those in AICU, RR and PICU, the aim being to produce workload profiles approximating those shown in Figure 3.16 below for RR and AICU (workload measures are those used in Resource Allocation Evaluation Study).

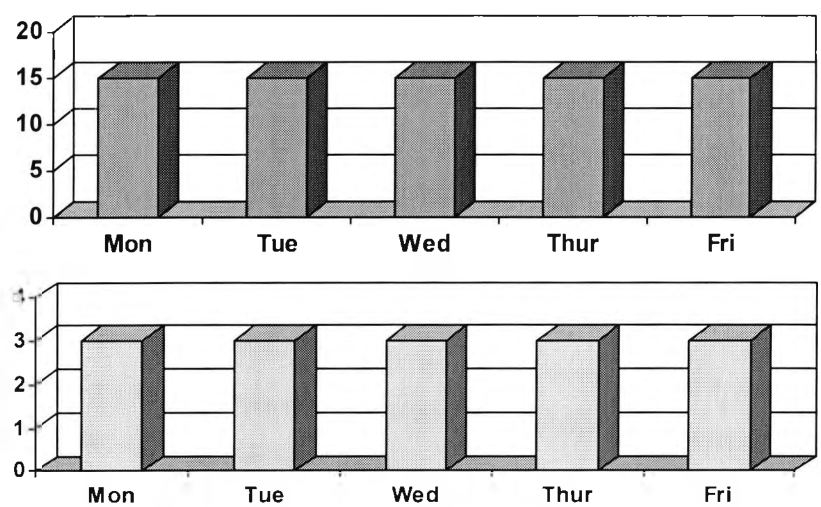

Figure 3.16. Ideal daily workload profiles for RR (top) and AICU (bottom)

The degree to which the situation depicted in Figure 3.16 can be realised is, as argued in Chapter 2, dependent on the degree to which resource allocation is controllable. In the case of the allocation of AICU and RR bed-slots to patients with primary admission to OR, this depends on the extent to which Master Scheduler is able to control the allocation of OR bed-slots.

The Resource Allocation Evaluation Study presented in Appendix 9 and summarised above showed that there are both epistemological and non-epistemological control-limiting factors evident in the performance evaluation of the RBH resource allocation system. This is clearly seen in Figure 3.17

below which depicts the actual daily workload for RR and AICU for a typical week. In this figure there is large variation in workload for both units during the week, with bed-slots being under-utilised some days and on other days utilised to near maximum capacity.

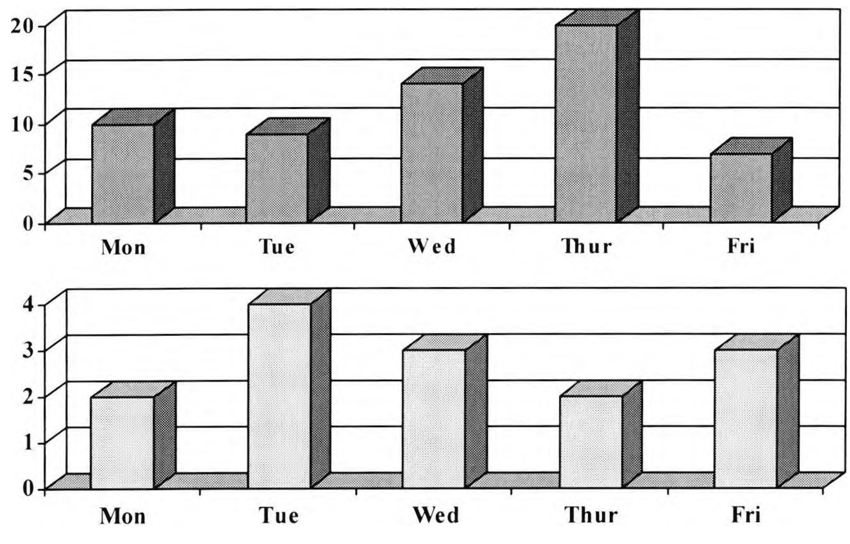

Figure 3.17. Actual daily workload profiles for RR (top) and AICU (bottom)

It may be hypothesised that the example weekly schedule template shown in Figure 3.15 provides some explanation for the variations in workload for AICU and RR shown in Figure 3.17. If it is assumed that there are no epistemological or non-epistemological control-limiting factors – that is, using the terminology adopted in Resource Allocation Evaluation Study, that those bed-slot allocations which appear in the non-monotonic scheduling status database also appear in the monotonic scheduling status database and vice-versa – then the resulting daily workloads for AICU and RR would still show large variations throughout the working week. For example, the weekly template shown in Figure 3.15 allocates a number of RR bed-slots on a Tuesday equating to a workload of 18; on a Wednesday, however, the corresponding workload is 10. The resulting workload for both AICU and RR according to this line of reasoning is shown in Figure 3.18 below.

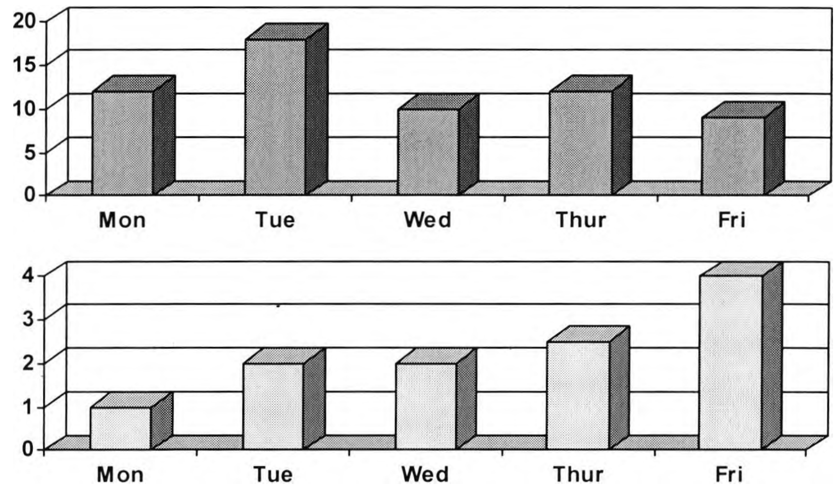

Figure 3.18. Resulting workload for RR (top) and AICU (bottom) for the weekly template of Figure 3.22

To test the hypothesis that at least some of the variation in mean daily workload figures is correlated with the implied workload in the weekly schedule template of Figure 3.15, the two data sets (using the same data as was used in Resource Allocation Evaluation Study) were compared using bivariate correlation. The hypothesis was confirmed in the case of RR (Pearson Correlation Coefficient = 0.842, p = 0.002) and was not confirmed in the case of AICU (Pearson Correlation Coefficient = 0.105, p = 0.182).

These results have to be taken with some degree of caution however. The implicit assumption is made in the above analysis that the same weekly template is used throughout the study period. In reality this is not the case as templates tend to change from month to month as the composition of surgical teams change. These changes are relatively minor, although the incremental affect may be sufficient to invalidate the above results. Unfortunately, no permanent records of historical weekly schedule templates are kept to verify this.

The above results suggest that the variation in the weekly scheduling template workloads is more important in determining the variation in the actual variations in workloads for RR than AICU. In the above analysis, the data used for the actual bed-slot allocations was from the monotonic scheduling status database. To add further support to the above results, therefore, the weekly template schedule workload was compared to the workload data for the bed-slot allocations from the non-monotonic scheduling status database on the assumption that if these show a greater correlation between the two AICU workload sets of data and a similar correlation between the two RR workload sets of data, then variation in AICU workload data for the monotonic scheduling status workload data is more dependent on informational and organisational control-limiting factors than variations implicit in the weekly scheduling template. As with the above analysis, the two sets of data were compared using bivariate correlation and the hypothesis was confirmed for both RR (Pearson Correlation Coefficient = 0.239, p = 0.003) and AICU (Pearson Correlation Coefficient = 0.196, p = 0.013).

These results are as would expected – AICU workload implied by the weekly scheduling template is more susceptible to the detrimental influence of control-limiting factors than RR workload. There are three reasons for making this claim. First, urgent cases admitted to OR are more likely to require an AICU bed-slot allocation than an RR one. Second, much of the RR workload is for fast-track or overnight recovery patients. However, approximately 13% of these patients are so-called failed fast-tracks. That is, patients who were originally classified as requiring an RR bed-slot allocation upon discharge from OR, but were subsequently re-classified – either in OR or RR – as requiring an AICU bed-slot allocation.

Finally, the allocation of bed-slots in RR is very different from that in AICU since AICU also allocated bed-slots to patients other than those whose primary admission is to OR. Therefore, since these admissions to OR are typically more urgent than those whose primary admission is to OR, the allocation of AICU bed-slots to patients with primary admission to OR must be secondary to the allocation of AICU bed-slots to other patients. Moreover, the allocation of AICU bed-slots must take place in the context of an already admitted population of patients within AICU and an estimate of the availability of AICU bed-slots during the next 24 hours, both of which limit the possible AICU bed-slot allocations which can be made for patients with primary admission to OR. Neither of these are considerations for RR. In the first place, no patients stay in RR for more than 24 hours (any overnight recovery patients who are not able to be discharged to an intermediate-care unit before OR opens the following morning are discharged instead to AICU), in which case the allocation of RR bed-slots starts with a blank page each day. Second, the ability to predict the bed-slot requirements for patients admitted to RR is much easier as the length of time required in RR is determined mainly by the anaesthetics used during surgery and the operative category of the patient, rather than the presence of some other extraneous risk factors or serious physiological derangement. Moreover, in those cases where the length of time required in RR is in excess of the estimate, there is always the possibility of transferring the patient to AICU or PICU instead.

Despite these considerations, the allocation of RR bed-slots remains a crucial factor in the cost-effective operation of the whole system and there are several factors which Master Scheduler needs to consider when allocating RR bed-slots. Chief among these is the limiting of the number of fast-track and overnight recovery patients per day, and as importantly, the controlled admission of them to RR. Whereas thoracic patients may only take one hour to recover and be discharged to a ward, fast-track and overnight recovery patients may take anywhere up to 24 hours to recover. This is an important consideration for Master Scheduler, since RR only has five bed-slots available at any one time, in which case admission of these patients has to be limited to a maximum of 3 per day, and spread throughout the day as much as possible to avoid RR becoming a system bottle-neck.

The foregoing discussion, together with the Operating Theatre Utilisation Analysis Study presented in Appendix 11 and the AICU Chronicity Study presented in Appendix 9, demonstrates that resource allocation in the kind of interdependent, high-dependency healthcare systems exemplified by RBH, is a complex optimisation problem. Operating Theatre Utilisation Analysis Study demonstrated that at least part of the solution to this problem involves the removing of epistemological control-limiting factors. It is the central hypothesis of this dissertation that the introduction of a computerised information system would help remove these epistemological control-limiting factors and thus improve the cost-effectiveness of healthcare delivery within the empirical domain at RBH. The objective of the following two chapters is to develop the specifications of such an information system. In the next chapter a modelling formalism will be developed which is capable of depicting an operational model of the empirical domain at RBH together with a proposal for an information system capable of providing a solution to the sub-optimisation of resources.

4. Modelling Approach and Formalism

4.1. Introduction

The objective of this chapter is to detail a modelling approach and formalism for the subsequent development of the models in the next chapter. The term 'modelling approach' is intended to designate a sequence of ordered developmental stages in a modelling exercise where each stage builds upon the previous stage according to a method which complements the overall aim of the modelling exercise. The term 'modelling formalism' is intended to designate a formal modelling language consisting of a graphical or textual set of symbols and construction rules that operate on those symbols which may be used to develop models of the empirical domain.

The modelling approach that will be developed in this chapter is proposed as being especially suited for the introduction of a management information system into complex business environments such as the high-dependency environment of the Royal Brompton Hospital . It will do this by allowing a comparative evaluation to be made between the existing empirical domain in terms of its structure processes and components and how the empirical domain would be once the proposed management information is introduced. This comparison will be made between two models of the environment – the current operational model and the proposed operational model. The approach developed for defining these models is proposed as a hybrid between existing approaches adopted in business process re-engineering exercises, and approaches adopted in software engineering for the design and implementation of software solutions.

It will be argued that the complementary nature of these two approaches implies a need for a hybrid alternative to ensure the successful implementation of software solutions and the corresponding re-engineering of the empirical domain. It will be further argued that existing approaches proposed for the re-engineering of business processes are currently not nearly well developed enough to provide the basis of informing the process of software engineering. In this regard, a formalised approach for business-process re-engineering will be proposed that takes a control-theoretic perspective and makes a fundamental yet often ignored distinction between the re-engineering of business-processes in terms of changing the processes themselves and the implementation of existing processes in different a system which is the characterising quality of the process of computerisation.

As a consequence of adopting this hybrid approach, the modelling formalism will be similarly proposed as a hybrid between those formalisms typically used for the design and coding of software solutions on the one hand and the re-engineering of the empirical domain to facilitate the implementation of those software solutions. It will be argued that a formalism is needed which is capable of being used at every design stage involved in the complementary processes of designing a management information system and re-designing the empirical domain to both inform the design of the management information system and to ensure its successful implementation and integration into the empirical domain. The formalism which will be proposed is an object-oriented formalism based on Petri nets [AJM89], [AJM89], [AJM95], [AUD95], [HOL89], [JEN88], [JEN89], [LOP95], [PED94], [PET81], [PTR62], [ZUR94] and Adaptive Reference Technology [MAK94]. This formalism will be contrasted against the industry standard object-oriented formalism, the Unified Modeling Language (UML) [RAT97a], [RAT97b], which, it will be claimed, fails to satisfy many basic criteria necessary for both the design of software solutions and the re-engineering of business-processes.

The contents of this chapter are structured according to the principle that the choice of modelling formalism should be informed by the modelling approach adopted, which in turn should be informed by the aims of the modelling exercise. Thus, in the following section the modelling approach will be developed, with the formalism being developed in the subsequent section. In the final section there will be a brief summary of the main features of the approach and formalism developed in this chapter and a discussion of how these relate to the development of the operational models developed in the next chapter.

4.2. Modelling Approach

The term 'modelling approach' is intended to designate a sequence of ordered developmental stages in a modelling exercise where each stage builds upon the previous stage according to a method which complements the overall aim of the modelling exercise. The aim of the modelling exercise for this project is to design a management information system which supports the primary hypothesis of this dissertation – that such a management information system is able to increase the cost-effectiveness of healthcare delivery in a progressive-care system through increasing the effectiveness of control over resource allocation.

When many software engineers encounter the term 'information system', they consider it to apply only to a software component of that system, with other components of the system being interpreted as 'users'. While this narrow conception of the modelling domain is understandable from the perspective of the software engineer, it will be argued here that it is only a suitable approach to adopt during the later stages of the software engineering process.

The software engineering process may be decomposed into five main developmental stages as shown in Figure 4.01 below.

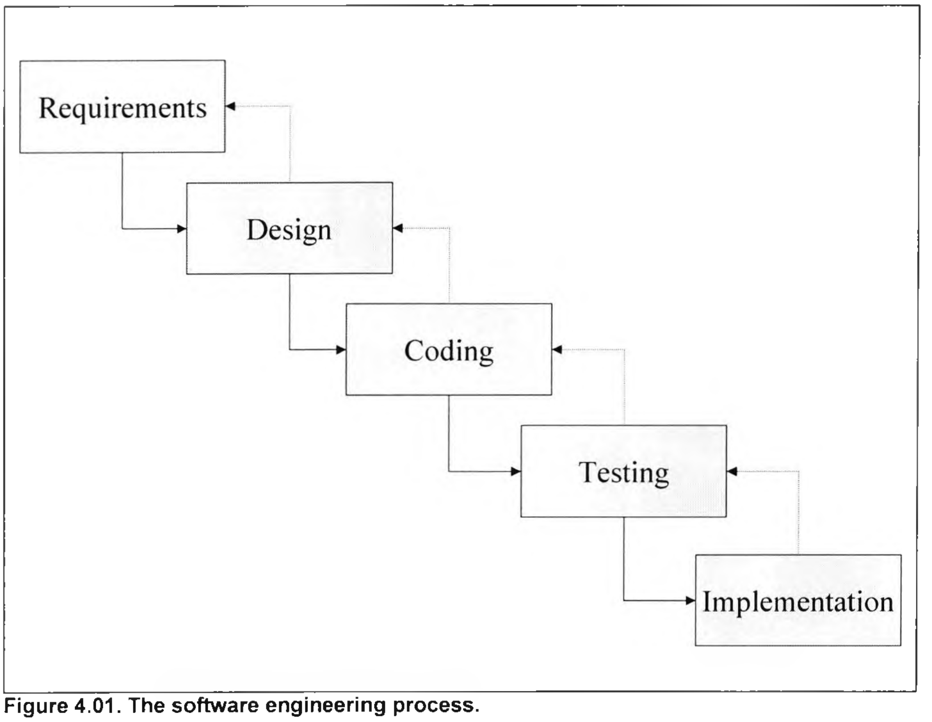

Figure 4.01. The software engineering process.

The initial stage of the software engineering process is Requirements [DVI90]. The objective of Requirements is to determine the requirements of the software solution to be engineered. This is often as interpreted as the determination of the user-level functionality. The term 'user-level functionality' used here refers to the behaviour that is generated at the interface between the system and the user. Thus, the specification of requirements is the specification of the output the system must make as a response to an input from the user and vice-versa. It does not specify how either the user of the system generates those outputs in terms of the internal workings of either party. For example, the requirements for a television remote control device would be specified according to rules such as 'press button A, and the volume is muted'. The requirement of the high-dependency resource management and patient scheduling tool that is to be modelled in the next chapter is, in its most general formulation, be 'the system will generate information for enhancing the level of control over resource utilisation and patient scheduling by inputting data regarding patient clinical and demographic characteristics and resource availability'.

This interpretation of Requirements as a description of the user-level functionality is problematic, however, since in some instances it may not be at all clear which agent or agents are the users of the system and what is the system. This is the case in what may be called autonomous systems where no user of the system exists. For example, consider the case of a software solution for a system of traffic lights. The description of the requirements for such a system would be the sequence of colours of each set of traffic lights, how long they stay on each colour and the co-ordination of the changes of colours between each set of lights. Yet, because there is normally no interaction in such a system between the traffic lights and the traffic which they control, it is implausible (and in any case pointless) to posit the traffic as the user of the system.

Nevertheless, in those cases where there is a clear distinction between the user and the software, it is at first sight an intuitively reasonable approach to deploy this distinction as the basis of Requirements from the perspective of the software engineer. However, it will be argued here that a wider perspective should be adopted in the development of software solutions and further, that even from the perspective of the software engineer, the distinction between user and software perpetuates a blurring of the boundaries between Requirements and later developmental stages of the software engineering process which is detrimental to the successful development of software solutions and that a more viable basis of Requirements is needed.

A non-autonomous software system does not, by definition, exist in isolation from the wider system of which it is an integrated component. The introduction of a software subsystem into a system cannot normally done in the manner by which, for example, an electrical appliance can be plugged into a wall socket. The only instance where this can plausibly be achieved is where the functionality of the software subsystem is identical in every respect to the functionality of a subsystem which it replaces. In reality, the introduction of a software subsystem necessarily involves some degree of modification or re-engineering of the wider system. This re-engineering is obviously best planned in advance of the implementation of the software rather than post-implementation.

With regards to the implementation of the high-dependency resource management and patient scheduling tool to be modelled in the next chapter, it is clear that the software will be non-autonomous. Moreover, that inevitably it will necessitate some degree of re-engineering of the wider healthcare system.

If the implementation of a non-autonomous software subsystem involves the re-engineering of the wider system of which it is to be a component, then the use of the distinction between user and software as the basis of Requirements inevitably results in the need to re-engineer the system post-implementation of the software if the definition of the user requirements is based on the system design prior to the specifications of the software.

It may be argued that a possible way out of this problem which retains the utility of the distinction between user and software is to re-engineer the wider system, and by implication the user-level functionality, pre-implementation of the software, then simultaneously implement the software and the re-engineered design of the wider system. There is circularity in this argument, however. To re-engineer the wider system, and by implication the user-level functionality, to suit the software requires that there be a preliminary design of the software. But, on the interpretation of Requirements as involving the determination of user-level functionality, there being a preliminary design of the software presumes a co-present design of the user-level functionality.

The solution, therefore, is to use a parallel development approach, where the re-engineering of the wider system and the specification of the requirements that this implies for a software solution merge into one and the same, hybrid, process. Such a hybrid approach to Requirements is necessarily systemic in nature, needing to represent all aspects of functionality, rather than just those of the software subsystem and the interface with the users which are represented as atomic entities. As a modelling process, however, this hybrid approach to Requirements differs from the approach typically adopted in business process re-engineering (BPR) in its purest form. Typically, the BPR approach focuses on processes rather than their implementations. In the hybrid approach proposed here, by contrast, an explicit representation of process implementation is a fundamental requirement because of the need to discriminate between those processes to be implemented by the software subsystem and those processes to be implemented by other subsystems, i.e., those processes which constitute the users of the system under the interpretation of Requirements based on the distinction between software and user.

The interpretation of Requirements according to this hybrid approach is the systemic black-box modelling of the system alongside an explicit representation of the implementation of the processes which are thus modelled. Thus, the objective of Requirements is to define the functionality of the system in terms of identifying system processes, their inputs and outputs from and to other system processes and their respective implementations. A more detailed examination of this interpretation of Requirements will be provided below. For the moment it is worth briefly examining the other stages of the software engineering process depicted in Figure 4..01 above.

Once Requirements is complete, the next stage in the software engineering process is Design. Following the same line of reasoning as that given above for Requirements, then just as Requirements was interpreted as black-box modelling, so Design may be interpreted as white-box modelling [DVI90]. That is, to take the models generated at Requirements and determine how the functionality thus modelled in terms of inputs and outputs is to be generated. The distinction between these two modelling stages could be summarised by saying that Requirements determines the 'what' of system functionality, Design determines the 'how'. It should be noted, however, that Design has a narrower scope than does Requirements in the context of software engineering. Requirements necessitates the modelling of the whole system; Design, by contrast, needs only to model those system processes whose implementation is represented in Requirements models as being software, rather than, for example, human or some mechanical device.

The next stage after Design is Coding. Coding is the transcription of Design models as software code. This code is then tested in terms of its ability to implement the functionality required of it as specified in Requirements (and subsequently in Design) models that constitutes the next stage in the software engineering process. Finally, once Testing is complete, the final stage in the software engineering process as depicted in Figure 4.02 is Implementation, whereby the software code is implemented in the empirical domain and integrated with other systems comprising that domain.

The resulting sequence of developmental stages according to the hybrid approach is depicted in Figure 4.02 below. In this figure the first two developmental stages constitute the main modelling stages of the development process and thereby provide the basis for all future stages. It is these two modelling stages that will be the focus of both this dissertation and this discussion. The hybrid approach is concerned with these first and most important stages, leaving the remit of other stages largely unaffected.

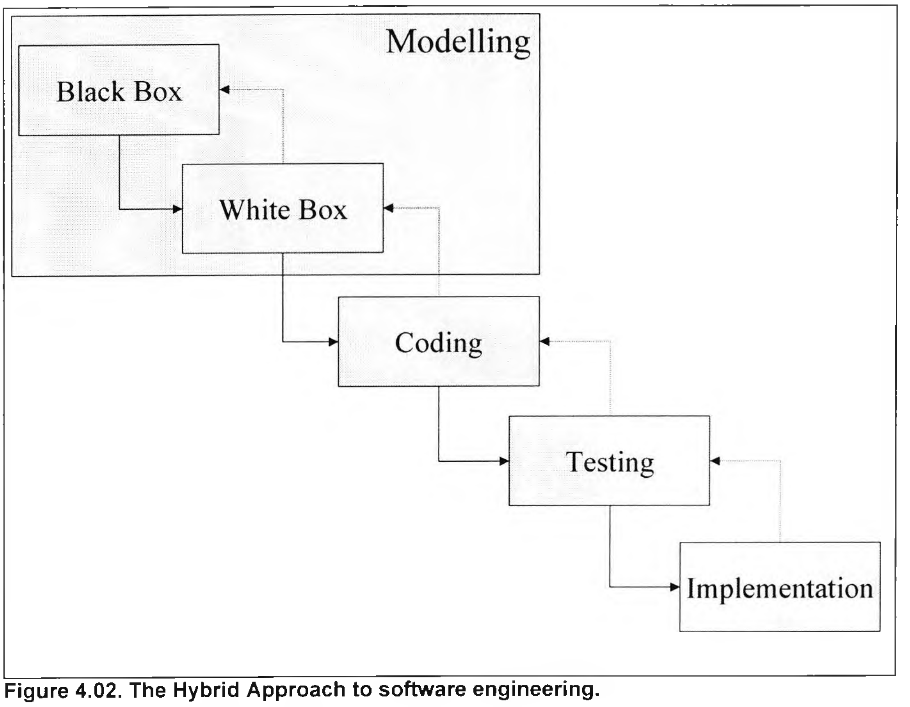

Figure 4.02. The Hybrid Approach to software engineering.

It was argued above that a parallel development approach is needed, where the re-engineering of the wider system and the specification of the requirements that this implies for a software solution merge into one and the same, hybrid, process. It will be argued here that this hybrid process should consist of specifying the requirements by making a comparison between two models referred to as the Current Operational Model (COP) and Proposed Operational Model (POP), where the latter is the outcome of applying principles of business process re-engineering (BPR) to the former. A high-level view of the approach that will be proposed is shown in Figure 4.03 below.

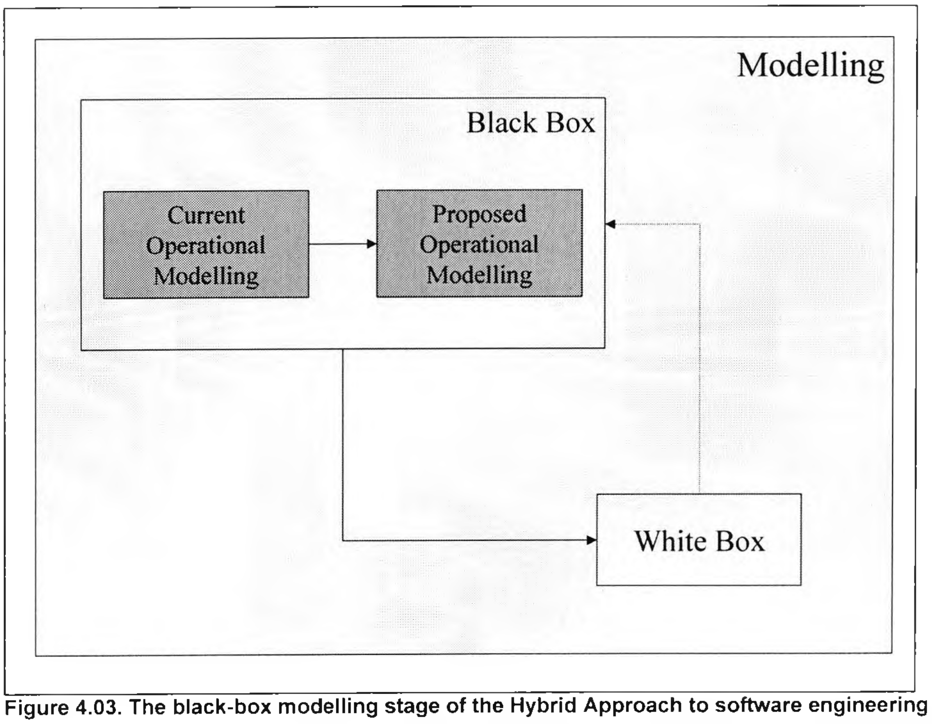

Figure 4.03. The black-box modelling stage of the Hybrid Approach to software engineering

The objective of the next chapter will be to develop operational models of the RBH high-dependency environment, both pre- and post-implementation of the high-dependency resource management and patient scheduling software tool. These models may be seen as corresponding to the black box modelling stage in Figure 4.02 above. Then, in Chapter 6, a preliminary set of high-level design models will be presented for the software tool itself. These models may be seen as corresponding to the white box modelling stage in Figure 4.02 above.

BPR has been proposed primarily as a business philosophy rather than a formal approach. The term was first used by Hammer [HMM90] in 1990 where the objective was to redesign the constituent flows of work from one processing unit to another, primarily with the aim of reducing downtime and thereby increase operating efficiency and quality of customer service. The main enabling idea adopted by Hammer was Case Management where, instead of having a linear workflow from one unit to another as typified by assembly lines, each processing unit works in parallel. The two situations are depicted in Figure 4.04 and Figure 4.05 below. In both of these figures, rectangular boxes represent business processes and circles represent work items with connecting arrows representing flows of work between business processes.

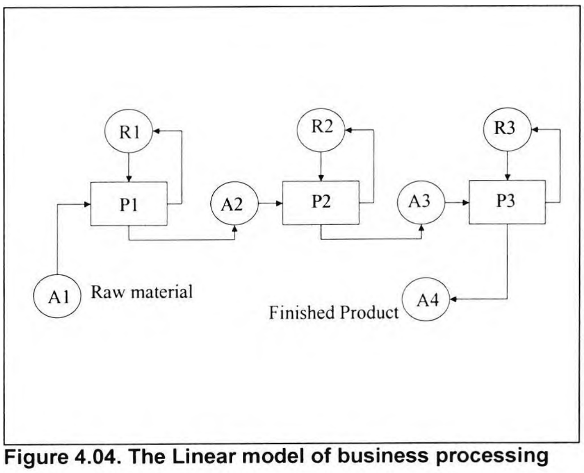

Figure 4.04. The Linear model of business processing

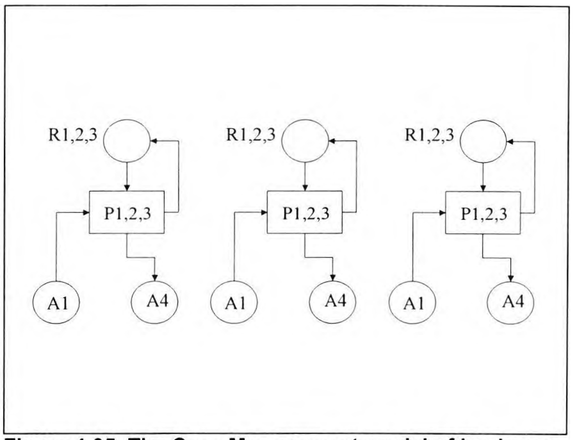

Figure 4.05. The Case Management model of business processing

The re-design of business processes according to the Case Management structure shown above is best seen as an instance of a wider conception of BPR rather than being identified with it. It seems appropriate, for example, to see the transition from a Case Management structuring of business process to an assembly line structure as involving the re-design of business processes, although of course it is the reverse process as that described by Hammer. Moreover, for the purposes of this discussion, it also seems reasonable to assume that the introduction of new information technology into the RBH high-dependency environment implies a need for some form of business process re-engineering in that environment.

The concept of BPR as being necessitated by advances in information technology has been developed by Kaplan and Murdoch [KAP90]. However, while Kaplan and Murdoch recognise the relationship between software design and BPR, their method of Core Process Redesign is assumed to involve a radical overhaul of the entire business. When seen in a historical context, this assumption seems unsurprising as information technology was being mass introduced into industry for the first time when Kaplan and Murdoch published their work. Now, however, the process of 'IT-initiated BPR' is best seen as being more incremental, with most new information technology being introduced as a replacement of older technology. This distinction between the holistic approach to BPR and the incremental approach is a departure from the original conception of BPR which was seen as the attempt to re-design business processes beginning from a carte blanche assumption. With the incremental approach, however, the process begins with the existing design and makes incremental changes to it to arrive at a new design.

From a modelling perspective, the incremental approach to BPR must assume that there is an existing model of the organisation that is used as the basis to derive a re-engineered model of the organisation. These two models will be referred to here as the Current Operational Model (COP) and the Proposed Operational Model (POP) as shown in Figure 4.03 above and as will be presented for the case of the RBH high-dependency environment in the next chapter. Thus, the hybrid approach can be described as the method for making the transition from COP to POP, where that method is based on the principles of an incremental conception of BPR seen in the context of a software engineering project. The resulting POP should therefore represent the basis of the requirements, not only for the software, but also for the wider organisational environment.

Before developing the method of deriving the POP from the COP, it is important to make a fundamental distinction that is usually overlooked in the literature on BPR. This is the distinction between what shall be referred to as topological re-engineering and implementation re-engineering. The example of Case Management shown in Figure 4.04 and Figure 4.05 above is an instance of topological re-engineering, where the dependency relationships between different processing units are re-engineered. Topological re-engineering is therefore the modification of the topological properties of a system that consists of the dependency relationships that obtain between different processing units.

In contrast, implementation re-engineering is the modification of the processing units themselves in terms of the resources that comprise them. The most common example of implementation re-engineering is the process of computerisation, where computerised processing units replace human processing units.

Implementation re-engineering and topological re-engineering are mutually independent. For example, the replacement of a human by a computer does not imply the need to make corresponding changes to processing topology beyond that which may be required due to the modified interface between the newly introduced computerised processing unit and other directly dependent processing units. Similarly, the introduction of, for example, Case Management does not imply the need to also extend the computerisation of system processing since essentially the same system functionality exists, it is simply that the relationships between different processing units has changed.

Given that a distinction may be made between the topological and implementation re-engineering of system processing, and that these two activities are mutually independent, it can be seen that Hammer and Kaplan and Murdoch are, in fact, talking about very different conceptions of BPR. Nevertheless, the two conceptions are not at odds to one another; they may be undertaken in parallel, and indeed may be considered as complementary to one another insofar that the process of introducing new technology may facilitate a corresponding modification to system topology and vice-versa.

On the premise that, although mutually independent, there is nevertheless a certain degree of symbiosis between the implementation and topological components of BPR, the question arises as to how these two activities should be ordered in the overall modelling exercise. The most obvious answer to this is that that component which should come first is that which is the more fundamental in achieving the goals of the modelling exercise. In the case of software engineering, for example, it seems reasonable to suggest that the implementation re-engineering activity should come first, given that the overall aim of the software engineering process is normally one of computerisation. This line of reasoning, however, is contrary to the principle argued for above that the re-engineering of the wider system and the specification of the requirements that this implies for a software solution merge into one and the same, hybrid, process. The same line of reasoning applies to the idea of undertaking topological re-engineering first, followed by implementation re-engineering.

The solution is to adopt a to-and-fro approach, where an initial implementation or topological re-engineering process facilitates a subsequent re-engineering process, and so on until a satisfactory black-box model is produced that can then be used as input to the next stage of the overall modelling process. The resulting decomposition of the black-box modelling stage is depicted in Figure 4.06 below.

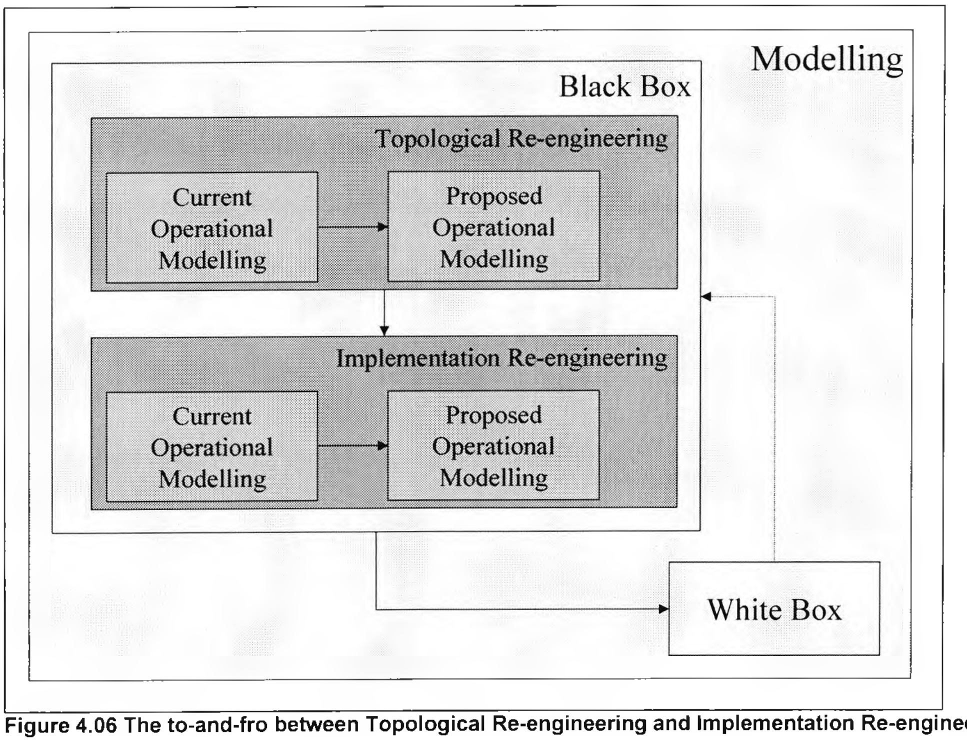

Figure 4.06 The to-and-fro between Topological Re-engineering and Implementation Re-engineering

Considering now the actual re-engineering techniques used in deriving the POP from the COP of an organisation, using the complementary processes of implementation and topological re-engineering. Clearly, one technique is needed to guide the implementation re-engineering process, and another technique needed to guide the topological re-engineering process. Neither of these two types of technique have been proposed in the literature on BPR, save for some superficial comments on the management of BPR projects.

Although different in nature, both implementation re-engineering and topological re-engineering are directed towards the same goal of increasing the cost-effectiveness of production. In the case of computerisation, therefore, it follows that computerisation qua BPR is directed towards that same goal of increasing cost-effectiveness. It was argued in Chapter 2 that in the case of healthcare delivery, increasing the cost-effectiveness of healthcare delivery is facilitated by increasing the effectiveness of control over the process of resource allocation. Moreover, that increasing the effectiveness of control may be achieved at least in part by reducing the effect of epistemological control-limiting factors. This specific hypothesis for the healthcare scenario will be generalised here to any type of organisation. This will allow for an abstract formulation of the mechanisms behind successful BPR projects presented in Appendix 12.

4.3. Modelling Formalism

In the previous section the objective was to develop a modelling approach that is especially suited for the introduction of a management information system into complex business environments exemplified by the RBH high-dependency environment. In this regard, a formalised approach for business-process re-engineering was proposed that takes a control-theoretic perspective and makes a distinction between the re-engineering of business-processes in terms of changing the processes themselves and the implementation of existing processes in different systems, which is the characterising quality of the process of computerisation.

As a consequence of adopting this hybrid approach, the modelling formalism that will be developed in this section will also be a hybrid between those formalisms typically used for the design and coding of software solutions on the one hand and the re-engineering of the empirical domain to facilitate the implementation of those software solutions. It will be argued that a formalism is needed which is capable of being used at every design stage involved in the complementary processes of designing a management information system and re-designing the empirical domain to both inform the design of the management information system and to ensure its successful implementation and integration into the empirical domain. The formalism which will be proposed is an object-oriented formalism based on Petri nets [AJM89], [AJM89], [AJM95], [AUD95], [HOL89], [JEN88], [JEN89], [LOP95], [PED94], [PET81], [PTR62], [ZUR94] and Adaptive Reference Technology [MAK94]. This formalism will be contrasted against the industry standard object-oriented formalism, the Unified Modeling Language (UML) [RAT97a], [RAT97b], which, it will be claimed, fails to satisfy many basic criteria necessary for both the design of software solutions and the re-engineering of business-processes.

A modelling formalism has to complement the objective of the modelling exercise. The objective of any modelling exercise is to build a model. This is true by definition, however the same modelling domain could be represented in a multitude of different models, where each model represents a particular viewpoint of the modelling domain. Thus, for example, an electrician's model of a house will be very different from an architect's model of the same house. Thus, when the purpose of the model is to facilitate the interrelated processes of re-engineering the modelling domain and designing a management information system, the model has to represent those elements of the domain that would be involved in those processes. Moreover, the model should represent only those processes that would be involved, if it is to be an effective abstraction and simplification.

In traditional conceptions of the models used in information system engineering, the domain is assumed to consist of functions that operate on sets of data and the sets of data themselves, which will be referred to in this discussion as datasets. These two types of entity are all that is necessary to define a system according to the traditional conception. From functions and datasets, it is possible to define system states as an ordered n-tuple composed of exactly one value from each of the n datasets. System behaviour over a period of time t may then be defined as a t-tuple composed of n-tuples.

This so-called function-oriented type of modelling formalism offers a simple and intuitive way of analysing modelling domains. However, as the modelling domain becomes increasingly complex, so the model too becomes proportionately more complex in models constructed using the function-oriented type of formalism. This increase in complexity does not become problematic for the analysis of relatively simple system, although when the complexity of the system increases to the size which typifies many of those which need to be modelled for the development of many management information systems, the complexity of the model can make it too cumbersome to be effective. This represents problems at every stage of the development process depicted in Figure 4.01 above.

In the initial process of black-box modelling, it can result in a long list of functions and datasets, each one requiring to be identified and related to at least one other system entity. This comes even more of a problem at the white-box modelling stage, where the number of functions and datasets is multiplied through the inclusion of the inside machinery of the black-boxes into the analysis.

The solution to this increase in complexity is to include another type of entity to allow for the intuitive organisation of functions and datasets. These new types of entity, called object classes, form the basis of so-called object-oriented modelling. The defining difference, therefore, between function-oriented modelling and object-oriented modelling is the presence or absence of object classes in the analysis of the system.

It is important to note here the distinction between object-oriented modelling and object-oriented programming. Evidently, as the name suggests, they both originate from the same school of thought, but object-oriented programming involves a far more expansive definition of object than does object-oriented modelling. For modelling, all that needs to be considered in the differentiation between the function-oriented approach and the object-oriented approach is the way that the system is decomposed into sets of system components: for the function-oriented approach, the system is decomposed into sets of components consisting of one function and one or more datasets; for the object-oriented approach, the system is decomposed into sets (normally) consisting of one object class and one or more function and one or more dataset. This statement shall be referred to as the Decomposition Hypothesis.

There are many claims made for object-oriented programming, and by extension object-oriented modelling, over its function-oriented equivalent. Many of these claims are difficult to assess, involving ill-defined or subjective psychometric measures, such as intuitive comprehensibility or degree of similarity to our perception of complexity in the real world. Before any of these claims can be examined in more detail, however, it is worth looking more closely at the concept of an object class.

According to the Decomposition Hypothesis, an object class is simply an organising entity, allowing functions and datasets to be grouped together into a set. It therefore plays only an abbreviating role in any formal decomposition of a system, as opposed to an extending role. That is, the behaviour of the system may be completely and consistently defined without object classes; functions and datasets are conjointly both necessary and sufficient to define system behaviour. It follows, therefore, that object classes must introduce into the decomposition process some additional degree of simplifications that could not be introduced with the use of functions and datasets alone. The origin of this simplification is often referred to as encapsulation.

Encapsulation is the notion that the constituent functions and datasets of object classes may not only mirror our perception of types of entity in the real world, but also encapsulate the functions and datasets that define those perceived real world entities into a single abstraction – an object class. For example, if we define an object class called *Cat*, then we can associate all of the functions and datasets that we use to define *Cat* processes (which can safely be called *cats*) as components of the object class *Cat*. This has proven to have enormous advantages in object-oriented programming because of the ease with which it allows for collaborative programming through programming tasks being able to be allocated and co-ordinated between programmers, where those programming tasks can be based on the encoding of object classes, which can largely be done without extensive knowledge of the code being generated for other object classes. It also allows for re-usability of code, since much of the code will be encapsulated into object classes and may be more easily plugged into other programs or deleted or edited within existing programs. With regards to object-oriented modelling, these advantages being brought about by encapsulation are equally applicable, although of course instead of 'programs' one should instead read 'models'.

So, how is it possible to measure this simplification in the model that is assumed to arise through the process of encapsulation? The obvious answer to this question is to use the Shannon concept of information by making the assumption that the 'simplification' referred to in the question can be considered synonymous with 'decrease in the amount of information necessary to specify the model." This assumption is used in the development of a method for estimating the amount of information necessary to specify a model of a system using both the function-oriented modelling approach and the object-oriented modelling approach included as the Appendix Information-Theoretic Evaluation of Object-Oriented System Representations.

The result of the study indicates that the object-oriented system representation is more efficient in modelling the complexity of systems in terms of the amount of information necessary to specify the system. However, this result only holds for systems beyond a certain level of complexity. For simpler systems, the function-oriented system representation is more efficient.

The question now arises as to whether or not most systems that are currently modelled using object-oriented formalisms fit within that class of system where object-oriented specification of the system is simpler in the manner defined above than a function-oriented specification.

There is an intuitive argument for this being the case: that it would be a co-incidence beyond reasonable explanation if systems had a 'naturally' even distribution of datasets amongst functions, in the case of function-oriented specifications, or an even distribution of datasets and functions amongst object classes in the case of object-oriented specifications.

Of course, the problem with this argument is the use of the term 'naturally'. There are as many ways to decompose a system as there are combinations of datasets and functions. Why should any one of these possible combinations be the one which is 'right' – that one which is 'naturally' occurring? The most that can be assumed in such circumstances without delving into the territory of metaphysics is that humans come equipped with a cognitive capacity which is evolutionary adapted to its environment, and that such evolutionary adaptation implies a particular way of decomposing the world and categorising its contents into distinct entities, and those distinct entities into associated functions and datasets.

Although there is, as demonstrated, an objective means for evaluating the two modelling paradigms, there is little chance of using any such objective method in evaluating modelling formalisms.

The industry-standard modelling formalism for object-oriented modelling is the Unified Modelling Language [RAT97a], [RAT97b], or UML. UML may be seen as a hybrid of various preceding formalisms [BOO91], [EMB92], [GRA93], [LEP94], [LOS94], all of which were based on the use of an object-class diagram as the most fundamental type of diagram. The object-class diagram is composed of a series of interconnected tables, where each table defines an object-class by listing its component datasets and functions, as well as the possible range of values that each dataset may be have. The links between the object-class tables represent any one of a various range of relationships. These relationships are normally as follows:

1. Mereological Relationship
2. Taxological Relationship
3. Associative Relationship

Conjointly, these three relationship types provide a powerful tool in the decomposition of a system. In particular, the taxological relationship allows for the concept of inheritance, whereby a child object-class is assumed to inherit all of the properties of its parent, without having to specify that this is the case. This is a very useful abbreviation, both in object-oriented modelling as well as object-oriented design. The relationship between the object-class vehicle and the object-class car in Figure 4.07 below is an example of such a relationship.

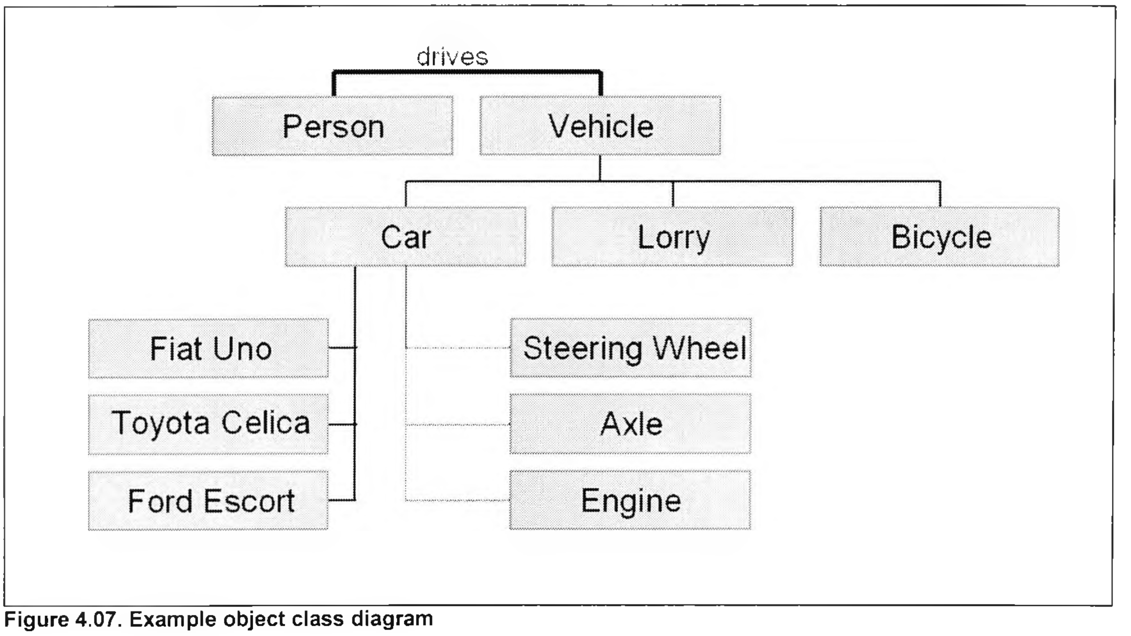

Figure 4.07. Example object class diagram

There are problems with object-class diagrams, however. First, they do not easily represent the dynamics of a system. Thus, for example, in the above diagram, how would one determine the process of making a journey from A to B? Because in object-oriented modelling, all processing is encapsulated within object classes, the dynamics of the system must be controlled by some other means than the output of one function being the input to another. The solution to this is to have message passing between objects. When one object has completed an operation using its data, it then passes a message to another object instructing it to perform another operation that logically follows from the completion of the first operation. Thus, the dynamics of the system are defined through message passing. The problem with this is that the dynamics of the system are more difficult to represent graphically, at least using the UML formalism.

In UML, system dynamics are represented by a series of diagrams; each one depicts the system dynamics from a particular perspective. One such UML diagram is the sequence diagram. The sequence diagram represents the system dynamics as messages passed between objects, where these objects have time lines so that the chronological ordering of the message passing can be tracked more easily. An example of a sequence diagram for the process of starting a car is shown in Figure 4.08 below.

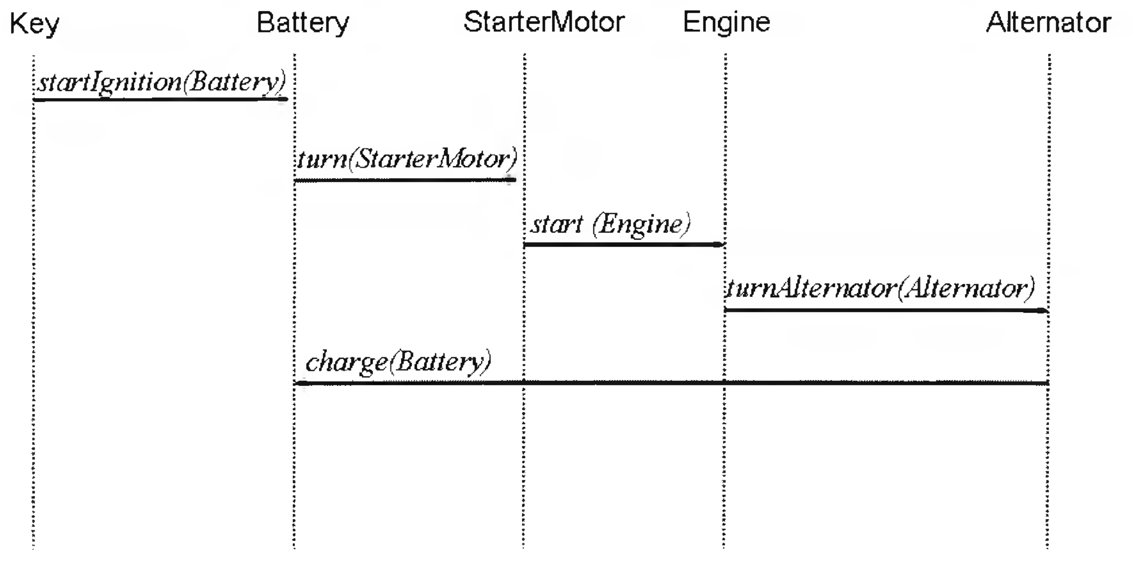

Figure 4.08. Example UML Sequence Diagram

There are other diagram types within the UML toolbox to represent dynamics, but they all suffer to some degree to some of the same issues of usability as the sequence diagram. One of the main issues of usability is the problem of readability. It is true that, in English at least, we read from left to right, and so it makes sense to have the flow of messages going in that direction. But the truth is that man instances of system dynamics are far more complex than the simple example above. In these cases, the diagram soon becomes unreadable, with message labels and arrows soon cluttering the diagram and reducing readability. This problem is increased as message passing becomes a conditional event, with the different conditions necessary for a message to be passed to one object rather than another needing to be included in the diagram as labels on arrows.

Usability also becomes a problem due to a more fundamental reason. With function-oriented programming it is possible to use the principle of functional decomposition to break the complexity of the system dynamics down into more simple components. Thus, for example, a system can be composed of one top-level function, such as, in the case of the healthcare system, treating patients. This function of treating patients is composed of several lower level functions, such as diagnosing a disease, administering treatment and monitoring the patient. And each one of these lower level functions may be themselves decomposed into more simple functions, and so on.

With object-oriented programming, there is no functional decomposition because functions are not the components used to define the system – object classes perform that role. So, it might be asked why it is not possible to have object class decomposition. Well, in fact there is object-class decomposition, and it is inheritance as described above, where a child object class inherits properties from a parent class. But, crucially, object class decomposition does not decompose the dynamics of the system since the exact same functions are inherited between child and parent. That is, the functions found in the child object class are not simplified versions of the functions found in the parent object class.

Thus, the UML approach faces a problem, which it is yet to be resolved. The answer lies in foregoing the concept of an object-oriented version of functional decomposition in favour of a more straightforward method of handling complexity. The functional decomposition of a system is what may be called a *vertical* form of complexity management, where complex objects are decomposed into more simple objects. But it is also possible to consider a *horizontal* form of complexity management, where the system components are all simple, but categorised into belonging to certain processing *loops*. And it is these loops that are then able to form the basis of managing complexity.

Loops may be simply defined as ordered sequences of functions and messages. Mathematically, there is no reason for a particular decomposition of a system composed of functions and messages into one set of component loops rather than another, just so long as each loop consists of contiguous functions and messages. Semantically, however, the system can be decomposed into loops that perform particular, higher-level, functions or areas of related functionality. For example, consider the process of starting the car, as depicted by a sequence diagram in Figure 4.08 above. Now, reading from Figure 4.07 above, we know that axle, wheel, and engine are all part of the locomotion system of the car; starter motor and ignition and battery are all part of the ignition system of the car. So, this gives us a basis for categorising the different loops involved in starting a car as to whether they are an ignition loop or a locomotion loop. Then, so long as we have a means of defining the interfaces between loops of different categories, we have a system of complexity management for object-oriented modelling.

This system of complexity management does not exist in UML sequence diagrams, nor any other dynamic diagram type within the UML toolbox, at least in any explicit or usable form. UML instead adopts the notion of Use Cases, which may be thought of as high-level system functions, such as, for example Start Car, for the car example above. The important issue regarding use cases, however, is that they are defined as user functions. Thus, the complexity of any Use Case cannot be further broken down in any systematic way into different sequence diagrams, for example. Moreover, since the decomposition of UML is based on use cases, which are defined in terms of users of the system, the modelling enterprise is necessarily restricted to adopting a particular modelling philosophy which, as argued above, is incompatible with a unified business process re-engineering and software engineering modelling formalism. It is reasonable, therefore, to adopt an alternative modelling formalism than UML to represent and manage the complexity of the system dynamics of object-oriented systems. This alternative modelling formalism will be proposed in the next section and will be Object-oriented Petri-nets.

There is still one further problem with the UML system. This is the problem of *implementation modelling*. The fact is that if we model a cat using UML or any other type of modelling formalism, it is probably because we want to encode the behaviour of the cat into some form of software application. Now, a piece of software is not a cat, no matter how accurate the simulation of the behaviour of the cat is by the underlying software code. It remains, ultimately, programming code that is implemented on a computer, just as the behaviour of the cat remains as certain electrochemical reactions that are implemented in a cat's body. The different is obviously crucial in any process of computerisation. The issue of implementation is included in UML, but it is not modelled in any dynamic diagram type. Thus, relating the process to its implementation(s) becomes difficult, having to correlate the information contained in two different diagrams. This problem is overcome in the alternative modelling formalism of Object-oriented Petri-nets, which is proposed in the next section.

4.3.1. Object-Oriented Petri-Nets

The modelling formalism which is proposed here shall be referred to as Object-oriented Petri-nets (OOPN). OOPN retains the object class as the primary definition of the system's internal data structure and thus is to be considered as an object-oriented modelling formalism. Unlike UML, however, it does not have a whole array of different diagram types to model the dynamics and implementation of the system – it has only two diagram types, both of which are based on Petri-nets. It uses the concept of loops to form the basis of a system of complexity management, and thus, unlike UML, does not suffer from the restrictions of Use Cases in terms of their ability to have their internal processing structure revealed and broken down into simpler components. Nevertheless, it still retains some of the properties of use cases, with the highest-level loops obviously corresponding to what would otherwise be use cases in a UML model of the system.

Petri –nets were first introduced in 1962 by Carl Petri [PTR62]. In their initial formulation, they consisted of a state-transition formalism with markings to generate behavioural properties. Since then, however, various extensions and abbreviations have been introduced to make Petri-nets a powerful modelling tool. An introduction to Petri-nets is provided in the Appendix 14.

In the formalism that will be proposed here, and subsequently used to model the computer-assisted patient scheduling system in the next chapter, there are four types of diagram:

- Process Diagram (Object View)
- Process Diagram (Processor View)
- Object Class Diagram
- Object Class Relationship Diagram

The latter two of these diagram types can also be found in large degree within the UML toolbox, as well as many other object-oriented formalisms. The first two diagram types, however, are specific to OOPN and utilise the formalism of Petri-nets to depict system dynamics from different perspectives.

4.3.2. Process Diagrams

Process diagrams utilise the Petri-net formalism as consequently comprise three basic modelling elements:

- **Places**, represented by blue rectangles with a curved edge, denote messages between object classes;
- **Transitions**, represented by rectangles with, denote object class processes;
- **Connecting Arcs**, represented by arrows connecting a place to a transition, or vice-versa, denote a message being passed between two object classes.

In addition to these three basic elements, two more elements are introduced, as follows:

- **Objects**, represented by blue rectangles and forming the background to a Petri-net or part of a Petri-net composed of the three basic elements, denote component system objects within which the processing represented by the Petri-net or part of a Petri-net of which it forms the background;
- **Processors**, represented by green rectangles and forming the background to a Petri-net or part of a Petri-net composed of the three basic elements, denote component system

processors that perform the processing represented by the Petri-net or part of a Petri-net of which it forms the background;

The purpose of process diagrams is to model system dynamics. Their closest UML equivalent is sequence diagrams. As with sequence diagrams they model system dynamics through showing messages that are passed between or within objects. Unlike sequence diagrams, however, they utilise the Petri-net formalism by representing the messages that control processing flow as places, and object class processes as transitions. In addition, they represent the different object classes and processors where the processing is being performed. This is done by placing parts of the Petri-net (or the entire Petri net) within rectangular representation of either objects or processors. This results in two different versions of process diagrams, depending on whether they represent the different object classes performing the processing or the different processors.

This inclusion of explicit representations of object classes and processors in system dynamic diagrams is important in two respects. First, it recognises that the consideration of *what* does the processing is important in BPR and software engineering; second, it means that the processing of data within a single object class need not be performed by the same processor, and vice-versa. This latter consideration is especially important for the design of information systems since the process of computerisation often involves enhancing system performance by integrating processing within single processing units, rather than distributed amongst a multiplicity of human and non-human processors.

Figure 4.09 below shows a process diagram from an object view – that is, by aggregating processing by object classes, rather than processors. In Figure 4.09, all processing occurs within the object class Object 1. In the first step, the object process labelled 1.01/P01 sends a message (1.01/01) to 1.01/R01, which then sends a return message (1.01/02) to 1.01/P01 which then sends another message (1.01/03) to the process 1.01/W01 which then sends a return message (1.01/04) to 1.01/P01. At that point, 1.01/P01 sends a message (1.01/05) to the object process 1.02/P02, and so on.

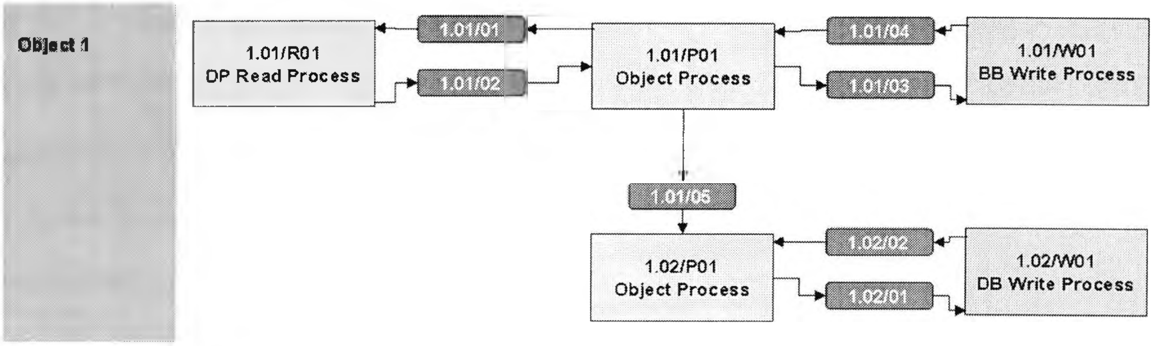

Figure 4.09. Example Process Diagram (Object View) representing a single-object process

There are several points to note from Figure 4.09 above. First, the labelling schema adopted for messages is not only able to identify processing loops within each Petri-net, but moreover can easily represent the processing flow and dependencies. The schema has three components:

1. The use case identifier
2. The process identifier
3. The message sequence

The use case identifier identifies within which use case the message occurs. In this context, a use case means much the same thing as it does in UML – a high-level system process or group of processes – and is represented by a Petri-net. The process identifier identifies which process generates the message. The message sequence number places the message within the ordered set of messages that defines the processing flow. Thus, for example, the message 1.01/02 is a message within the first use case ('1') and is generated by the first process within that use case ('01') and is the second message to be generated within the processing loop of the first process.

The second point to note is the distinction made above between object class processes represented by yellow boxes and database processes represented by orange boxes. Although this distinction has little meaning in UML, it is an important consideration when performing BPR exercises or designing information systems. The reason behind this claim lies in the fact that many processes in any information system are really just reading from or writing to a database. These processes are of course important, but do not represent any domain specificity since such processes are duplicated across many different information systems. Moreover, such processes are always the antecedent or consequent of the other type of processes shown in Figure 4.09 above – object processes. That is, they do not perform any *creative* tasks involving the creation or modification or deletion of data.

The inclusion of database processes is important because often the impetus behind BPR exercises or the design of new information systems is to reduce the degree to which the same object process required data to be read from or written to a multiplicity of different databases. This advantage will be highlighted in process diagrams from a processor view below.

Another point to note from Figure 4.09 above is that the actual dynamics of the system is represented not only by the numbering of the messages, but also with the layout of the diagram itself. Just as with sequence diagrams in UML, the flow of time is read from left to right and from top to bottom across and down the diagram. This makes the diagram much easier to read in terms of comprehending system dynamics and the processing dependencies involved.

Finally, and of critical importance in any large-scale systems engineering project, is the ability to incorporate complexity management tools into the system modelling process. This is achieved through the formalism allowing system analysis into Petri-net loops, with the message labelling schema being able to identify loops and provide the basis of cataloguing them into a library of system dynamics.

As with UML sequence diagrams, process diagrams can represent system dynamics as it occurs between multiple objects. Figure 4.09 above shows a single object class being responsible for all system dynamics. Figure 4.10 below is an example process diagram where system processing involves two objects.

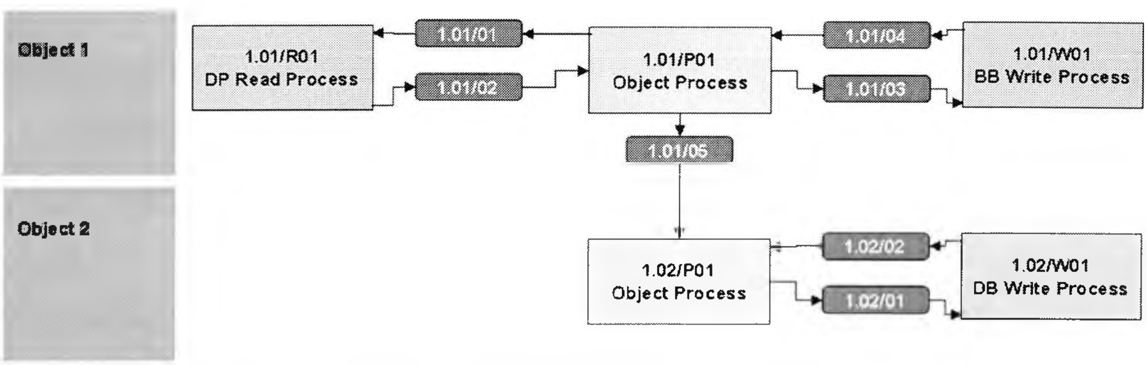

Figure 4.10. Example Process Diagram (Object View) representing a multiple-object process

In modelling processing flow, it is often necessary to include Boolean conjunctives where processing branches. In this respect, there are three basic possibilities with processing branching according to an OR, XOR or AND logic gate. In the case of an OR gate, the branch in processing will be simply represented by two unconnected arcs originating from the object process where the bifurcation occurs. This situation is shown in Figure 4.11 below, where process 1.01/P01 can either send the message 1.01/05.1 and/or the message 1.01/05.2. The extension to the labelling schema for messages required for the modelling of bifurcations is to label each branch sequentially after the message sequence number.

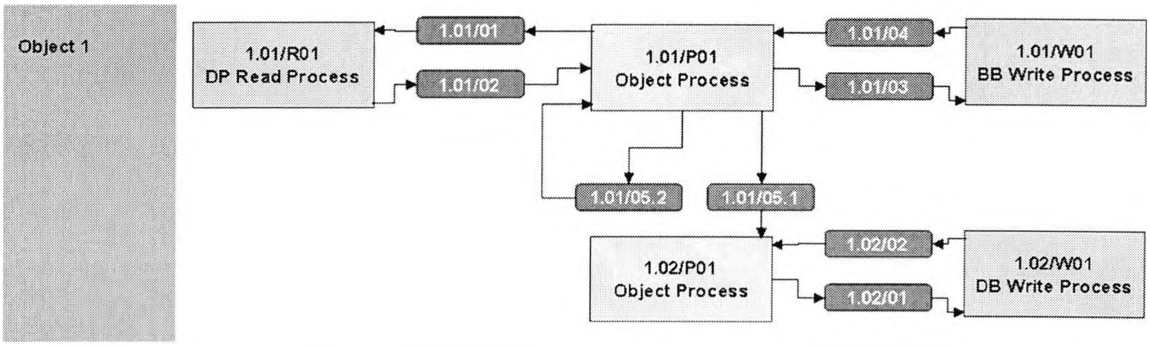

Figure 4.11. Example Process Diagram (Object View) representing a process with an OR logic gate

In the case of an AND logic gate, the conjoined messages are denoted by connecting their input arcs by a solid line as shown in Figure 4.12 below.

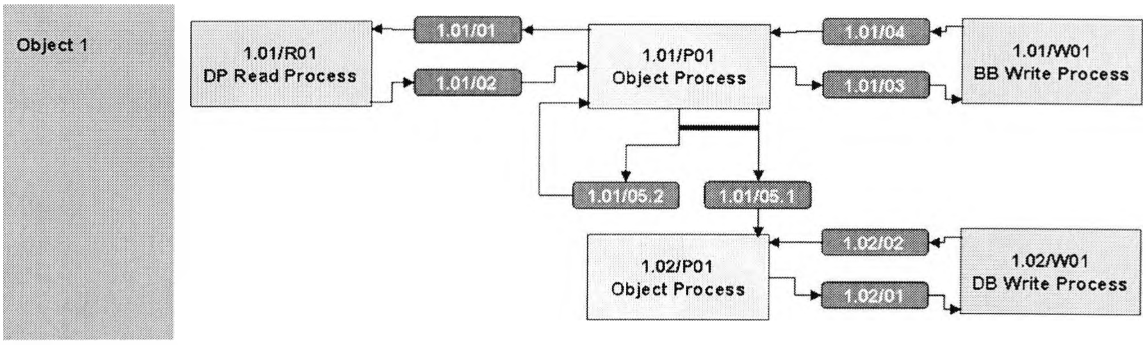

Figure 4.12. Example Process Diagram (Object View) representing a process with an AND logic gate

In the case of an XOR logic gate, the conjoined messages are denoted by connecting their input arcs by a dashed line as shown in Figure 4.13 below.

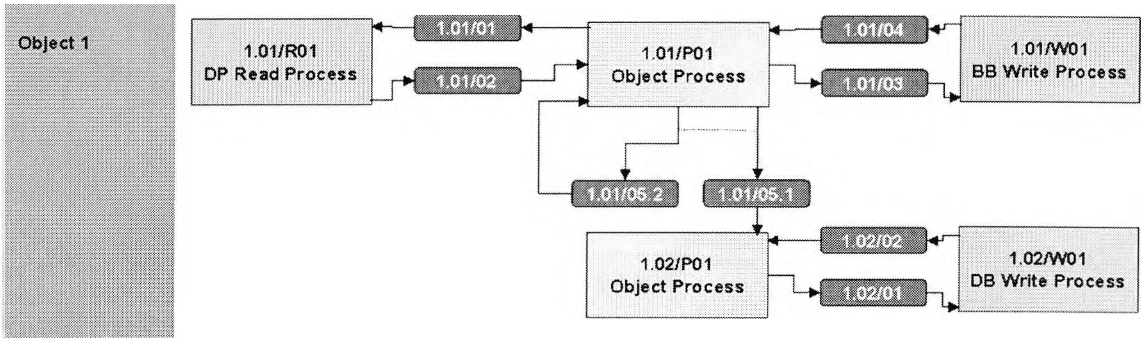

Figure 4.13. Example Process Diagram (Object View) representing a process with an XOR logic gate

Apart from being able to represent system dynamics in terms of the object classes whose data and processes comprise the processing, it is also possible to represent system dynamics in terms of the processors performing the processing. Process diagrams from a processor view involve the same Petri-net formalism and labelling schema as process diagrams from a object view with the exception that representations of processors replace that of objects. Processors are represented by green boxes as opposed to the blue boxes representing objects. An example of a process diagram from a processor view is shown in Figure 4.14 below.

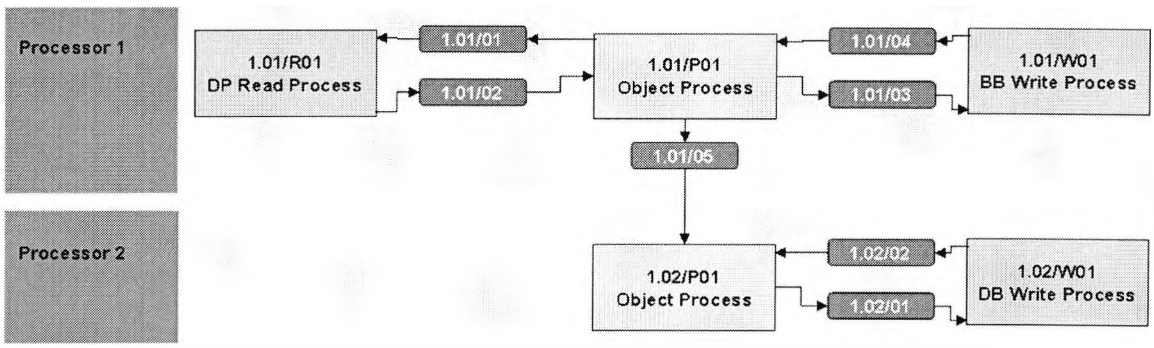

Figure 4.14. Example Process Diagram (Processor View) representing a multiple-processor process

An important point to note between the two different views of system dynamics represented by the two variants of process diagrams is that there need not be a co-extension between object and processors between corresponding process diagrams. That is, for any object, there is no requirement that its component processes be performed by a single processor, and vice-versa. It is this lack of necessary correspondence between object and processor which allows OOPN to be used in BPR exercises and information system design, since the identification of non-correspondence between object and processor can indicate fragmentation and lack of integration in system design.

4.3.3. Object Class Diagrams

The modelling formalism that will be used differs from UML only in so far as the representation of system dynamics is concerned. Thus, object class diagrams will be included as a type of diagram and will appear as the example given in Figure 4.15 above.

Object Name	
Attribute 01	Text String
Attribute 02	Category
Attribute 03	Binary
Attribute 04	Number
Attribute 05	Association
ProcessID1: processName1(ObjectClass)	

Figure 4.15. Example Object Class Diagram

As with UML and other object-oriented modelling formalisms, the object class diagram represents an object class in tabular form, with the header column being the name of the object, and the object classes attributes listed by name with the data properties of each attribute indicated. Finally, at the bottom of the table the object class's processes are listed with the process name and the object name on whose data the process modifies or uses placed in brackets after the process name.

4.3.4. Object Relationship Diagrams

Object relationship diagrams depict the numerical mapping relationships between different object classes. These relationships are as follows:

1. **0,1** = An object in the origin object class maps on to 0 or 1 objects in the destination object class
2. **1** = An object in the origin object class maps on to exactly 1 object in the destination object class
3. **0..*** = An object in the origin object class maps on to any number of objects in the destination object class
4. **1..*** = An object in the origin object class maps on to at least 1 objects in the destination object class

In object class relationship diagrams, object classes are represented by blue boxes and numerical mappings between object classes by connecting arcs with the numerical mappings between the two object classes placed at the ends of the arc. An example object class relationship diagram is shown in Figure 4.16 below.

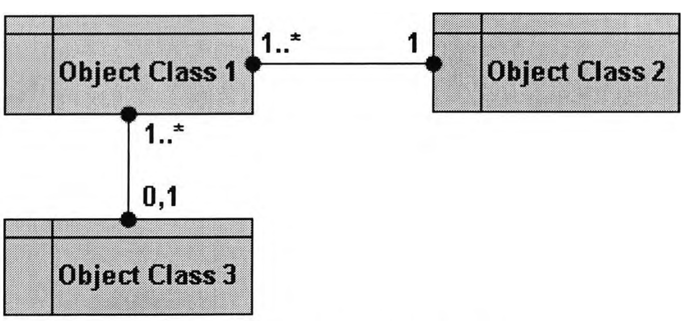

Figure 4.16. Example Object Class Relationship diagram

In Figure 4.16 above, an object in Object Class 1 maps on to exactly one object in Object Class 2 and 0 or 1 object in Object Class 3. An object in Object Class 2 maps on to at least 1 object in Object Class 1. An object in Object Class 3 maps on to at least 1 object in Object Class 1.

A useful extension to the object class relationship diagram is the inclusion of relationships between object classes and data stores. That is, where the data for each object class if stored. This consideration is something missing from formalisms such as UML since it is designed only for the design of software. When considering entire systems, however, it is often unrealistic to expect that the data defining a single object would reside in a single storage medium, let alone within a single data storage device.

The degree to which the data defining objects is distributed amongst multiple data storage devices may be considered as a good measure of the degree of system integration – the more data storage devices have to be accessed by an object to be able to process its data, then inevitably the slower and more inefficient that processing is likely to be. This is especially the case when there are multiple media used for data storage, such as paper-based filing systems, computer hard-disks and databases, or even human brains. For the purposes of system engineering, the term 'integration' therefore has a much more encompassing significance than it would have to a software engineer. For example, two database management systems (DMBS) that are able to communicate with each other and synchronise their records and tables automatically would be considered as an integrated system by the software engineer. But for the purpose of this discussion, the two DMBS are, in effect, a single DMBS.

Incorporation of data stores into object class relationship diagrams is straightforward, with each data storage device being represented by an orange box. A directed arc connecting an object to a data store represents that all or some of the data defining the object resides in the data store. Evidently, two directed arcs emanating from an object and connected to two different data stores represents a distribution of the object's data between the two data stores.

An example extended object relationship diagram is shown in Figure 4.17 below. In this diagram, for example, the data defining the Bed Slot object is distributed amongst the Scheduling DB and Census Data data stores, and so on.

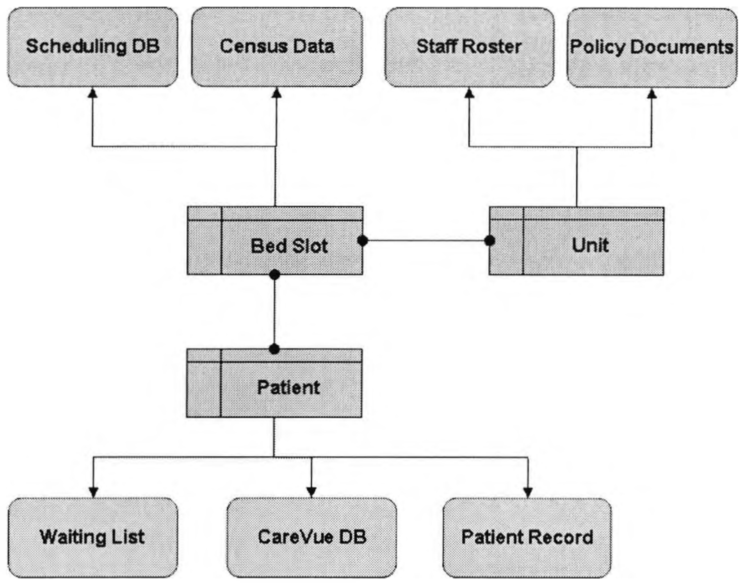

Figure 4.17. Example extended Object Class Relationship diagram

4.3.5. System Metrics

In software engineering it is common practice to calculate software metrics for the evaluation of code or productivity metrics for the evaluation of programmer productivity. In both these cases, the metric provides an objective and calculable measurement of performance.

In the case of taking measurements of the performance of systems, however, no well-defined objective measures exist. With the OOPN formalism, however, system metrics may be easily calculated and used objectively to compare different system designs and evaluate projected system performance and efficiency.

Two system metrics will be defined and will be used in the comparison of the models developed in the next chapter. These metrics are not integral to OOPN, but rather may be used alongside any systems modelling formalism that adopts an object-oriented approach.

Before these metrics are defined, some basic system parameters are introduced as follows:

1. D^n = The total number of data stores defined in system n
2. d^n_{MR} = The total number of manual database read processes defined in system n
3. d^n_{AR} = The total number of automated database read processes defined in system n
4. d^n_{MW} = The total number of manual database write processes defined in system n

5. d^n_{AW} = The total number of automated database write processes defined in system n

6. O^n = The total number of object classes defined in system n

7. R^n = The number of relations between data stores and object classes in system n

The first system metric to be defined is the DB Process Type Profile and measures the degree of automation in the system. It calculates the ratio between manual database processes and automated database processes and then multiplies this ratio by the sum of all database processes. This latter step in the calculation is important since a system may be redesigned to have a very low proportion of manual database processes, but only at the cost of dramatically increasing the total number of database processes. The equation for DB Process Type Profile is shown in Equation 4.01 below

Eq.4.01
$$\text{DB Process Type Profile} = \left(\frac{d^n_{AR} + d^n_{AW}}{d^n_{MR} + d^n_{MW}}\right)\left(d^n_{AR} + d^n_{MR} + d^n_{AW} + d^n_{MW}\right)$$

The second metric to be defined is the Object Class Fragmentation Rate and measures the degree to which the data of the system is distributed amongst different databases. Moreover, it also considers the degree to which this data distribution fragments the data of individual object classes through having the data defining an object being located in multiple data stores. The Object Class Fragmentation Rate is calculated as the ratio between the number of objects and the product of the number of relations and the number of data stores. The equation for the Object Class Fragmentation Ratio is shown in Equation 4.02 below

Eq.4.02
$$\text{Object Class Fragmentation Ratio} = \frac{O^n}{D^n R^n}$$

4.4. Summary

The objective of this chapter has been to develop and introduce a modelling approach and formalism which will be especially suited to the task of designing a resource management and patient scheduling software tool to be implemented in the RBH high-dependency environment.

The main arguments proposed in this chapter were that any modelling approach and formalism must recognise the fact that the design of the software cannot be done in isolation from considerations of the environment into which it is to be deployed. The reasoning for this conclusion was that the deployment of a new software system inevitably involves the modification of existing work practices and the ways in which the users of the software interact with it and incorporate it into their daily routine. Thus, the process of designing software should be considered in parallel with the process of redesigning the whole business environment and the modelling formalism used must therefore accommodate both objectives.

The industry standard formalism of UML was considered and rejected as a potential modelling formalism on the grounds that it did not recognise the need for a hybrid approach to software design and business process re-engineering. This was especially the case for the modelling of system dynamics using UML diagrams. Instead a new formalism was developed which was informed by the needs of the proposed hybrid modelling approach. This new formalism was Object-oriented Petri-nets (OOPN).

In the next chapter the discussions of this chapter will come to fruition in the form of two models of the RBH high-dependency environment. In accordance with the hybrid modelling approach, the first model will be the current operational model; the second will be the proposed operational model. Both models will be developed using OOPN.

5. Operational Models

The objective of this chapter is to use the modelling approach and formalism of Chapter 4 to develop a set of operational models that specify a computerised information system for use in the resource allocation process in the RBH high-dependency environment, and where the design of this information system supports the hypothesis that enhancing the level of control that healthcare managers have over the resource allocation process has the potential to increase the cost-effectiveness of healthcare delivery.

In Chapter 2, a mathematical model was developed which demonstrated the relationship between the level of control over of healthcare resource allocation and the ability to increase the cost-effectiveness of healthcare delivery. In Chapter 3 the empirical domain of the RBH High-Dependency Environment was examined in terms of its clinical characteristics and an evaluation of its current degree of its control over the resource allocation process. This evaluation was based on a classification of different types of control deficiencies that provided an empirical basis for the claim that the level of control over resource allocation in the RBH high-dependency environment is sub-optimal, and that therefore any information system which enhances the level of control over resource allocation will, if the hypothesis that increased level of control over resource allocation allows for an increased cost-effectiveness of healthcare delivery is correct, increase the cost-effectiveness of healthcare delivery.

In Chapter 4, a modelling approach and formalism was developed which was argued to be especially suited to the task of modelling an information system which is capable of supporting the hypothesis that greater levels of control over healthcare resource allocation is capable of increasing the cost-effectiveness of healthcare delivery. This chapter will present two models using the modelling approach and formalism developed in Chapter 4 to design such a system.

This chapter is divided into three sections. The first two sections present the current and proposed operational models of RBH, respectively. That is, the current process of patient scheduling, and the process as it would be following the deployment of a computer-assisted patient scheduling system (CAPSS). Both models are structured in the same manner and use the same formalism to aid in comparison. The first model – the current operational model – presents the problem domain as it currently exists. That is, before the deployment of any information system designed to enhance cost-effectiveness. The second model – the proposed operational model – presents the problem domain as it would exist post-deployment of an information system designed to enhance cost-effectiveness.

The final section of this chapter provides a discussion of the pertinent points from each model, followed by a comparative evaluation of the two models.

The presentation of each model begins with the definition of the component object classes. This is followed by a listing and overview summary of all of the different databases which are identified in the model. The component processes of each object class are then exhaustively listed and described both in plain language, as well as in terms of the data flows which they comprise. Finally, the process diagrams are shown for each process. In each case, the object view of the process diagram is presented followed by the processor view of the process diagram. The listing and description of all the component attributes are provided in plain language terms in Appendix 11.

5.1. RBH Current Operational Model (COP)

5.1.1. Introduction

The purpose of this model is to represent the object classes, databases, and component processes and data sets of the problem domain as it currently exists. In so doing the basis will be formed for comparing this model with the proposed operational model to be developed in the next section.

5.1.2. Object Classes

Bed Slot

The Bed Slot object class contains all data and processes related to the individual bed slot within the high-dependency environment. This includes associations to the patient object to which the bed slot is allocated, and the unit object of which the bed slot object is a component.

The Bed Slot object class appears in both RBH COP and RBH POP models. It is shown below in standard object oriented representation.

COP Bed Slot	
COP Actual Bed Slot Status Time = [T]	Category
COP Actual Labour Component [N] Status Time = [T]	Binary
COP Actual TISS Component [N] Status Time = [T]	Binary
COP Bed Slot Unit Name	Association
COP Patient Scheduled Monotonic Bed Slot Time = [T]	Association
COP Patient Scheduled Non-Monotonic Bed Slot Time = [T]	Association
COP Patient Type [P] Admissible	Binary
COP Projected Bed Slot Status Time = [T]	Category
COP Projected Labour Component [N] Status Time = [T]	Binary
COP Bed Slot ID	Text String
COP2.04: schedulePreAdmissionPatient(Bed Slot)	
COP2.06: allocateUnAllocatedBedSlot(Bed Slot)	
COP3.02: updateAllocatedBedSlot(Bed Slot)	
COP4.02: deallocateAllocatedBedSlot(Bed Slot)	

Figure 5.01. COP Bed Slot object class diagram

PATIENT

The Patient object class contains all data and processes related to the individual patient. This includes both clinical, demographic and economic data.

The Patient object class appears in both RBH COP and RBH POP models. It is shown below in standard object oriented representation.

COP Patient	
COP Actual Admission Time Unit [U]	Date/Time
COP Actual Length of Stay Unit [U]	Number
COP Admitting Consultant	Text String
COP Admitting Surgeon	Text String
COP Patient Admission Diagnosis	Category
COP Patient Clinical Attribute [N]	Text String
COP Patient Clinical Attribute [N] Time = T	Number
COP Patient Current Diagnosis	Category
COP Patient Current Location	Text String
COP Patient Date of Birth	Date/Time
COP Patient Demographic Attribute [N]	Text String
COP Patient Home Address	Text String
COP Patient Hospital Admission Date	Date/Time
COP Patient Name	Text String
COP Patient Projected Discharge Time	Date/Time
COP Patient Scheduling Status	Category
COP Patient Hospital Number	Text String
COP2.01: createPreAdmissionPatientRecord(Patient)	
COP2.02: updatePreAdmissionPatentRecord(Patient)	
COP2.03: updateWaitingList(Patient)	
COP2.05: admitPreAdmissionPatient(Patient)	
COP3.01: updateTreatmentPatient(Patient)	
COP4.01: transferTreatmentPatient(Patient)	

Figure 5.02. COP Patient object class diagram

UNIT

The Unit object class contains all data and processes related to the individual healthcare unit within the high-dependency environment. The Unit object class is primarily a data object class, containing many derived variables related to the bed slot objects which partly compose the unit object through association.

The Unit object class appears in both RBH COP and RBH POP models. It is shown below in standard object oriented representation.

COP Unit	
COP Accept Patient Type [P] From Unit [U]	Binary
COP Actual Labour Component [N] per Bed Slot Time = [T]	Number
COP Actual Number Allocated-Occupied Bed Slots Time = [T]	Number
COP Actual Number Allocated-Unoccupied Bed Slots Time = [T]	Number
COP Actual Number Unallocated-Unoccupied Bed Slots Time = [T]	Number
COP Actual Occupancy Rate Time = [T]	Number
COP Actual TISS Component [N] per Bed Slot Time = [T]	Number
COP Clinical Director Name	Text String
COP Discharge Patient Type [P] To Unit [U]	Binary
COP Maximum Actual Number Bed Slots	Number
COP Maximum Actual Number Bed Slots Patient Type [P]	Number
COP Mean Actual Occupancy Rate Period = [P]	Number
COP Operational Manager Name	Text String
COP Projected Labour Component [N] per Bed Slot Time = [T]	Number
COP Target Labour Component [N] per Bed Slot	Number
COP Target TISS Component [N] per Bed Slot	Number
COP Unit Name	Text String
COP1.01: createPolicy(Unit)	

Figure 5.03. COP Unit object class diagram

The object relationship diagram for the three object classes is as shown in Figure 5.04 below:

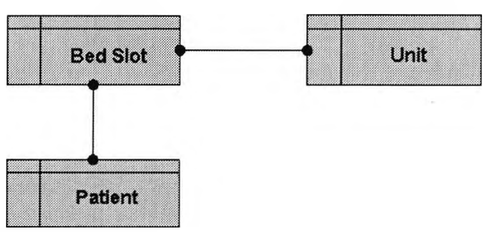

Figure 5.04. COP object class relationship diagram

5.1.3. Databases

CAREVUEDB

The CareVue Database is an integrated component of the CareVue medical information system. It records many fields of the patients medical condition, as well as many demographic and economic data fields. CareVueDB duplicates much of the data which is contained in PatientRecord.

CareVueDB is a component database of both the current and operational models.

CensusData

CensusData is an informally defined database consisting of various documents and files in both paper and computerised format. Responsibility for CesnsusData is distributed and ill-defined amongst members of both clinical and non-clinical members of staff.

CensusData is a component of only the current operational model.

PatientRecord

The Patient Record is the enduring source of each patient's medical history. It is purely paper based and duplicates much of the data which is also contained in other clinical and demographic databases. In particular, CareVueDB, the patient record from wh

PatientRecord is a component database of both the current and operational models.

PolicyDocuments

PolicyDocuments are not a database in any recognizable sense. They consist of various computer files and documents, often with no formal tabular representation. Conjointly, they define the operational constraints and parameters of the high-dependency environment.

PolicyDocuments is a component of only the current operational model.

SchedulingDB

SchedulingDB is a set of various databases and other files which record which patients will be, or were, admitted to each unit, and when.

SchedulingDB is a component of only the current operational model.

StaffRoster

StaffRoster is a set of rosters in tabular format which are normally maintained and updated from within each unit and specifies which members of clinical staff are on duty throughout the week.

StaffRoster is a component of only the current operational model.

WaitingList

WaitingList is really a set of different waiting lists maintained by each surgeon individually. Each list records which of the surgeon's patients are to undergo which surgical procedure, and in which order of priority. WaitingList is a component of only the current operational model.

The relationship between the data which defines the objects within each of the object classes and the databases where that data is recorded and stored, is as shown in Figure 5.05below:

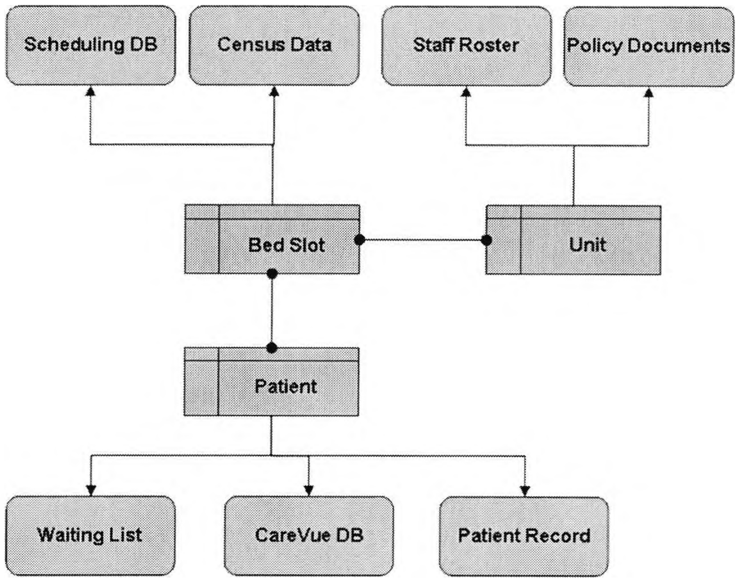

Figure 5.05. COP extended object class relationship diagram

5.1.4. Processes

COP1.01: CREATEPOLICY(UNIT)

COP1.01: createPolicy(Unit) is a component process of the Create Policy process group. It creates a new set of policy specifications or updates existing policy specifications determining parameters such as which patients may be admitted to the unit, the resource profiles of the unit, and so on. It is a component process of the Unit object class and is performed by the processor StrategicManager. It is composed of the following database processes:

Database Process	DB Process Type	Database Name
COP1.01/R01	Manual Read	CensusData
COP1.01/R02	Manual Read	PatientRecord
COP1.01/R03	Manual Read	StaffRoster
COP1.01/R04	Manual Read	PolicyDocuments
COP1.01/W01	Manual Read	StaffRoster
COP1.01/W02	Manual Write	PolicyDocuments

Table 5.01. COP1.01: createPolicy(Unit) database processes

COP2.01: CREATEPATIENTRECORD(PATIENT)

COP2.01: createPatientRecord(Patient) is a component process of the Schedule Patient process group. It creates the hospital patient record when the patient is first admitted to the hospital environment. It is a component process of the Patient object class and is performed by the processor ClinicalStaff. It is composed of the following database processes:

Database Process	DB Process Type	Database Name
COP2.01/W01	Manual Write	PatientRecord

Table 5.02. COP2.01: createPatientRecord(Patient) database processes

COP2.02: UPDATEPATENTRECORD(PATIENT)

COP2.02: updatePatentRecord(Patient) is a component process of the Schedule Patient process group. It updates the hospital patient record which was created when the patient was first admitted to the hospital environment. It is a component process of the Patient object class and is performed by the processor ClinicalStaff. It is composed of the following database processes:

Database Process	DB Process Type	Database Name
COP2.02/R01	Manual Read	PatientRecord
COP2.02/W01	Manual Write	PatientRecord

Table 5.03. COP2.02: updatePatientRecord(Patient) database processes

COP2.03: UPDATEWAITINGLIST(PATIENT)

COP2.03: updateWaitingList(Patient) is a component process of the Schedule Patient process group. It updates the waiting list by entering the patient onto the waiting list or updating the entry of an existing patient. It is a component process of the Patient object class and is performed by the processor ClinicalStaff. It is composed of the following database processes:

Database Process	DB Process Type	Database Name
COP2.03/R01	Manual Read	PatientRecord
COP2.03/W01	Manual Write	WaitingList

Table 5.04. COP2.03: updateWaitingList(Patient) database processes

COP2.04: SCHEDULEPATIENT(BED SLOT)

COP2.04: schedulePatient(Bed Slot) is a component process of the Schedule Patient process group. It schedules a patient for admission to the high dependency environment by non-monotonically allocating an unallocated bed slot to a patient. It is a component process of the Bed Slot object class and is performed by the processor OperationalManager. It is composed of the following database processes:

Database Process	DB Process Type	Database Name
COP2.04/R01	Manual Read	WaitingList
COP2.04/R02	Manual Read	PatientRecord
COP2.04/R03	Manual Read	PolicyDocuments
COP2.04/R04	Manual Read	CensusData
COP2.04/R05	Manual Read	StaffRoster
COP2.04/W01	Manual Write	SchedulingDB

Table 5.05. COP2.04: schedulePatient(Bed Slot) database processes

COP2.05: ADMITPATIENT(PATIENT)

COP2.05: admitPatient(Patient) is a component process of the Schedule Patient process group. It enters the patients details on the CareVue medical information system once the patient is admitted to the high dependency environment. It is a component process of the Patient object class and is performed by the processor ClinicalStaff. It is composed of the following database processes:

Database Process	DB Process Type	Database Name
COP2.05/R01	Manual Read	PatientRecord
COP2.05/W01	Manual Write	CareVueDB

Table 5.06. COP2.05: admitPatient(Patient) database processes

COP2.06: ALLOCATEBEDSLOT(BED SLOT)

COP2.06: allocateBedSlot(Bed Slot) is a component process of the Schedule Patient process group. It updates the bed slot object by changing the allocation status of the bed slot to being monotonically allocated to a patient. It is a component process of the Bed Slot object class and is performed by the processor ClinicalStaff. It is composed of the following database processes:

Database Process	DB Process Type	Database Name
COP2.06/R01	Manual Read	SchedulingDB
COP2.06/W01	Manual Write	CensusData

Table 5.07. COP2.06: allocateBedSlot(Bed Slot) database processes

COP3.01: UPDATEPATIENT(PATIENT)

COP3.01: updatePatient(Patient) is a component process of the Treat Patient process group. It continually updates both the paper-based and CareVue medical records of the patient as the treatment process continues. It is a component process of the Patient object class and is performed by the processor ClinicalStaff. It is composed of the following database processes:

Database Process	DB Process Type	Database Name

Database Process	DB Process Type	Database Name
COP3.01/R01	Manual Read	CareVueDB
COP3.01/R02	Manual Read	PatientRecord
COP3.01/W01	Manual Write	CareVueDB
COP3.01/W02	Manual Write	PatientRecord
COP3.01/W03	Automatic Write	CareVueDB

Table 5.08. COP3.01: updatePatient(Patient) database processes

COP3.02: updateBedSlot(Bed Slot)

COP3.02: updateBedSlot(Bed Slot) is a component process of the Treat Patient process group. It continually updates the bed slot status according to the projected duration of the allocation or recording any change in status. It is a component process of the Bed Slot object class and is performed by the processor ClinicalStaff. It is composed of the following database processes:

Database Process	DB Process Type	Database Name
COP3.02/R01	Manual Read	PatientRecord
COP3.02/R02	Manual Read	CareVueDB
COP3.02/R03	Manual Read	SchedulingDB
COP3.02/W01	Manual Write	CensusData

Table 5.09. COP3.02: updateBedSlot(Bed Slot) database processes

COP4.01: transferPatient(Patient)

COP4.01: transferPatient(Patient) is a component process of the Transfer Patient process group. It updates both paper-based and CareVue medical records to reflect the transfer of the patient from one unit to another. It is a component process of the Patient object class and is performed by the processor ClinicalStaff. It is composed of the following database processes:

Database Process	DB Process Type	Database Name
COP4.01/R01	Manual Read	SchedulingDB
COP4.01/R02	Manual Read	CensusData
COP4.01/R03	Manual Read	StaffRoster
COP4.01/W01	Manual Write	PatientRecord
COP4.01/W02	Manual Write	CareVueDB

Table 5.10. COP4.01: transferPatient(Patient) database processes

COP4.02: deallocateBedSlot(Bed Slot)

COP4.02: deallocateBedSlot(Bed Slot) is a component process of the Transfer Patient process group. It updates the status of the bed slot to reflect the transfer of the patient from the unit, thus de-allocating the bed slot to the patient, and possibly allocating it to another patient. It is a component process of the Bed Slot object class and is performed by the processor ClinicalStaff. It is composed of the following database processes:

Database Process	DB Process Type	Database Name
COP4.02/R02	Manual Read	CensusData
COP4.02/W02	Manual Write	SchedulingDB

Table 5.11. COP4.02: deallocateBedSlot(Bed Slot) database processes

5.1.5. Process Diagrams

RBH COP: CREATE POLICY (OBJECT VIEW)

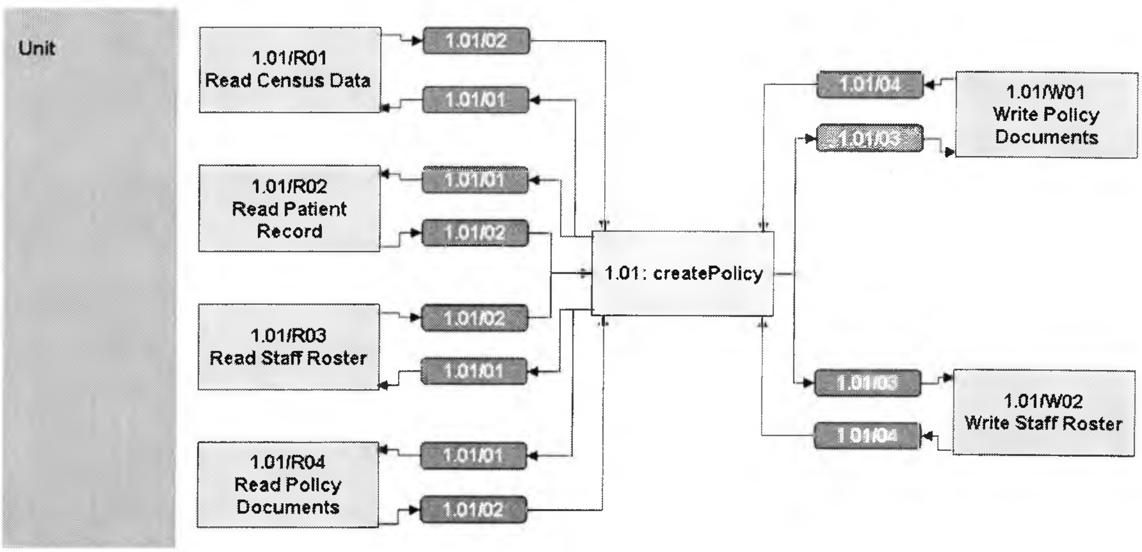

Figure 5.06. RBH COP: Create Policy process diagram (object view)

RBH COP: CREATE POLICY (PROCESSOR VIEW)

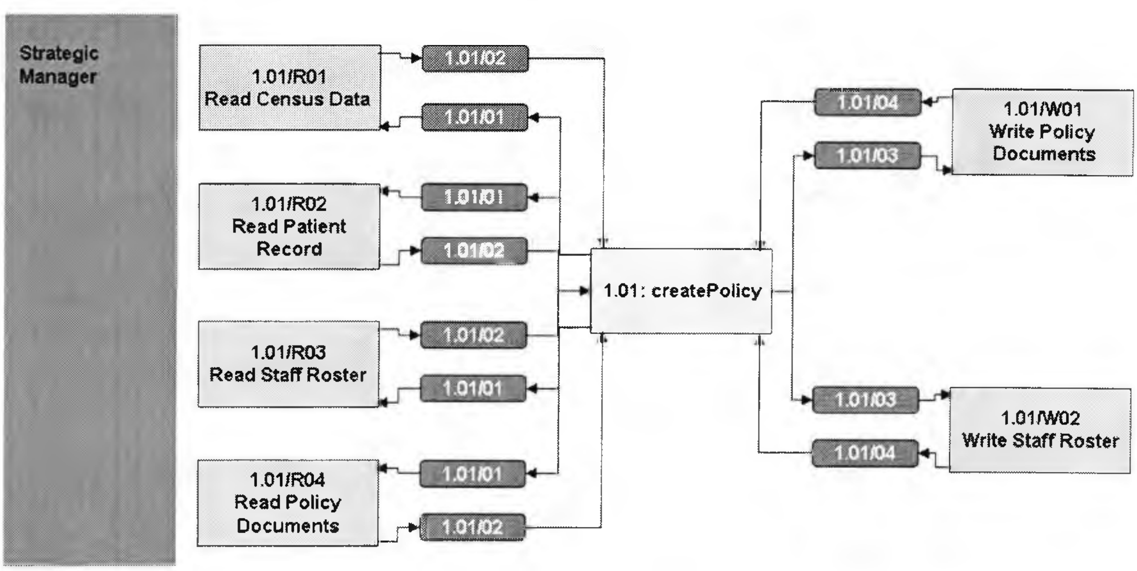

Figure 5.07. RBH COP: Create Policy process diagram (processor view)

RBH COP: Schedule Patient (Object View)

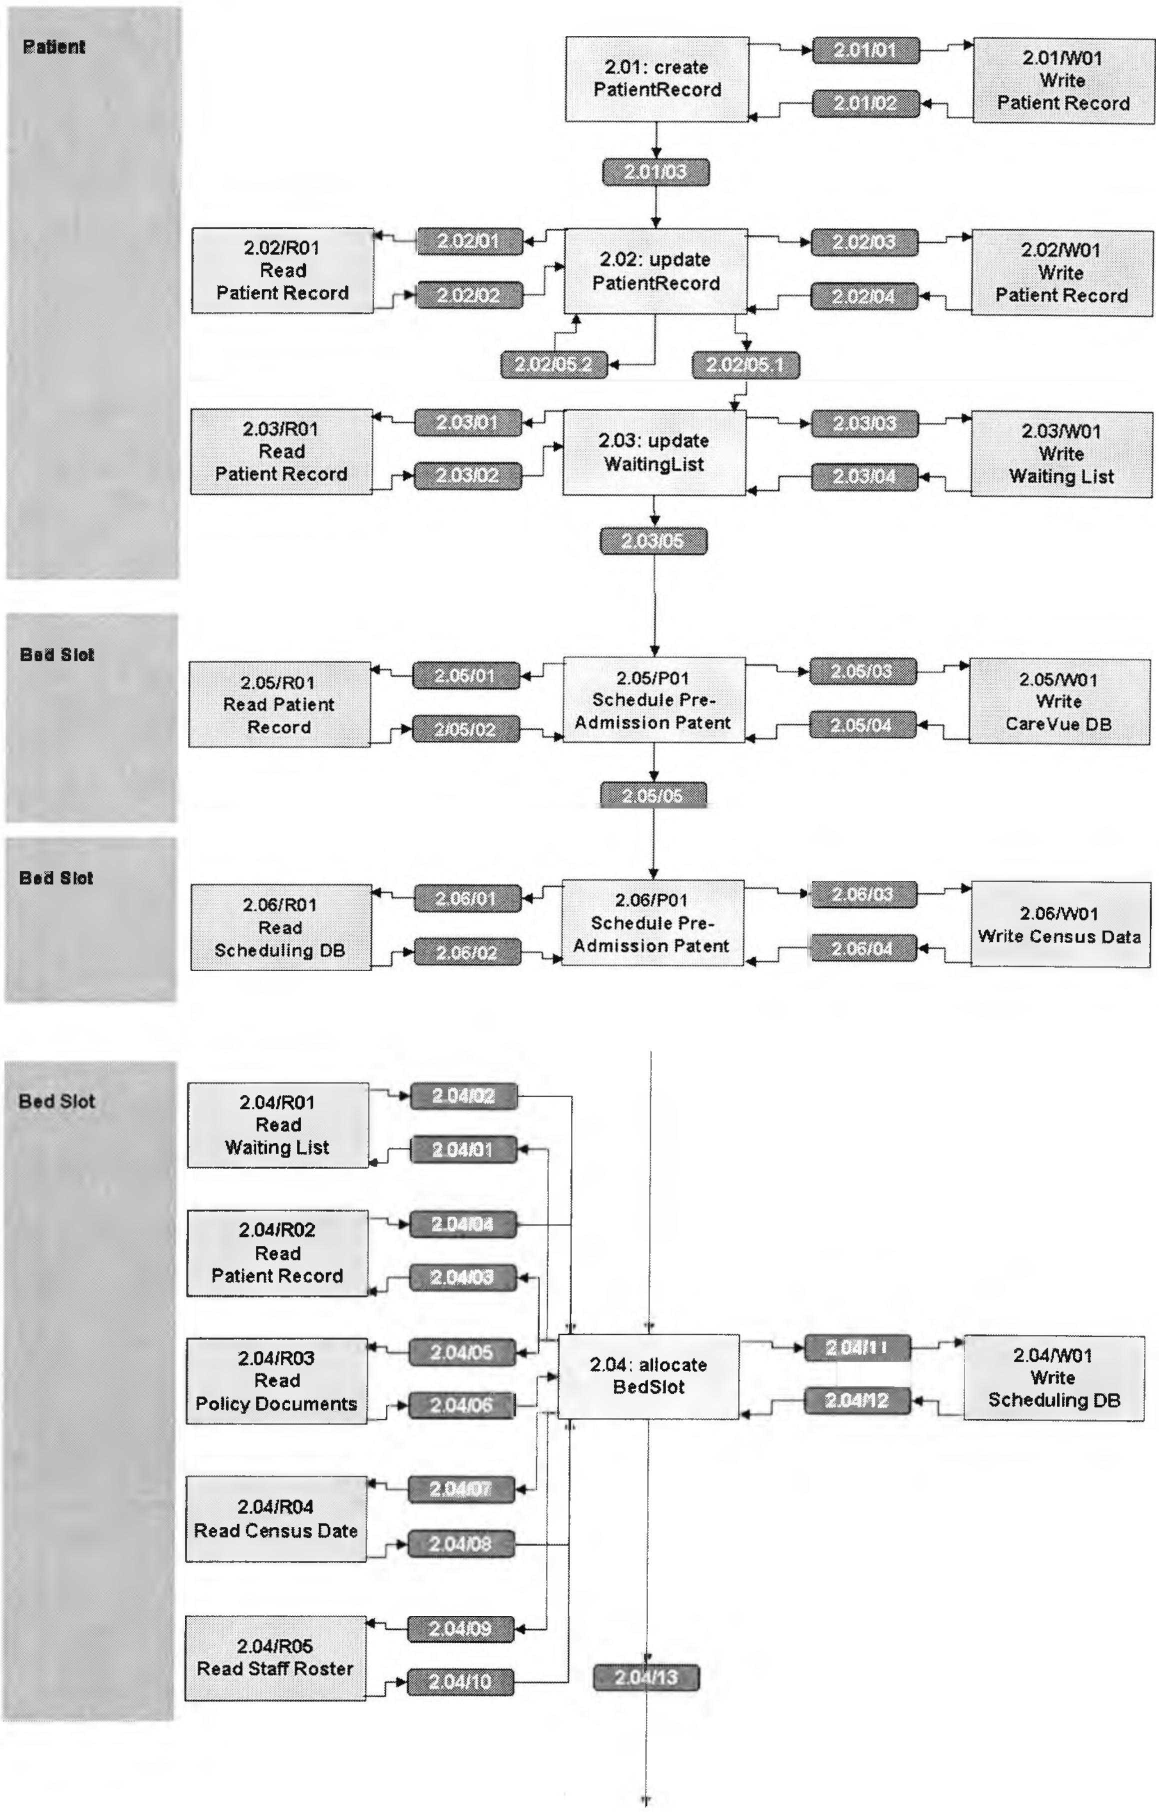

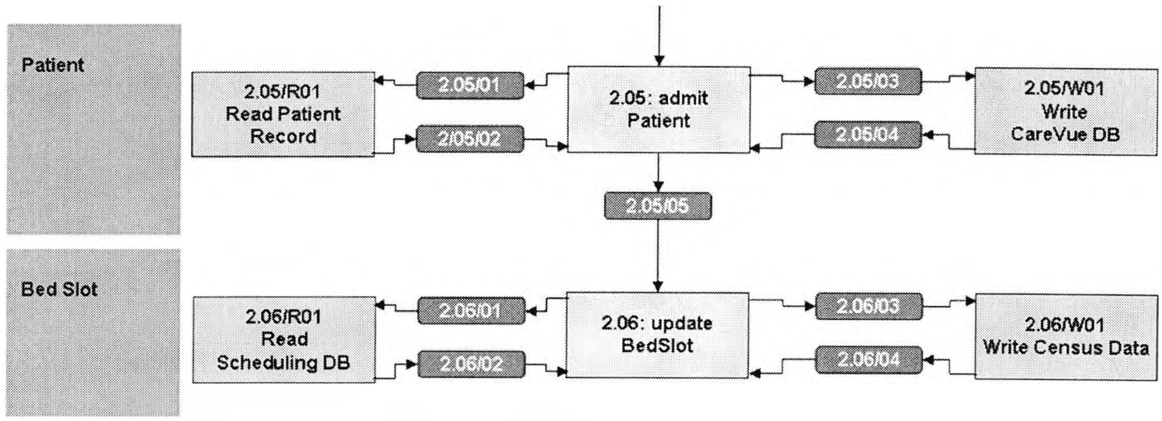

Figure 5.08. RBH COP: Schedule Patient process diagram (object view)

RBH COP: SCHEDULE PATIENT (PROCESSOR VIEW)

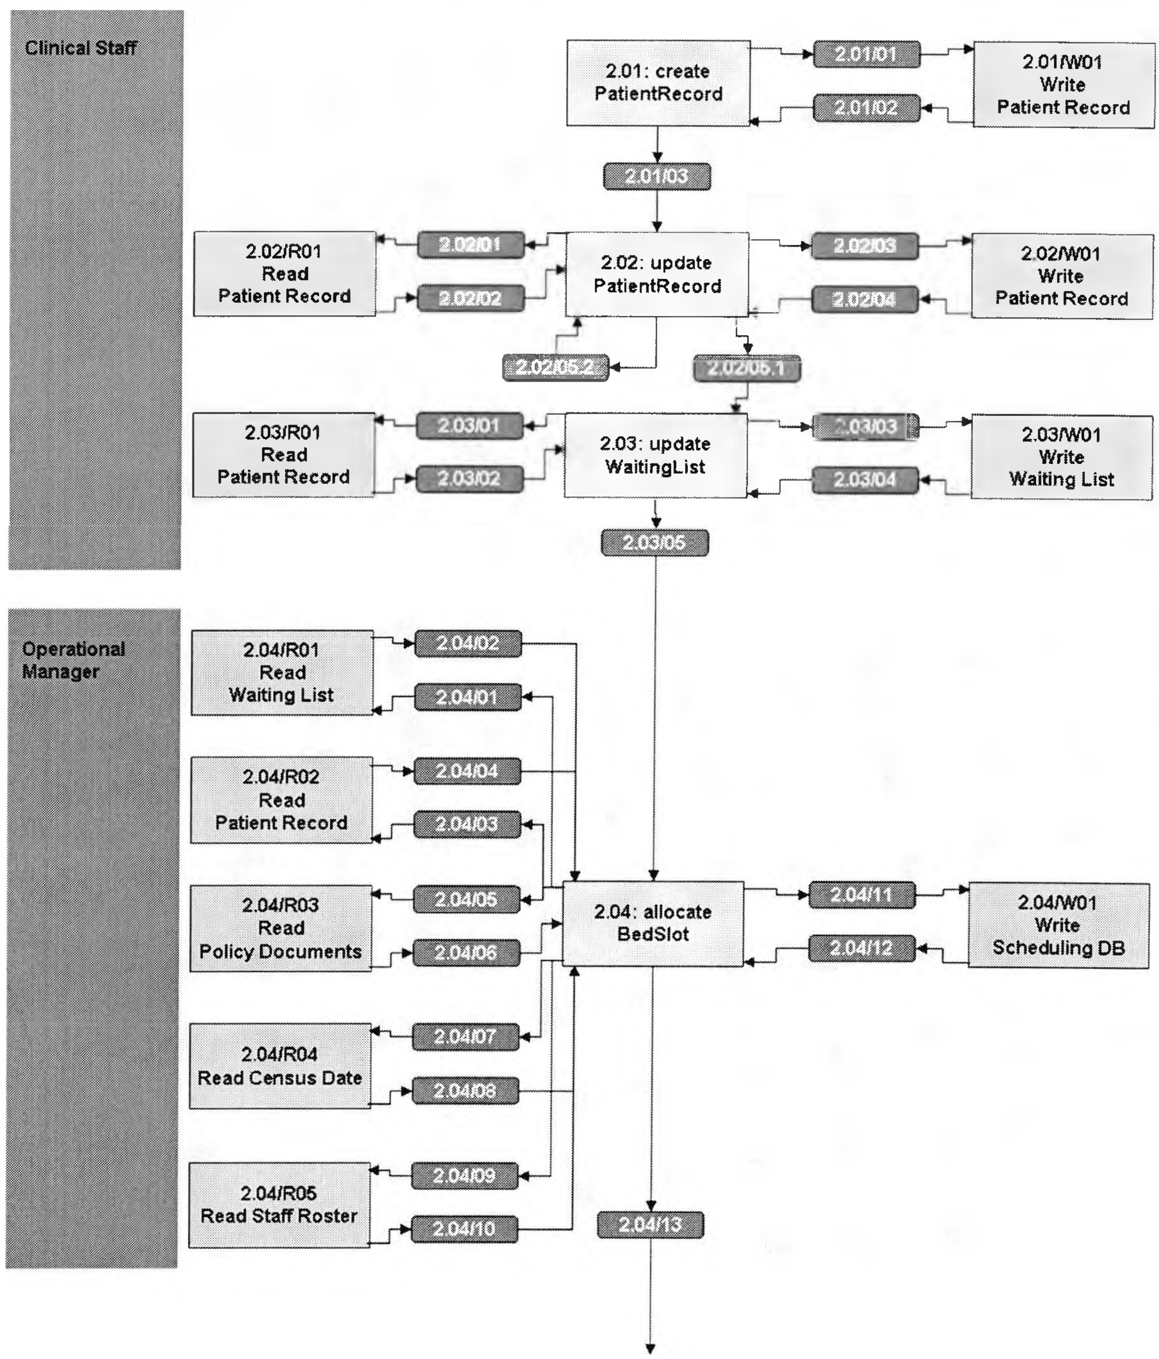

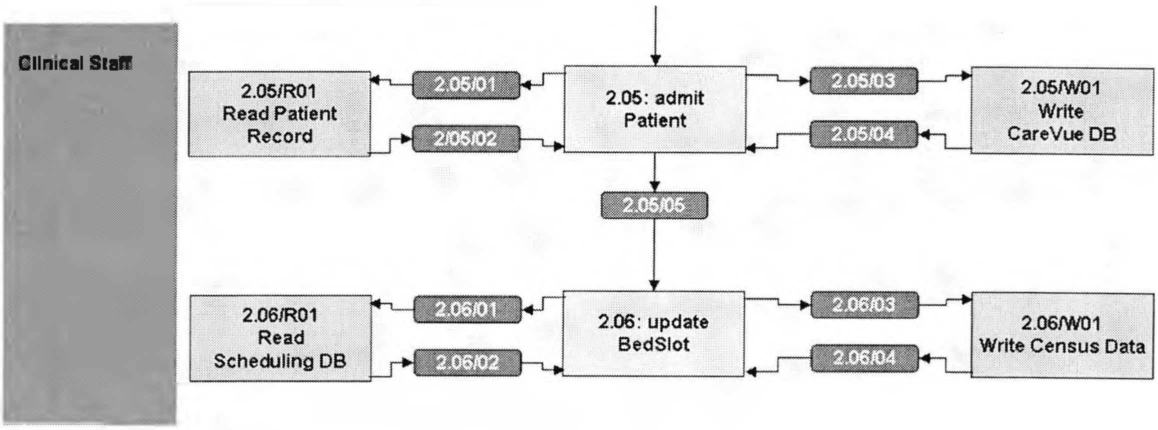

Figure 5.09. RBH COP: Schedule Patient process diagram (processor view)

RBH COP: Treat Patient (Object View)

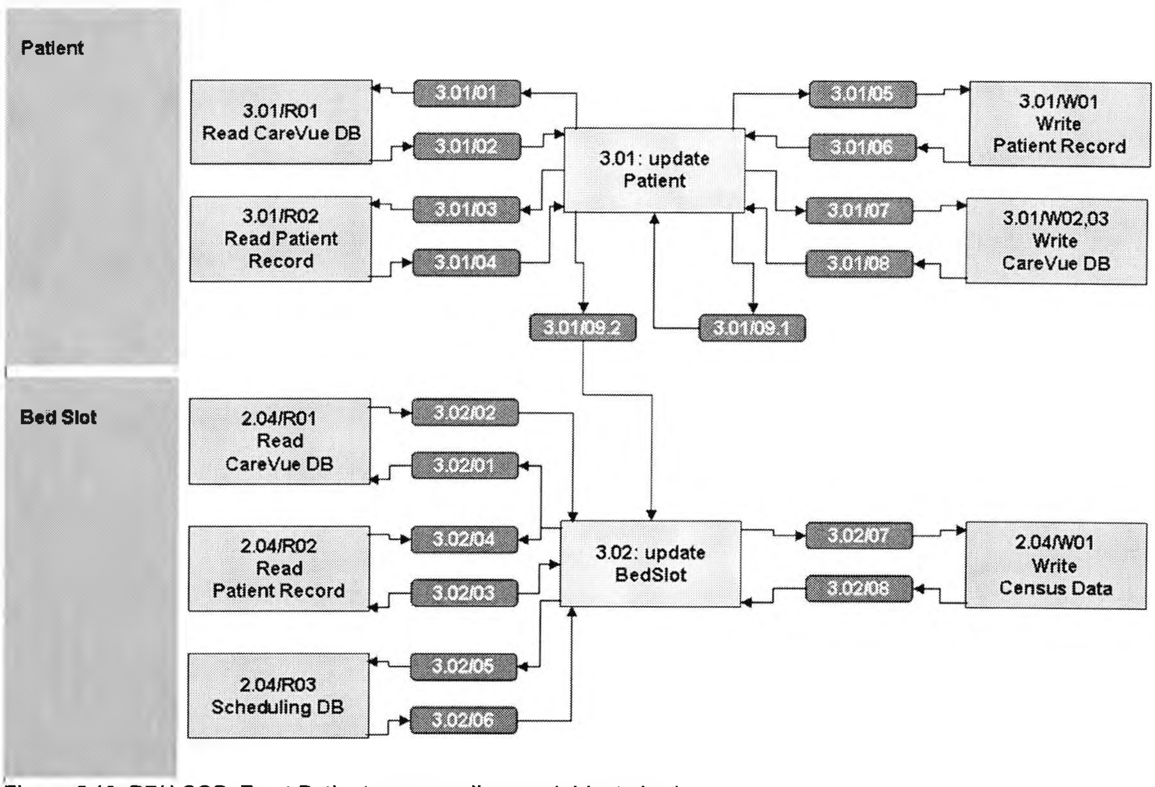

Figure 5.10. RBH COP: Treat Patient process diagram (object view)

RBH COP: Treat Patient (Processor View)

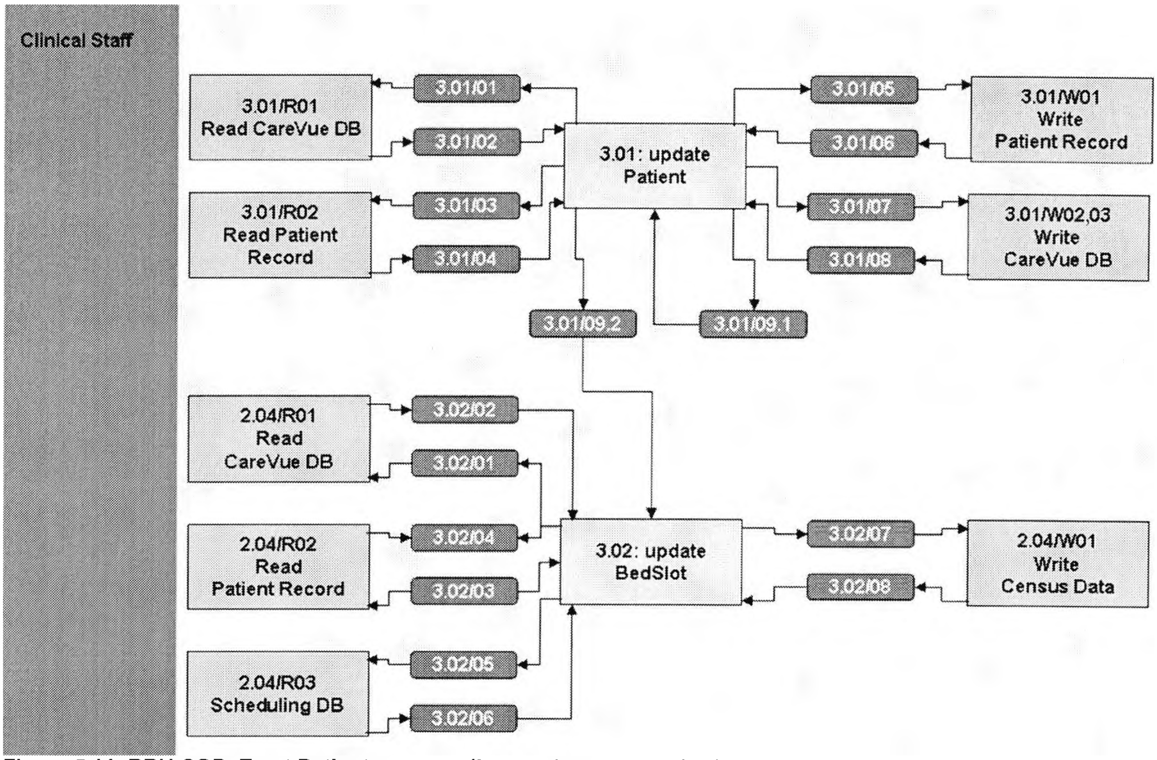

Figure 5.11. RBH COP: Treat Patient process diagram (processor view)

RBH COP: Transfer Patient (Object View)

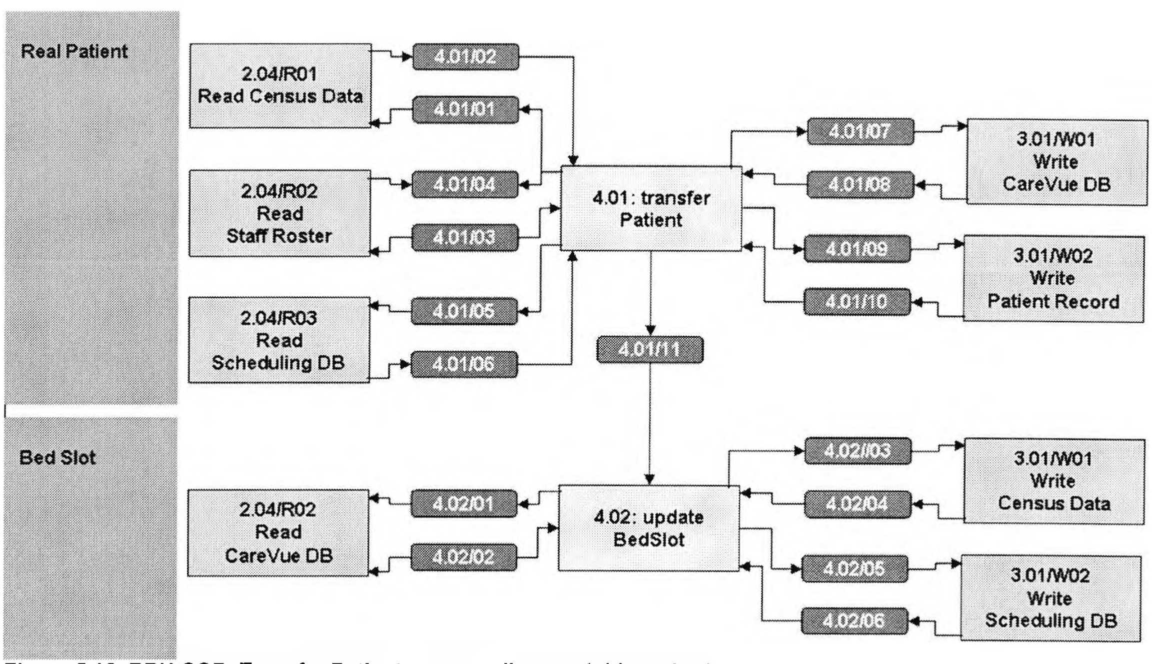

Figure 5.12. RBH COP: Transfer Patient process diagram (object view)

RBH COP: TRANSFER PATIENT (PROCESSOR VIEW)

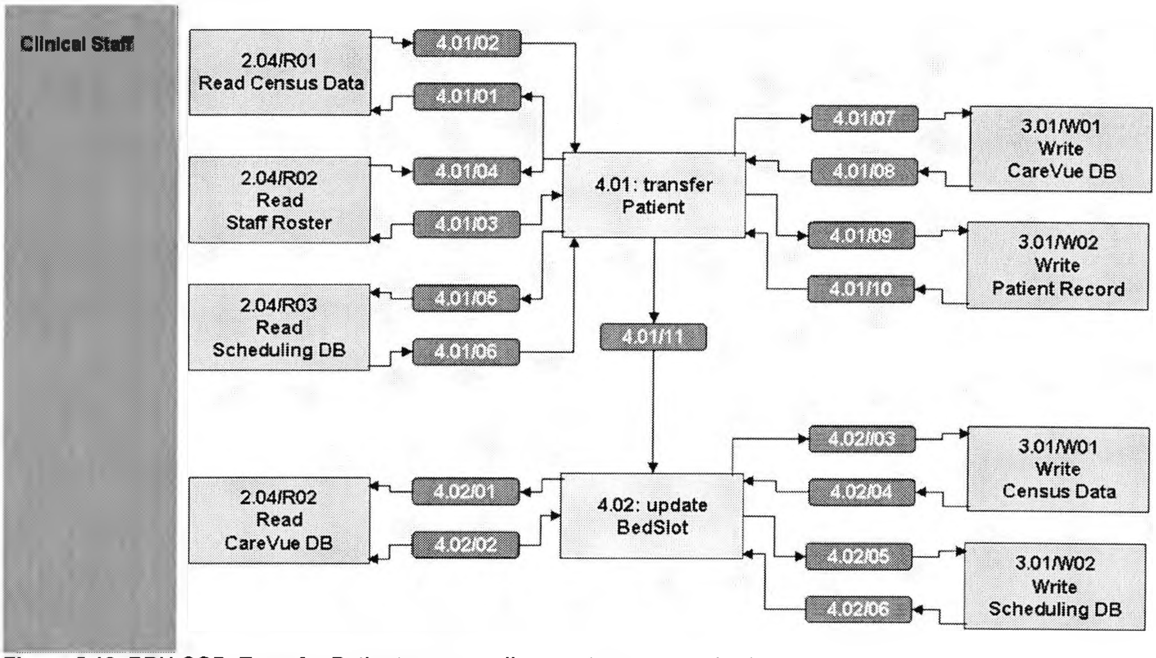

Figure 5.13. RBH COP: Transfer Patient process diagram (processor view)

5.2. RBH Proposed Operational Model (POP)

5.2.1. Introduction

The purpose of this model is to represent the object classes, databases, and component processes and data sets of the problem domain as it would exist following the deployment of the CAPSS information system. In conjunction with the current operational model developed in the preceding section, this model may then be used to not only present the problem domain post-deployment of CAPSS, but also to identify what needs to be changed in order to deploy CAPSS. Moreover, it also constitutes the preliminary requirements model for the specification of the functional requirements of CAPSS which may then be used as the basis for the design modelling phase of development.

5.2.2. Objects

BED SLOT

The Bed Slot object class contains all data and processes related to the individual bed slot within the high-dependency environment. This includes associations to the patient object to which the bed slot is allocated, and the unit object of which the bed slot object is a component.

The Bed Slot object class appears in both RBH COP and RBH POP models. It is shown below in standard object oriented representation.

POP Bed Slot	
POP Bed Slot ID	Text String
POP Actual Bed Slot Status Time = [T]	Category
POP Actual Labour Component [N] Status Time = [T]	Binary
POP Actual TISS Component [N] Status Time = [T]	Binary
POP Bed Slot Unit Name	Association
POP Patient Scheduled Monotonic Bed Slot Time = [T]	Association
POP Patient Scheduled Non-Monotonic Bed Slot Time = [T]	Association
POP Patient Type [P] Admissible	Binary
POP Projected Bed Slot Status Time = [T]	Category
POP Projected Labour Component [N] Status Time = [T]	Binary
POP Projected TISS Component [N] Status Time = [T]	Binary
POP2.04: schedulePreAdmissionPatient(Bed Slot)	
POP2.06: allocateUnAllocatedBedSlot(Bed Slot)	
POP3.02: updateAllocatedBedSlot(Bed Slot)	
POP4.02: deallocateAllocatedBedSlot(Bed Slot)	

Figure 5.14. POP Bed Slot object class diagram

PATIENT

The Patient object class contains all data and processes related to the individual patient. This includes both clinical, demographic and economic data.

The Patient object class appears in both RBH COP and RBH POP models. It is shown below in standard object oriented representation.

POP Patient	
POP Patient Hospital Number	Text String
POP Actual Admission Time Unit [U]	Date/Time
POP Actual Length of Stay Unit [U]	Number
POP Admitting Consultant	Text String
POP Admitting Surgeon	Text String
POP Discrepancy Admission Time Unit [U] Time = [T]	Number
POP Discrepancy Length of Stay Unit [U] Time = [T]	Number
POP Patient Admission Diagnosis	Category
POP Patient Clinical Attribute [N]	Text String
POP Patient Clinical Attribute [N] Time = T	Number
POP Patient Current Diagnosis	Category
POP Patient Current Location	Text String
POP Patient Date of Birth	Date/Time
POP Patient Demographic Attribute [N]	Text String
POP Patient Home Address	Text String
POP Patient Hospital Admission Date	Date/Time
POP Patient Name	Text String
POP Patient Projected Discharge Time	Date/Time
POP Patient Scheduling Status	Category
POP Projected Admission Time Unit [U] Time = [T]	Date/Time
POP Projected Length of Stay Unit [U] Time = [T]	Number
POP2.01: createPreAdmissionPatientRecord(Patient)	
POP2.02: updatePreAdmissionPatentRecord(Patient)	
POP2.03: updateWaitingList(Patient)	
POP2.05: admitPreAdmissionPatient(Patient)	
POP2.07: predictResourceConsumption(Patient)	
POP3.01: updateTreatmentPatient(Patient)	
POP4.01: transferTreatmentPatient(Patient)	

Figure 5.15. Patient object class diagram

UNIT

The Unit object class contains all data and processes related to the individual healthcare unit within the high-dependency environment. The Unit object class is primarily a data object class, containing many derived variables related to the bed slot objects which partly compose the unit object through association.

The Unit object class appears in both RBH COP and RBH POP models. It is shown below in standard object oriented representation.

POP Unit	
POP Unit Name	Text String
POP Accept Patient Type [P] From Unit [U]	Binary
POP Actual Labour Component [N] per Bed Slot Time = [T]	Number
POP Actual Number Allocated-Occupied Bed Slots Time = [T]	Number
POP Actual Number Allocated-Unoccupied Bed Slots Time = [T]	Number
POP Actual Number Unallocated-Unoccupied Bed Slots Time = [T]	Number
POP Actual Occupancy Rate Time = [T]	Number
POP Actual TISS Component [N] per Bed Slot Time = [T]	Number
POP Allocated-Occupied Bed Slots Discrepancy Time = [T]	Number
POP Allocated-Unoccupied Bed Slots Discrepancy Time = [T]	Number
POP Clinical Director Name	Text String
POP Discharge Patient Type [P] To Unit [U]	Binary
POP Labour Component [N] per Bed Slot Discrepancy Time = [T]	Number
POP Maximum Actual Number Bed Slots	Number
POP Maximum Actual Number Bed Slots Patient Type [P]	Number
POP Mean Actual Occupancy Rate Period = [P]	Number
POP Mean Occupancy Rate Discrepancy Period = [P]	Number
POP Mean Projected Occupancy Rate Period = [P]	Number
POP Occupancy Rate Discrepancy Time = [T]	Number
POP Operational Manager Name	Text String
POP Projected Labour Component [N] per Bed Slot Time = [T]	Number
POP Projected Number Allocated-Occupied Bed Slots Time = [T]	Number
POP Projected Number Allocated-Unoccupied Bed Slots Time = [T]	Number
POP Projected Number Unallocated-Unoccupied Bed Slots Time = [T]	Number
POP Projected Occupancy Rate Time = [T]	Number
POP Projected TISS Component [N] per Bed Slot Time = [T]	Number
POP Target Labour Component [N] per Bed Slot	Number
POP Target TISS Component [N] per Bed Slot	Number
POP TISS Component [N] per Bed Slot Discrepancy Time = [T]	Number
POP Unallocated-Unoccupied Bed Slots Discrepancy Time = [T]	Number
POP1.01: createPolicy(Unit)	
POP2.08: createScheduleEvaluation(Unit)	

Figure 5.16. Unit object class diagram

The object relationship diagram for the three object classes is as shown in Figure 5.17 below:

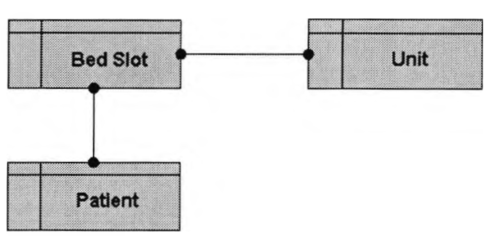

Figure 5.17. POP object class relationship diagram

5.2.3. Databases

CAREVUEDB

The CareVue Database is an integrated component of the CareVue medical information system. It records many fields of the patients medical condition, as well as many demographic and economic data fields. CareVueDB duplicates much of the data which is contained in PatientRecord. CareVueDB is a component database of both the current and operational models.

CAPSS DB

The CAPSS Database is an integrated component of the CAPSS computer system. It contains both clinical and economic data and includes all of the fields contained in CensusData, Policy Documents, SchedulingDB, Staff Roster and WaitingList databases of the current operational model. CAPSS DB is a component of only the proposed operational model.

PATIENTRECORD

The Patient Record is the enduring source of each patient's medical history. It is purely paper based and duplicates much of the data which is also contained in other clinical and demographic databases. PatientRecord is a component database of both the current and operational models.

The relationship between the data which defines the objects within each of the object classes and the databases where that data is recorded and stored, is as shown in Figure 5.18 below:

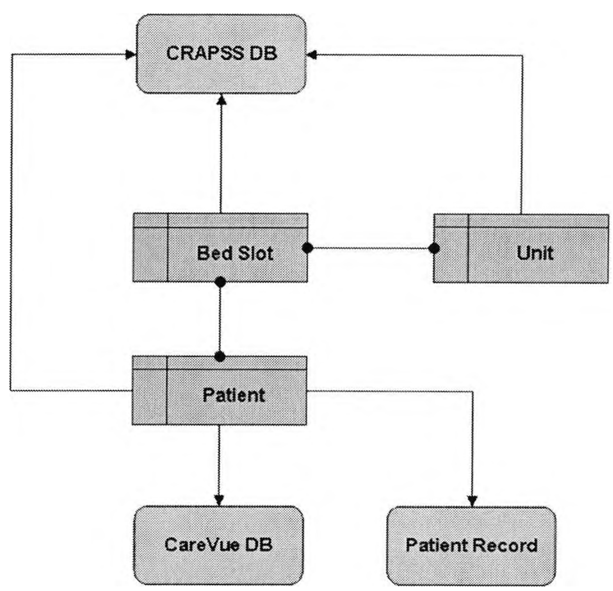

Figure 5.18. POP extended object class relationship diagram

5.2.4. Processes

POP1.01: CREATEPOLICY(UNIT)

POP1.01: createPolicy(Unit) is a component process of the Create Policy process group. It creates a new set of policy specifications or updates existing policy specifications determining parameters such as which patients may be admitted to the unit, the resource profiles of the unit, and so on. It is a component process of the Unit object class and is performed by the processor StrategicManager. It is composed of the following database processes:

Database Process	DB Process Type	Database Name
POP1.01/R01	Manual Read	CAPSS DB
POP1.01/W01	Manual Write	CAPSS DB

Table 5.12. POP1.01: createPolicy(Unit) database processes

POP2.01: CREATEPATIENTRECORD(PATIENT)

POP2.01: createPatientRecord(Patient) is a component process of the Schedule Patient process group. It creates the hospital patient record when the patient is first admitted to the hospital environment. It is a component process of the Patient object class and is performed by the processor ClinicalStaff. It is composed of the following database processes:

Database Process	DB Process Type	Database Name
POP2.01/R01	Manual Write	PatientRecord

Table 5.13. POP2.01: createPatientRecord(Patient) database processes

POP2.02: UPDATEPATENTRECORD(PATIENT)

POP2.02: updatePatentRecord(Patient) is a component process of the Schedule Patient process group. It updates the hospital patient record which was created when the patient was first admitted to the hospital environment. It is a component process of the Patient object class and is performed by the processor ClinicalStaff. It is composed of the following database processes:

Database Process	DB Process Type	Database Name
POP2.02/R01	Manual Read	PatientRecord
POP2.02/W01	Manual Write	PatientRecord

Table 5.14. POP2.01: updatePatientRecord(Patient) database processes

POP2.03: UPDATEWAITINGLIST(PATIENT)

POP2.03: updateWaitingList(Patient) is a component process of the Schedule Patient process group. It updates the waiting list by entering the patient onto the waiting list or updating the entry of an existing patient. It is a component process of the Patient object class and is performed by the processor ClinicalStaff. It is composed of the following database processes:

Database Process	DB Process Type	Database Name
POP2.03/R01	Manual Read	PatientRecord
POP2.03/W01	Manual Write	CAPSS DB

Table 5.15. POP2.03: updateWaitingList(Patient) database processes

POP2.04: SCHEDULEPATIENT(BED SLOT)

POP2.04: schedulePatient(Bed Slot) is a component process of the Schedule Patient process group. It schedules a patient for admission to the high dependency environment by non-monotonically allocating an unallocated bed slot to a patient. It is a component process of the Bed Slot object class and is performed by the processor OperationalManager. It is composed of the following database processes:

Database Process	DB Process Type	Database Name
POP2.04/R01	Manual Read	CAPSS DB
POP2.04/R02	Manual Read	PatientRecord
POP2.04/W01	Manual Write	CAPSS DB

Table 5.16. POP2.04: schedulePatient(Bed Slot) database processes

POP2.05: ADMITPATIENT(PATIENT)

POP2.05: admitPatient(Patient) is a component process of the Schedule Patient process group. It enters the patients details on the CareVue medical information system once the patient is admitted to the high dependency environment. It is a component process of the Patient object class and is performed by the processor ClinicalStaff. It is composed of the following database processes:

Database Process	DB Process Type	Database Name
POP2.05/R01	Manual Read	PatientRecord
POP2.05/W01	Manual Write	CareVueDB

Table 5.17. POP2.05: admitPatient(Patient) database processes

POP2.06: ALLOCATEBEDSLOT(BED SLOT)

POP2.06: allocateBedSlot(Bed Slot) is a component process of the Schedule Patient process group. It updates the bed slot object by changing the allocation status of the bed slot to being monotonically allocated to a patient. It is a component process of the Bed Slot object class and is performed by the processor ClinicalStaff. It is composed of the following database processes:

Database Process	DB Process Type	Database Name
POP2.06/R01	Automatic Read	CareVueDB
POP2.06/W01	Automatic Write	CAPSS DB

Table 5.18. POP2.06: allocateBedSlot(Bed Slot) database processes

POP2.07: PREDICTRESOURCECONSUMPTION(PATIENT)

POP2.07: predictResourceConsumption(Patient) is a component process of the Schedule Patient process group. It makes projections of the resource requirements of the patient prior to admission to the high-dependency environment, including to which units the patient will require admission, and when. It is a component process of the Patient object class and is performed by the processor CAPSS. It is composed of the following database processes:

Database Process	DB Process Type	Database Name
POP2.07/R01	Automatic Read	CAPSS DB
POP2.07/R02	Manual Read	PatientRecord
POP2.07/W01	Automatic Write	CAPSS DB

Table 5.19. POP2.07: predictResourceConsumption(Patient) database processes

POP2.08: CREATESCHEDULEEVALUATION(UNIT)

POP2.08: createScheduleEvaluation(Unit) is a component process of the Schedule Patient process group. It evaluates the proposed admission schedule by generating economic and clinical performance statistics of each unit that are projected to result from admitting the patients according to the proposed schedule. It is a component process of the Unit object class and is performed by the processor CAPSS. It is composed of the following database processes:

Database Process	DB Process Type	Database Name
POP2.08/R01	Automatic Read	CareVueDB
POP2.08/R02	Automatic Read	CAPSS DB
POP2.08/W01	Automatic Write	CAPSS DB

Table 5.20. POP2.08: createScheduleEvaluation(Unit) database processes

POP3.01: UPDATEPATIENT(PATIENT)

POP3.01: updatePatient(Patient) is a component process of the Treat Patient process group. It continually updates both the paper-based and CareVue medical records of the patient as the treatment process continues. It is a component process of the Patient object class and is performed by the processor ClinicalStaff. It is composed of the following database processes:

Database Process	DB Process Type	Database Name
POP3.01/R01	Manual Read	PatientRecord
POP3.01/R02	Manual Read	CareVueDB
POP3.01/W01	Manual Write	PatientRecord
POP3.01/W02	Manual Write	CareVueDB
POP3.01/W03	Automatic Write	CareVueDB

Table 5.21. POP3.01: updatePatient(Patient) database processes

POP3.02: UPDATEBEDSLOT(BED SLOT)

POP3.02: updateBedSlot(Bed Slot) is a component process of the Treat Patient process group. It continually updates the bed slot status according to the projected duration of the allocation or recording any change in status. It is a component process of the Bed Slot object class and is performed by the processor ClinicalStaff. It is composed of the following database processes:

Database Process	DB Process Type	Database Name
POP3.02/R01	Automatic Read	CareVueDB
POP3.02/W01	Automatic Write	CAPSS DB

Table 5.22. POP3.02: updateBedSlot(Bed Slot) database processes

POP4.01: TRANSFERPATIENT(PATIENT)

POP4.01: transferPatient(Patient) is a component process of the Transfer Patient process group. It updates both paper-based and CareVue medical records to reflect the transfer of the patient from one unit to another. It is a component process of the Patient object class and is performed by the processor ClinicalStaff. It is composed of the following database processes:

Database Process	DB Process Type	Database Name
POP4.01/R01	Manual Read	CAPSS DB
POP4.01/W01	Manual Write	CareVueDB
POP4.01/W02	Manual Write	PatientRecord

Table 5.23. POP4.01: transferPatient(Patient) database processes

POP4.02: DEALLOCATEBEDSLOT(BED SLOT)

POP4.02: deallocateBedSlot(Bed Slot) is a component process of the Transfer Patient process group. It updates the status of the bed slot to reflect the transfer of the patient from the unit, thus de-allocating the bed slot to the patient, and possibly allocating it to another patient. It is a component process of the Bed Slot object class and is performed by the processor ClinicalStaff. It is composed of the following database processes:

Database Process	DB Process Type	Database Name
POP4.02/R01	Automatic Read	CareVueDB
POP4.02/W01	Automatic Write	CAPSS DB

Table 5.24. POP4.02: deallocateBedSlot(Bed Slot) database processes

5.2.5. Process Diagrams

RBH POP: CREATE POLICY (OBJECT VIEW)

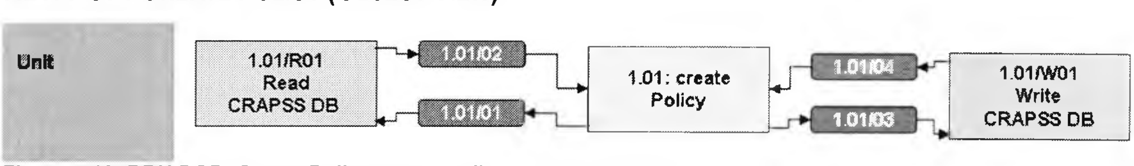

Figure 5.19. RBH POP: Create Policy process diagram (object view)

RBH POP: Create Policy (Processor View)

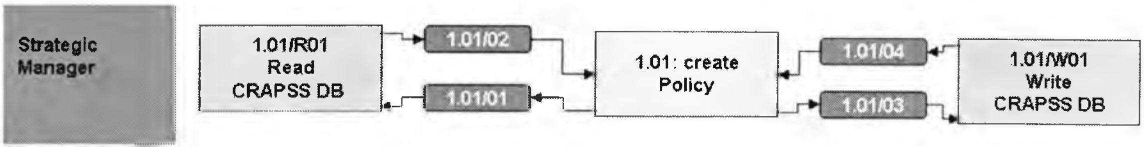

Figure 5.20. RBH POP: Create Policy process diagram (process view)

RBH POP: Schedule Patient (Object View)

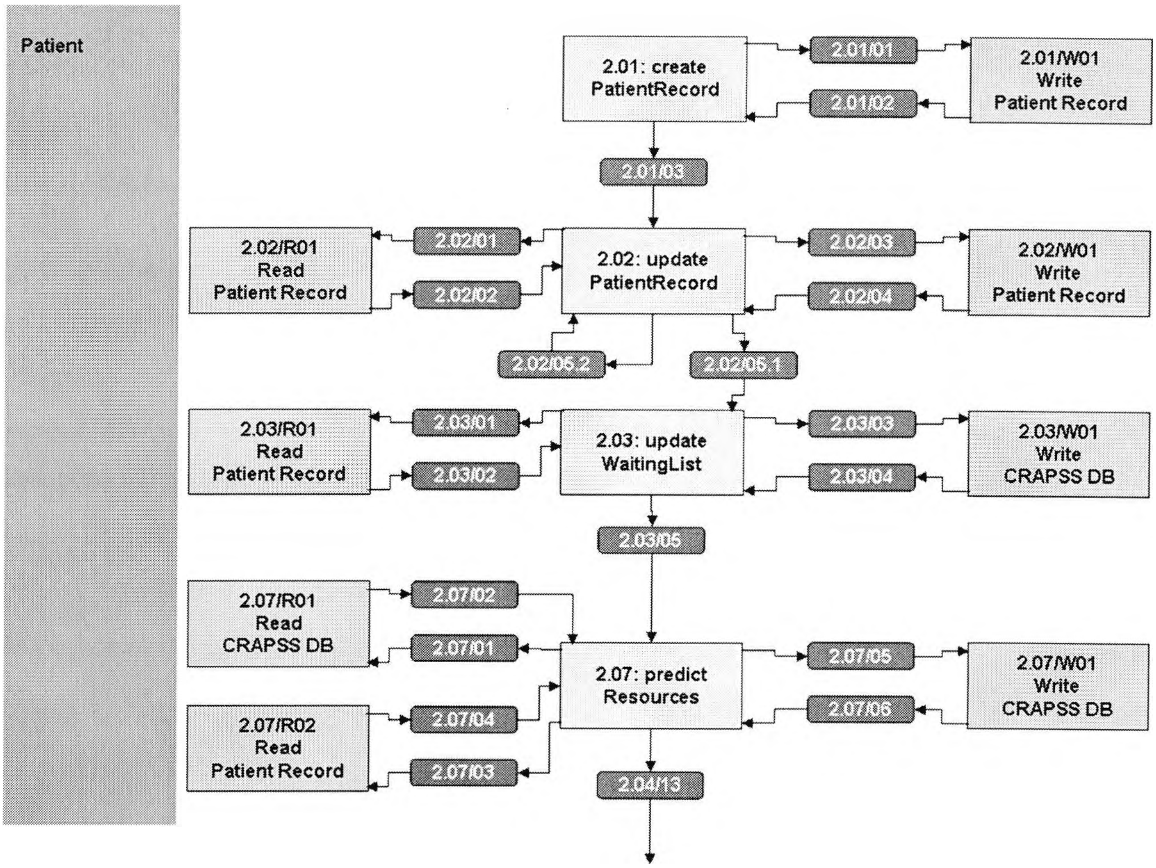

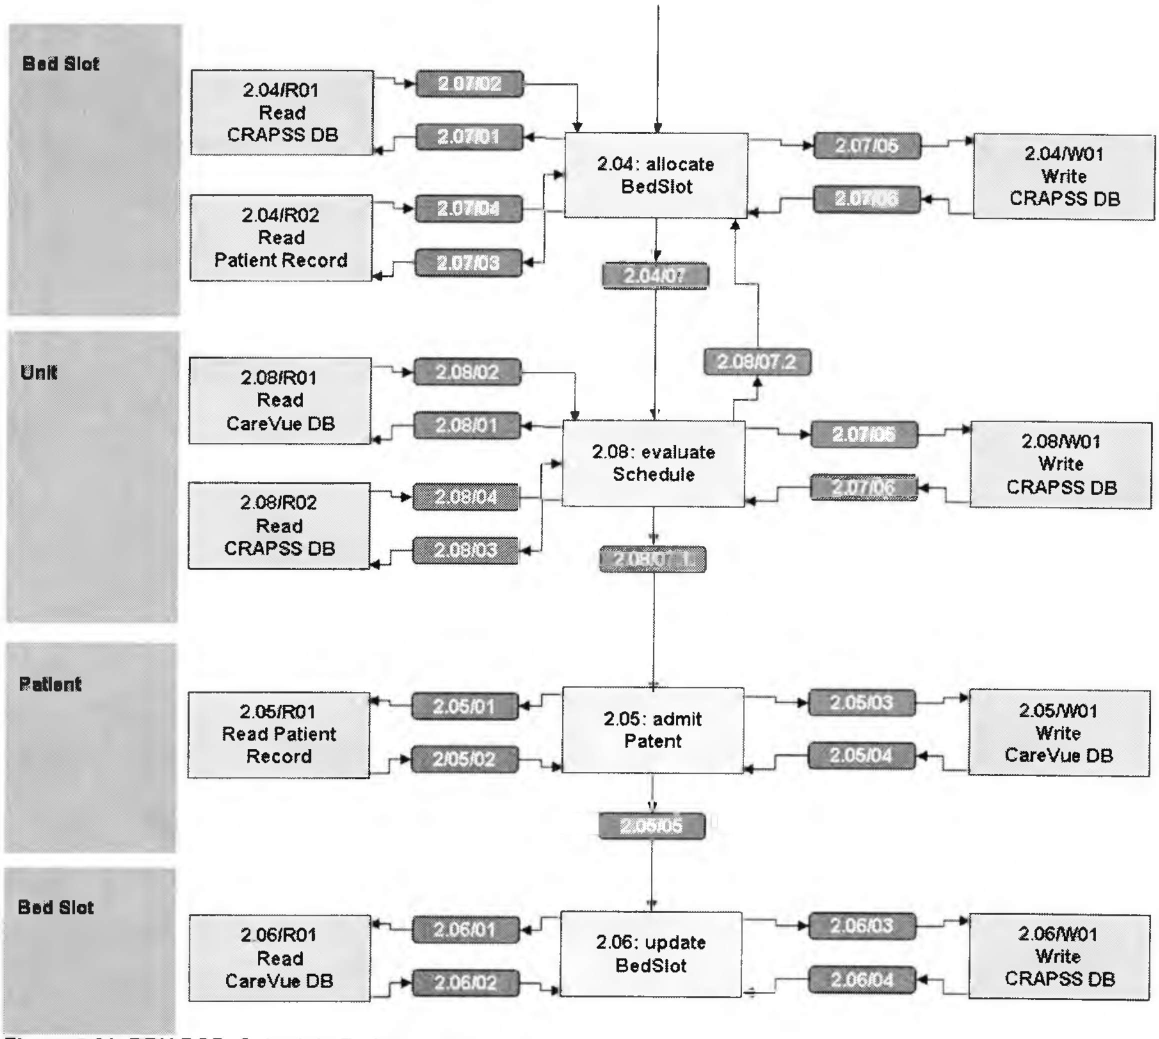

Figure 5.21. RBH POP: Schedule Patient process diagram (object view)

RBH POP: Schedule Patient (Processor View)

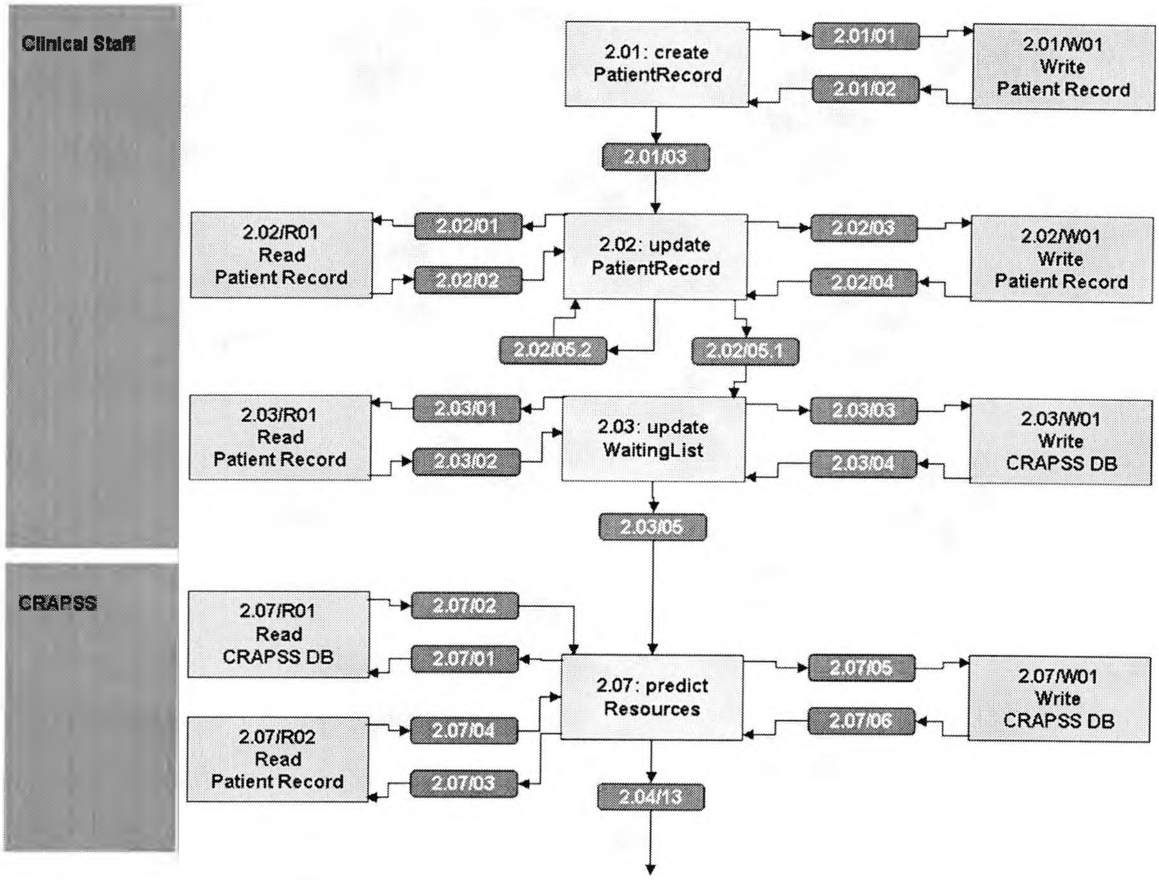

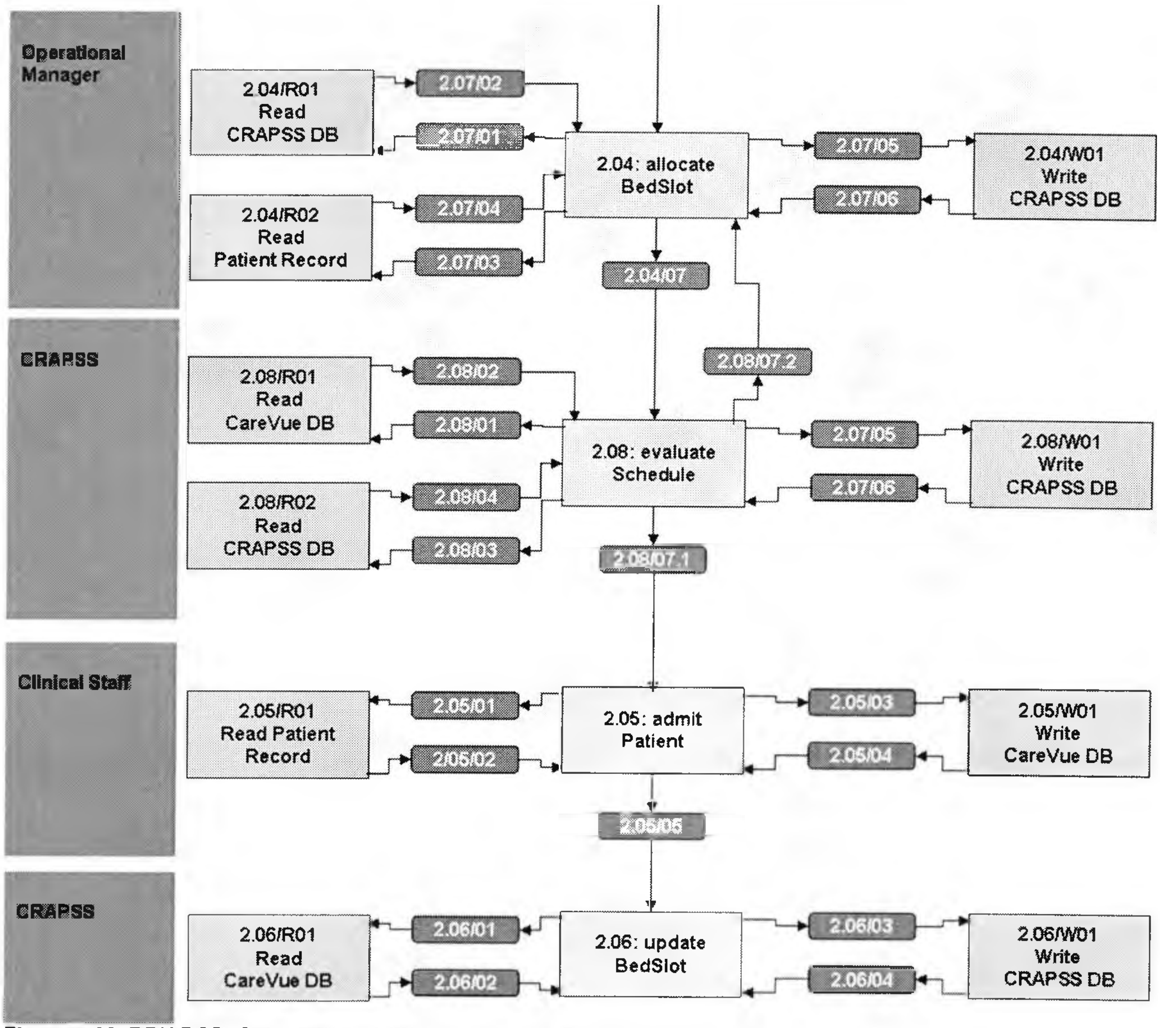

Figure 5.22. RBH POP: Schedule Patient process diagram (processor view)

RBH POP: TREAT PATIENT (OBJECT VIEW)

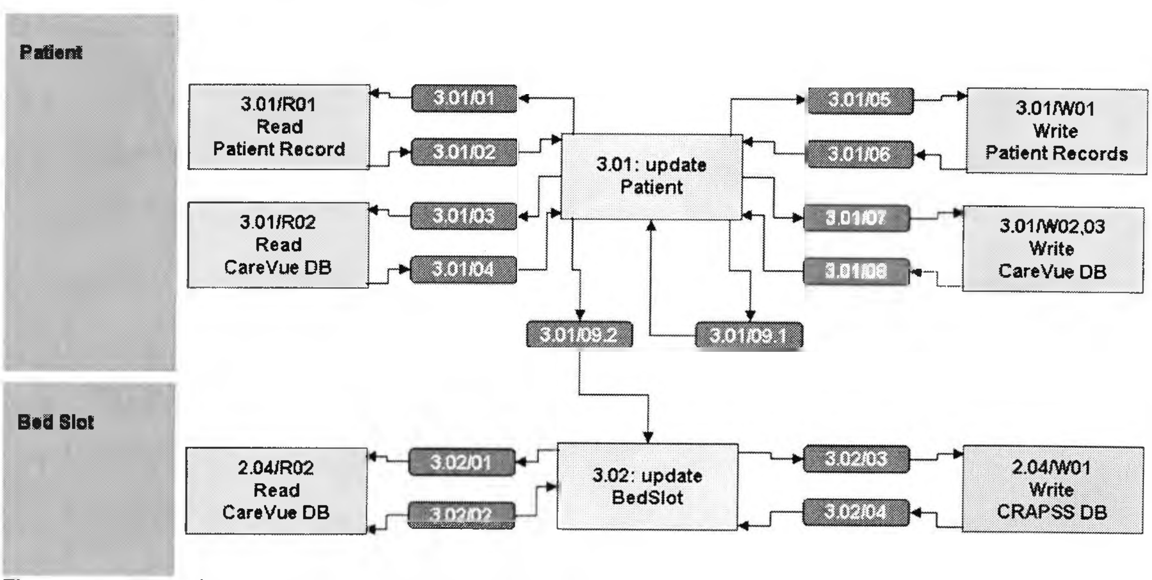

Figure 5.23. RBH POP: Treat Patient process diagram (object view)

RBH POP: TREAT PATIENT (PROCESSOR VIEW)

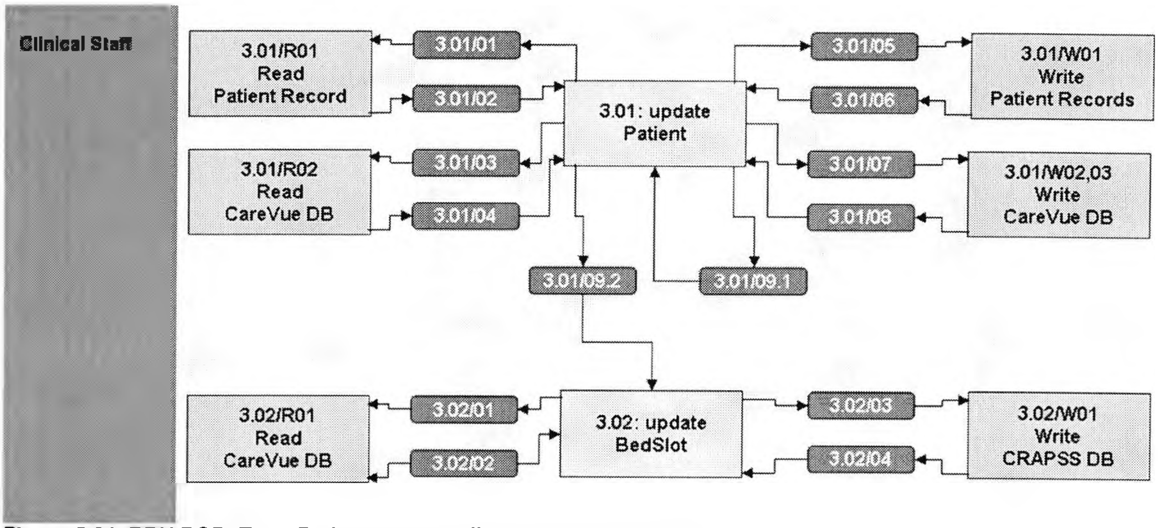

Figure 5.24. RBH POP: Treat Patient process diagram (processor view)

RBH POP: TRANSFER PATIENT (OBJECT VIEW)

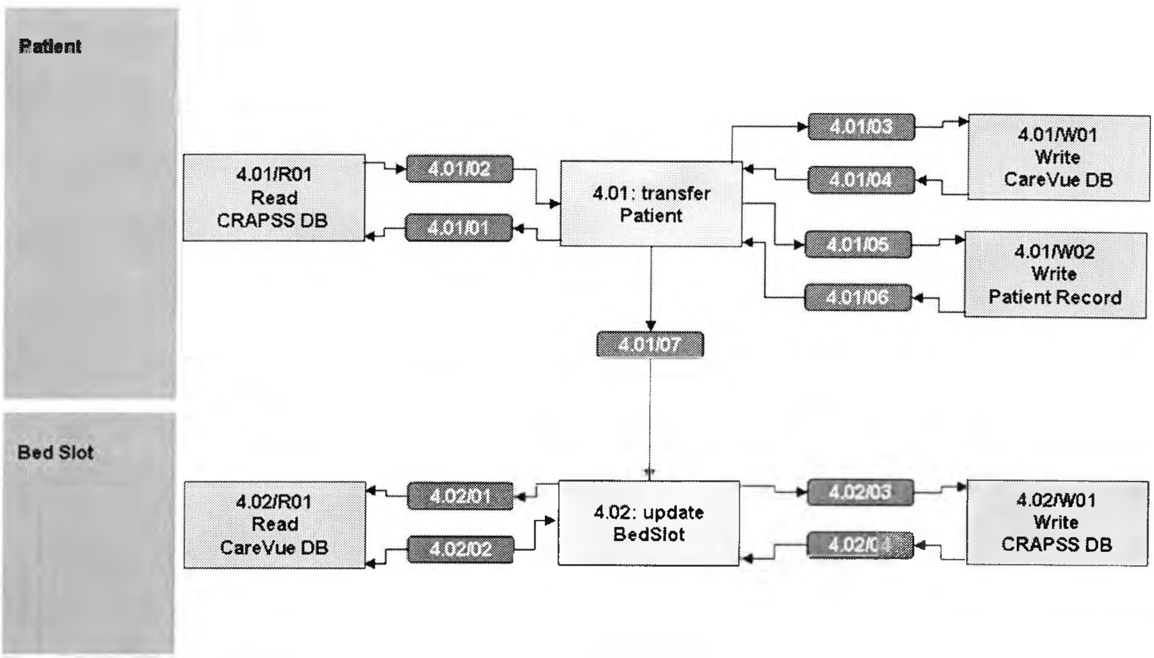

Figure 5.25. RBH POP: Transfer Patient process diagram (object view)

RBH POP: TRANSFER PATIENT (PROCESSOR VIEW)

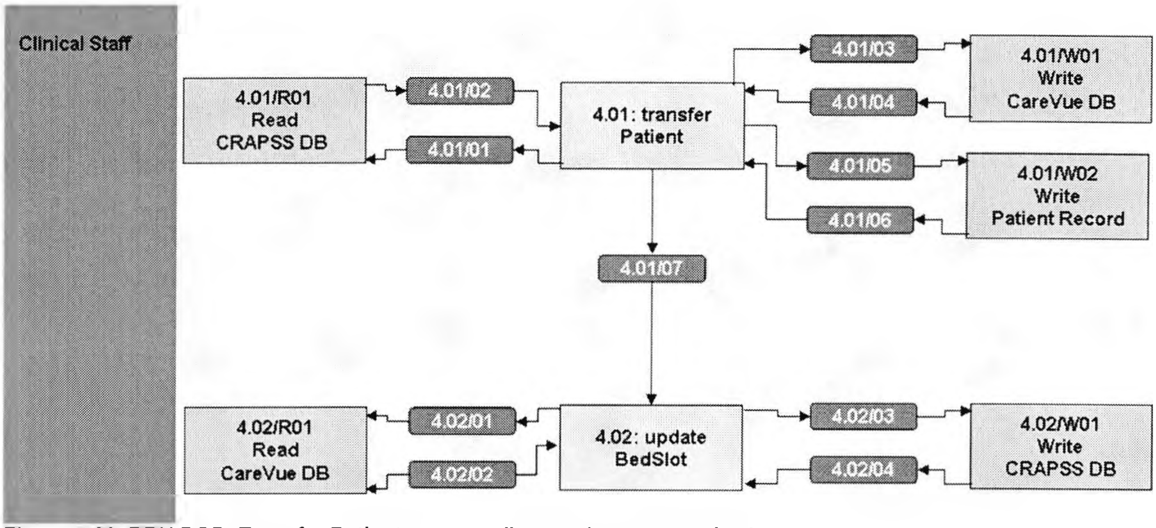

Figure 5.26. RBH POP: Transfer Patient process diagram (processor view)

5.3. Discussion

5.3.1. Model Comparison

In this chapter, two operational models have been presented; the first represents the current operational system at the Royal Brompton and Harefield NHS Trust's High-Dependency Environment (HDE); the second the proposed operational system at HDE that incorporates a Computer-assisted patient scheduling system (CAPSS) aimed at increasing the cost-effective operation of HDE by enhancing the level of control over patient scheduling. Both the current and proposed operational models were developed and presented using the Petri-net based object-oriented modelling approach and formalism presented in the previous chapter.

The main points of each model are as follows:

1. The problem domain was assumed to comprise three distinct object classes in both the current operational model (COP) and the proposed operational model (POP). These three classes represented the healthcare unit, the individual bed slot and the individual patient.
2. The operational processes of patient scheduling were grouped into four distinct process groups in both COP and POP. These three process groups were Policy Development, Patient Scheduling, Patient Treatment, and Patient Transfer.
3. In COP various databases were identified. These databases were interpreted very loosely as databases, often being nothing more than a loose collection of computer files and paper documents. In POP many of these databases were eliminated, being replaced by a single dedicated database called CAPSS DB.

4. In COP it was assumed that projected values of clinical and economic attributes, such as projected length of stay in a particular unit, were not included in the object classes due to these values not being recorded in any recognisable database.

The main points in comparing the two models are as follows:

1. In the current operational model (COP), there were many different databases, many of which were informal and paper-based with the data comprising the databases often distributed both geographically and in terms of being deployed in multiple media. The problem represented by this multiplicity of databases was compounded by a lack of integration between the databases, with no automatic read or write capabilities between the databases, and by the data defining objects being derived from one or more different data resources.

   The problems represented by the multiplicity of databases were largely resolved in the proposed operational model (POP) by replacing many of the non-computerised databases by a single database labelled 'CAPSS' in the model, which also reduced the amount of object data which was distributed amongst multiple data resources.

2. In COP there were no automated or systematic means for implementing effective control over resource allocation or patient scheduling. Those processes which were present in the problem domain were typically human based decision-making processes which lacked any standardisation or organisational learning capacity. Moreover, all such decision-making processes were not integrated into any data resource, thus dramatically reducing the capability of communication of control information between different processing resources.

   These problems were resolved in POP by introducing two new processes in the patient scheduling group of operational processes, both of which were fully integrated into the CAPSS computerised database, designed specifically for recording data necessary for increasing the cost-effectiveness of healthcare delivery through enhancing the level of control over patient scheduling. The first of these two new processes made computerised projections of patients' resourcing requirements, such as their projected length of stay in each component healthcare unit of HDE. These projections could then be used to inform the patient scheduling process performed by operational management and clinical staff. The resulting patient admission schedule and healthcare resource schedule could then be evaluated by the second of the two new processes introduced in the model by generating various performance statistics for each unit within the HDE.

3. In COP, because there were no systems in place for integrating the decision-making processes involved in patient scheduling into data resources, the capacity to learn from that

data is dramatically reduced, being left to the whims and frailties of human cognition and memory.

The foundations for resolving this problem were established in POP by integrating the decision making processes involved in patient scheduling into the CAPSS database. The bed slot and unit object classes data structures were also modified in POP by including error-measuring attributes that measure the discrepancies between the projected and actual projections made regarding both patient object class attributes, such as projected length of stay in each healthcare unit within the HDE, as well as for unit class attributes such as projected occupancy rates.

The two models may also be compared statistically by measuring and comparing the following variables from each model:

1. DB Process Type Profile, and
2. Object Class Fragmentation Rate

where each of these variables are defined in the preceding chapter.

The DB Process Type Profile for each model is shown in Table 5.25 below:

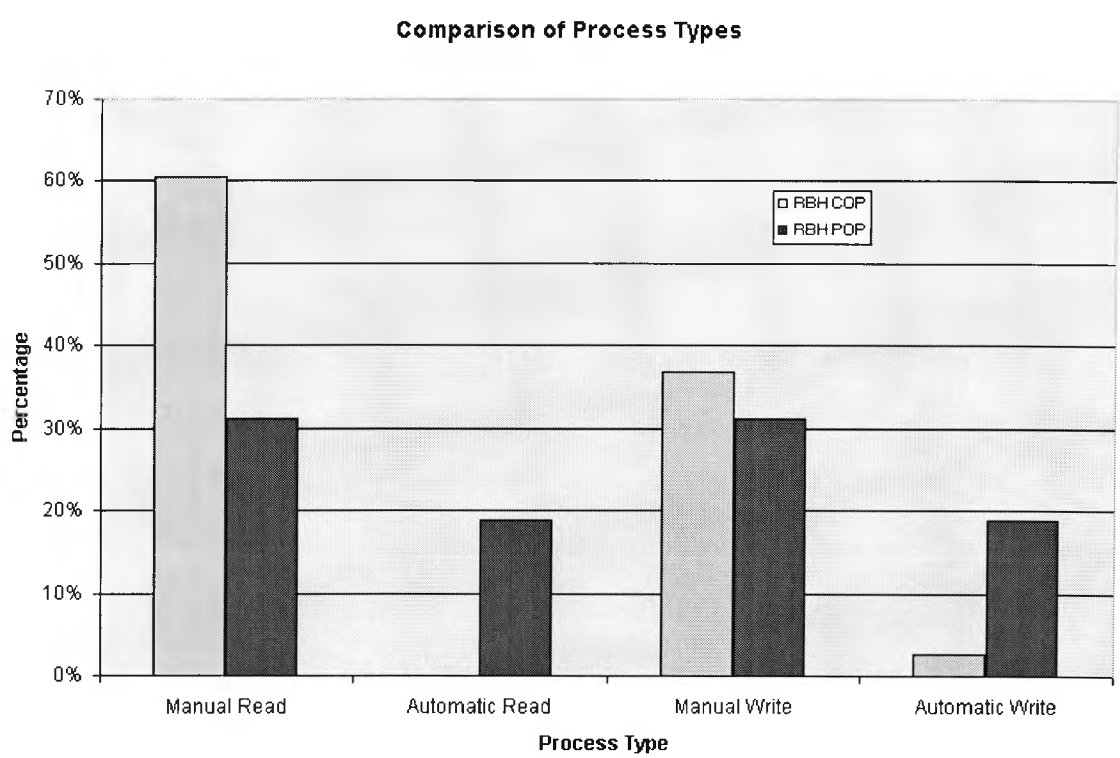

Table 5.25. Analysis of database process types between RBH COP and RBH POP models

Process Type	RBH COP	RBH POP

Manual Read	23	10
Automatic Read	0	6
Manual Write	14	10
Automatic Write	1	6
	38	**32**

Table 5.26. Numbers of each database process type between RBH COP and RBH POP models

From Table 5.26 above, it can be seen that RBH POP has 6 fewer database processes than does RBH COP. More importantly, however, whereas 37 out of 38 of RBH COP database processes (97%) are manual – i.e. either Manual Read or Manual Write process types – the corresponding figure for RBH POP is 20 out of 32 (62%), with all of the other database processes being automated. This demonstrates that, ceteris paribus, RBH POP is more cost-effective as a system of patient scheduling than is RBH COP.

With regards to the Object Class Fragmentation Rate, the two models are compared in Table 5.27 below:

Model	Objects	Relations	Databases	Ratio
RBH COP	3	7	7	0.06
RBH POP	3	5	3	0.20

Table 5.27. Object Class Fragmentation Rate between RBH COP and RBH POP models

In Table 5.27 above, it can be seen that, while the two models have the same number of object classes in each case (3), RBH COP has more relations between the data attributes which comprise those object classes' data structures and the databases in which the values of those attributes are recorded and updated (7 versus 5). Moreover, RBH COP has more than twice the number of databases than does RBH POP (7 versus 3). All of this together results in a very low Object Class Fragmentation Rate of 0.06 for RBH COP, against a figure of 0.20 for RBH POP.

It should be noted that the figure of 0.20 for RBH POP is still far from the optimal level of unity, however. The reason for this can be seen as originating primarily from a lack of database integration, rather than a distribution of object class data amongst a multiplicity of databases. Thus, for example, if all of the 3 databases in RBH POP were integrated into a single databases, the resulting Object Class Fragmentation Rate would increase to 0.6.

5.3.2. The Design Models

All of the discussion thus far has been hypothetical. A proposed model of patient scheduling has been presented as an improvement on the existing model of patient scheduling, primarily through the introduction of a new software subsystem called CAPSS. The argument for the proposed model being better than the existing model was based on abstract graph-theoretic measures which take no consideration of technical or empirical restrictions on the proposed model being actually possible, and even if it were to be possible, whether or not it would fulfil the expectations made of it.

6. Design Models

In the previous chapter a proposed operational model of the Royal Brompton and Harefield High-Dependency Environment (RBH HDE) was presented (RBH POP). RBH POP was proposed as combining two functions within one model. First, to demonstrate the impact on organisational structure and processing of implementing a computer-assisted patient scheduling system (CAPSS); second to indicate the preliminary data and functional requirements of the software component of such a system.

The main feature of RBH POP is the addition of two new organisational processes involved in scheduling a patient for admission to RBH HDE. The first of these, labelled '2.07: predictResources' and '2.08: evaluateSchedule'.

With the specifications of requirements as a black-box modelling phase, as described in Chapter 4, neither the internal workings of 2.07: predictResources or 2.08: evaluateSchedule were defined in RBH POP. Therefore, as a proof of concept that the computerisation of the patient scheduling and resource allocation process may increase the operational cost-effectiveness of RBH HDE and similar healthcare systems, it is necessary to propose white-box models of both of these processes. The development of these white-box models will be the objective of this chapter.

The Chapter is divided into three sections. The first and section sections will present the models of 2.07: predictResources and 2.08: evaluateSchedule respectively, with the final section being a discussion of the foregoing models and how they relate to RBH POP.

6.1. 2.07: predictResources

The objective of the process 2.07: predictResources, as shown in RBH POP, is to take as input various clinical and demographic characteristics of the patient population, and output various outcome predictions that reflect the projected consumption of healthcare resources of RBH HDE.

The output variables identified in RBH POP are labelled in RBH POP as the Patient object class attribute 'POP Projected Length of Stay Unit [U] Time = [T]'. That is, the patient's projected length of stay in unit U at time T. There are three things to note from how this attribute is defined in RBH POP. First, a different instance of the attribute is needed for each unit. Thus, for example, there will be a projection of the patient's length of stay in the operating theatre, another projection of the length of stay in the intensive care unit, and so on. Second, these projections are not defined as monotonic in RBH POP, so that their values may change over time. Thus, for each projection for each unit, there will be an additional time factor, so that, for example, the projection of length of stay in intensive care made at time T1 need not be the same as the projection made at time T2, and so on. Finally, RBH POP describes the attribute POP Projected Length of Stay Unit [U] Time = [T] as being a derived attribute, where the derivation is endogenous – calculated from the values of other attributes within the Patient object class. In other words, therefore, POP Projected Length of Stay Unit [U] Time = [T] is proposing a length of stay prediction model, where the predictions of length of stay are made i) for each unit within the RBH HDE, ii) on an ongoing basis throughout the patient's stay within the RBH HDE, and iii) on the basis of the patient's clinical and demographic characteristics.

The design of a computerised prediction model (actually, a suite of prediction models), to satisfy the above requirements would normally be of one of two different architectures. Either the prediction model could be a relatively simple algorithm based on statistical regression models, where, for example, the equation of the regression line is used to define the algorithm, with each predictive clinical and demographic characteristic of the patient being represented by one of the co-efficients of the regression equation. Or, the prediction model could use connectionist techniques such as artificial neural networks (ANNs), where the input nodes of an ANN would represent the values of different clinical or demographic variables, and the output node would represent the predicted length of stay.

Both of these types of prediction model have been developed in various contexts in the literature on clinical prediction models (although the vast majority have been of the former – statistical regression – kind). The actual computerisation of these kinds of model is relatively straightforward, with the main effort being directed towards the identification of those clinical and demographic variables to be included in the model and, in the case of the statistical models, their weighting in any algorithm. For this reason, this section will concentrate on a literature review of studies developing such algorithms and the identification of those variables to be included in them.

6.1.1. Literature Review

In the literature review the Medline database of clinical research journal papers was searched with the search string "Length of Stay Prediction". A total of 10 journal articles satisfied the criteria of i) having the objective of developing a length of stay prediction model, ii) being published after 1985, and iii) measuring outcome within a hospital environment. The 10 articles are listed below, with each given a code for ease of reference:

[BAR96a] Barie, P S et al, *Utility of illness severity scoring for prediction of prolonged surgical critical care*. J. Trauma, 1996, 40(4) pp.513-519.

[BEC95] Becker, RB et al, *The use of APACHE III to evaluate ICU length of stay, resource use, and mortality after coronary artery by-pass surgery.* J Cardiovasc Surg, 1995, 36(1) pp.35735.

[BUC94] Buchman, T G, et al, *A comparison of statitistical and connectionist models for the prediction of chronicity in a surgical intensive care unit.* Crit. Care Med., 1994, 22(5) pp.750-762.

[KAT88] Katz et al, *Predictors of Length of Hospitalization after Cardiac Surgery*, Ann. Thorac. Surg., 1988 (45)

[KNA93] Knaus, W A, et al, *Variations in mortality and length of stay in intensive care units*. Annals Internal Med., 1993, 118(10) pp.753-761.

[MAR95] Marshall, J C, et al, *The multiple organ dysfunction (MOD) score: A reliable indicator of a complex clinical outcome*. Crit. Care Med., 1995, 23(10) pp.1638-1652.

[MOU95] Mounsey, J P, et al, *Determinants of the length of stay in intensive care and in hospital after coronary artery surgery.* Br. Heart J., 1995, 73 pp.92-98.

[TUM92] Tuman, K J, et al, *Morbidity and Duration of ICU Stay after Cardiac Surgery*, Chest, 1992, 102 pp.36-44.

[TUJ92] Tu, J V, et al, *Use of a neural network as a predictive instrument for length of stay in the intensive care unit following cardiac surgery*. Proc. Ann. Symp. Computer Apps in Med. Care, 1993, N/A pp.666-672.

[TUJ94] Tu, J V, et al, *A predictive index for length of stay in the intensive care unit following cardiac surgery.* Can. Med. Assoc. J., 1994, 151(2) pp.177-185.

Each of the above papers was classified into the outcome measures which were predicted by the model or models that were developed. These different outcome measures were categorised as follows:

- **Pre-operative Length of Stay**. This outcome measure predicts the length of stay in a post-operative healthcare unit before the patient is admitted to surgery.

- **Post-operative Length of Stay**. This outcome measure predicts the length of stay in a post-operative healthcare unit after the patient has been discharged from the operating theatre.
- **Mortality**. The outcome measure predicts the likelihood of whether or not the patient will die within either the high-dependency environment or the wider hospital setting as a consequence of either the treatment or the disease process.
- **Morbidity**. This outcome measure predicts the extent of morbidity that the patient suffers as a consequence of either the treatment or the disease process at a fixed period of time after discharge from the hospital setting.
- **ICU Length of Stay**. This outcome measure predicts the length of stay specifically within an intensive care unit.
- **Hospital Length of Stay**. This outcome measure predicts the length of stay within the hospital setting, including the length of stay within in high-dependency environment within the hospital setting.

The above categories are not mutually exclusive. For example, one prediction model can predict both post-operative length of stay and ICU length of stay. Table 6.01 below categorises each of the above 10 papers according to the outcome measures they model. In the case of [MOU95] there were two models developed, labelled as [MOU95]#1 and [MOU95]#2 in the table.

Reference	Pre-op LOS Prediction	Post-op LOS Prediction	Mortality Prediction	Morbidity Prediction	ICU LOS Prediction	Hosp. LOS Prediction
[MOU95]#2	No	Yes	No	No	No	Yes
[BAR96a]	No	Yes	Yes	No	Yes	No
[BEC95]	No	Yes	Yes	No	Yes	No
[KNA93]	No	Yes	Yes	No	Yes	No
[MAR95]	No	Yes	Yes	Yes	Yes	No
[KAT88]	Yes	No	No	No	No	Yes
[MOU95]#1	Yes	No	No	No	Yes	No
[TUJ93]	Yes	No	No	No	Yes	No
[TUJ94]	Yes	No	Yes	No	Yes	No
[TUM92]	Yes	No	No	Yes	Yes	No

Table 6.01. Categorisation of prediction models by outcome measure.

Apart from the outcome measures which are predicted by the models developed in the above studies, there are other important considerations in comparing the models. Many of these considerations relate to evaluating the methodological rigour of the studies, others relate to their applicability to the particular situation represented by RBH HDE in terms of its case mix. Both of these types of consideration are as follows:

Sample Size

This is an important consideration in any study. If the sample size is small, then any results cannot be validated through tests for statistical significance. The normal minimum sample size for prediction models is around 1000, although of course this is only a guideline as whether or not statistically significant results can be generated is dependent on the strength of the effect, as well as the sample size.

Mortalities

How the study deals with mortalities is especially important in the development of prediction models of length of stay. Those patients who are most likely to die within the high-dependency environment tend also to be those patients whose lengths of stay are the least predictable, as well as demonstrating greater variance than their surviving counterparts. Therefore, to exclude mortalities post hoc from the derivation of any length of stay prediction model tends to increase the accuracy of the prediction model. But in excluding mortalities from the sample, the prediction model thus developed is rendered useless by not being able to predict the length of stay of those patients whose outcome is of most concern, quite apart from the fact that a further mortality prediction model would be required to make the initial classification into those patients that would die and those that would survive.

Cardiac Status

The prediction of cardiac surgery outcomes is notoriously difficult [TUJ96], and has warranted specific outcome scores for cardiac patients. Apart from this being a consideration in comparing length of stay prediction studies, it is particularly important as it relates to RBH HDE due to its cardiothoracic case mix.

Surgical Status

Those patients who undergo surgery as part of their treatment tend – in the case of cardiothoracic high-dependency medicine at least – to have more predictable outcomes (both mortality, morbidity and length of stay measures) than those patients who do not undergo surgery. For this reason, whether or not a prediction model includes non-surgical patients in its study population has an effect on the accuracy of the predictions.

PROSPECTIVE STATUS

In deriving prediction models a study population is used. In validating the predictions of the model, however, it needs to be applied to a population different to the study population. This can either be done by dividing the population into a so-called test population and a training population in the case of connectionist models, where the model is derived using the training population and subsequently validated or otherwise using the test population. Or, in the case of statistical regression models it can be done by testing the model prospectively on patients as they are admitted to the healthcare system.

Table 6.02 below summarises all of the above 10 studies according to the above list of considerations. In addition to the above considerations, the list of operative procedures included in each study is given. The study [BAR96a] is divided into two studies labelled [BAR96a]#1 and [BAR96a]#2 because of differing case mixes in each study (one includes mortalities, the other study excludes them).

Reference	No. Patients	Including Mortalities	Cardiac only	Surgical only	Cardiac Procedures	Retro/Pros
[KNA93]	17105	Yes	No	No	Excludes CABG	Yes/Yes
[BAR96a]#1	2295	Yes	No	Yes	All except Cardiothoracic surgery	Yes/No
[BAR96a]#2	2295	No	No	Yes	All except Cardiothoracic surgery	Yes/No
[MAR95]	692	No	No	Yes	N/A	Yes/Yes
[BEC95]	2435	Yes	Yes	Yes	CABG	Yes/Yes
[MOU95]#1	431	Yes	Yes	Yes	CABG	Yes/Yes
[MOU95]#2	431	Yes	Yes	Yes	CABG	Yes/Yes
[TUM92]	3156	Yes	Yes	Yes	Multivalve, AVR, MVR, CABG, CABG+Valve	Yes/Yes
[KAT88]	1576	No	Yes	Yes	Valve, CABG, CABG+Valve(s)	Yes/No
[TUJ93]	713	No	Yes	Yes	Valve, CABG+Valve(s)	Yes/Yes
[TUJ94]	713	No	Yes	Yes	Valve, CABG+Valve(s)	Yes/Yes

Table 6.02. Summarisation of studies by factors affecting study validity

From the above table, it can be seen that the studies [BAR96a]#1, [BAR96a]#2 and [KAT88] do not have any prospective validation of the models; [BAR96a]#2, [MAR95], [KAT88], [TUJ93] and [TUJ94] do not include mortalities; [BAR96a]#1, [BAR96a]#2 and [MAR95] do not include cardiothoracic patients in the study population, and [KNA93] excludes the most common cardiac procedure performed at RBH HDE (CABG). With all of these factors, 6 out of the 10 studies would be excluded from the literature review if each factor was considered justification for exclusion. Pragmatically, therefore, all studies will be included, with the proviso that the above factors will be considered when drawing any firm conclusions.

The following summarises the method of each study.

[BAR96A]

All patients were admitted to an intensive care unit. Within the first 24 hours of admission, APACHE II and APACHE III scores were calculated for each patient. All patients were followed until hospital discharge or death. The multiple organ dysfunction (MOD) score [REF] was also calculated for each patient within the first 24 hours of admission to the intensive care unit and for each 24 hour period thereafter until discharge.

A set of clinical and demographic variables were selected based on prior research studies. The data for these variables were used to predict the probability of mortality, the length of stay in the intensive care unit and resource consumption (measured as TISS score during the first 72 hours in the intensive care unit) using statistical regression. A group jacknife procedure was used to validate the prediction models by randomly assigning each patient to one of ten groups and using one group as a test group to test the regression model generated by the remaining groups. This procedure was repeated by excluding each group in turn and using it as the test group.

The accuracy and explanatory power of the prediction models was measured by the area under the receiver-operator curve (ROC) and R-squared values for each predicted variable.

[BUC94]

The authors of this study took length of stay to be a binary measure called chronicity. If patients stayed in the intensive care unit for longer than 7 days, they were classified as chronic and non-chronic if the stay was less than 7 days. The patient variables selected for the prediction models were selected on the basis of whether there was a significant difference between those patients who (retrospectively) stayed longer than 7 days in the intensive care unit and those who stayed less than 7 days for the variable. To avoid the effect of patient's whose condition was so severe that they died before the 7 days threshold introducing bias into this measurement, those patients were excluded.

The authors used their data to develop four different models – one statistical regression model and four different neural network models: "a back propagation neural network with a single associative layer containing nine neurons, a generalized regression neural network, and a probabilistic neural network".

[KAT88]

The authors of this study performed a simple comparison of means for the predicted variable of length of hospitalisation for various patient variables. A forward-selection regression analysis was then performed on the data to "determine the relative importance of factors as independent predictors of length of hospitalization. Interactions between age and sex, NYHA Class, preoperative MI, hypertension, diabetes mellitus, and previous operation as well as between NYHA Class and type of operation, urgency, preoperative MI, and previous operation were analyzed."

A similar procedure was then performed for the predicted variable of hospital mortality. The variables selected for both analyses were selected on the basis of previous research.

[KNA93]

The method used in this study was the same as that used in [BAR96a] above, with the exception that the study was a multi-centre cohort study involving additional variables such as "geographic region, bed size, and teaching status".

[MAR95]

This study looked at the correlation between the Multiple Organ Dysfunction Score (MODS) to length of stay in an intensive care unit.

A logistic regression equation was developed which related the component MODS score of each of six different organ systems with mortality and length of stay in an intensive care unit. The sensitivity and specificity of MODS were calculated for different threshold values of predicting mortality and plotted on a Receiver-Operator Curve (ROC) which provides a graphical representation of the strength of the accuracy of MODS as a predictive tool with the area under the ROC indicating greater predictive value.

[TUJ93]

In this study an artificial neural network (ANN) was developed for predicting chronic length of stay in an intensive care unit. The variables thought to be predictive were encoded into the input layers of the ANN. The prediction of a prolonged length of stay was made by an output node of the ANN, with the output being a continuous value between 0 and 1, and being interpreted as a probability measure of the patient having a prolonged length of stay (1) or a short length of stay (0).

[TUJ94]

Patients admitted to an intensive care unit were divided into two sets according to whether they had a prolonged stay in the intensive care unit (greater than 2 days) or not. A multivariate logistic regression model was used on a set of patient clinical and demographic variables to develop a set of co-efficients. The clinical and demographic variables used in the logistic regression model were selected from a univariate analysis from each variable available in the intensive care unit's dataset, with only those variables with a correlation significant at p<0.01 with length of stay being included in the regression model.

A predictive index was created from the multivariate logistic regression model by taking the odds-ratios for each variable and rounding up to the nearest integer. The authors developed two multivariate logistic regression models – one which included mortalities, and one which excluded mortalities.

The model was validated by splitting the patient population into a test group and a training group. The accuracy of the predictive index was estimated by the area under a Receiver-Operator Curve.

[TUM92]

A univariate analysis of clinical and demographic variables was performed to evaluate their prediction of various forms of morbidity and mortality using pearson chi-squared statistic. Independent predictive variables were identified using a forward stepwise logistic regression model. The resulting predictive variables were each assigned weights according to their associated odd-ratio to predict various forms of morbidity and mortality to create a predictive score.

The predictive score was then correlated with intensive care unit length of stay by relating each patient and their respective length of stay to a points interval in the predictive score for morbidity and mortality. Lengths of stay between different points intervals were tested for significance using analysis of variance.

6.1.2. Univariate Analyses

The union of all of the clinical and non-clinical variables represents a large set. Table 6.03 below summarises all of the variables which were used in the studies. In the table each variable is grouped into the following variable groups:

1. Demographics
2. Hospitalisation Variables
3. Cardiologic Variables
4. Operative Variables
5. Non-Cardiologic Variables
6. Component Scoring Systems

The variable group Component Scoring Systems represents variables which are themselves the scores of existing clinical scoring systems evaluating, for example, neurocognitive functioning in the case of the Glasgow Coma Score.

Each variable is categorised according to whether it may be evaluated pre-operatively or not. This is an important categorisation in developing prediction models of length of stay when those models are to be used in enhancing the level of control over operational cost-effectiveness, since those predictions must be made before any monotonic resource allocation decisions have been made.

For each variable, in the case of categorical variables, the number of categories comprising the variable is listed in Table 6.03. In those cases where the variable is continuous, the variable is listed in Table 6.03 as 'C'.

For those studies where a preliminary univariate analysis of the variable is performed in terms of its correlation with length of stay variables, the variable is listed in Table 6.03 according to whether the correlation was found to be significant at the p<0.05 or the p<0.005 levels (represented in Table 6.03 by the shading of the relevant cell)

Variable	**Pre-Op**	**BEC95**	**KNA93**	**TUM92**	**KAT88**	**BAR96 #1**	**BAR96 #2**
Demographics							
Age	Yes	6	7	3	7		
Family History	Yes				2		
Sex	Yes	2		2	2		
Smoking	Yes				2(1,6)		
Hospitalisation Variables							
Non ICU LOS	Yes		6				
Previous Hospital Location	Yes		7				
Reason for ICU Admission	Yes		78				
ICU Readmission	No		2				
Bed Size	Yes		5				
Hospital Location	Yes		4				
Teaching Status	Yes		3				
Cardiologic Variables							
Angina Grade	Yes						
Cardiovascular Function	Yes						
Chronic Pulmonary Disease	Yes			2			
Congestive Heart Failure	Yes			2			
Exercise Tolerance	Yes						
Hypertension	Yes				2		
LVEF	Yes			2			
LVEDP	Yes						

Variable	Pre-Op	BEC95	KNA93	TUM92	KAT88	BAR96 #1	BAR96 #2
Pre-op MI	Yes			3	5		
Pulmonary Hypertension	Yes			2			
Serious Arrythmias	Yes			2			
Unstable Angina	Yes			2			
Operative Variables							
Bypass Conduit	Yes						
Bypass Time	No						
Left Main Stem Stenosis	Yes						
Number Diseased Coronaries	Yes	5					
SVG	Yes						
Operative Category	Yes	1		5	3		
Previous Procedure	Yes	2		2	2		
Urgency/Operative Status	Yes	2	3	2	2	2	2
Non Cardiologic Variables							
Cerebrovascular Condition	Yes			2			
Diabetes	Yes			2	4		
Hematologic Function	Yes						
Hepatic Function	Yes						
Renal Function	Yes			2			
Respiratory Function	Yes						
Weight/BMI	Yes			2,2	4		
Component Scoring Systems							
APACHE II	No					C	C
APACHE III	No					C	C
APS (of APACHE III)	No	7	9				
Comorbidity	Yes		7				
Glasgow Coma Score	No						
MODS (Re: MAR95)	No					C	C
NYHA Class	Yes				4		

Factor	Pre-Op	MAR95	TUJ93	TUJ94	MOU95 #1	MOU95 #2
Demographics						
Age	Yes		3	3	C	C
Family History	Yes					
Sex	Yes		2	2	2	2
Smoking	Yes				2	2
Hospitalisation Variables						
Non ICU LOS	Yes					
Previous Hospital Location	Yes					
Reason for ICU Admission	Yes					
ICU Readmission	No					
Bed Size	Yes					
Hospital Location	Yes					

Factor	Pre-Op	MAR95	TUJ93	TUJ94	MOU95 #1	MOU95 #2
Teaching Status	Yes					
Cardiology						
Angina Grade	Yes				2	2
Cardiovascular Function	Yes	5				
Chronic Pulmonary Disease	Yes				2	2
Congestive Heart Failure	Yes					
Exercise Tolerance	Yes				2	2
Hypertension	Yes				2	2
LVEF	Yes		4	4	2,C	2,C
LVEDP	Yes				2,C	2,C
Pre-op MI	Yes				2	2
Pulmonary Hypertension	Yes					
Serious Arrythmias	Yes					
Unstable Angina	Yes				2	2
Operative Variables						
Bypass Conduit	Yes				2	2
Bypass Time	No				C	C
Left Main Stem Stenosis	Yes				2	2
Number Diseased Coronaries	Yes				2	2
SVG	Yes				2	2
Operative Category	Yes		3	3	1	1
Previous Procedure	Yes			2	2	2
Urgency/Operative Status	Yes		3	3	2	2
Non Cardiologic Variables						
Cerebrovascular Condition	Yes					
Diabetes	Yes				2	2
Hematologic Function	Yes	5				
Hepatic Function	Yes	5				
Renal Function	Yes	5			2	2
Respiratory Function	Yes	5				
Weight/BMI	Yes				2	2
Component Scoring Systems						
APACHE II	No					
APACHE III	No					
APS (of APACHE III)	No					
Comorbidity	Yes		2(1,7)	3		
Glasgow Coma Score	No	5				
MODS (Re: MAR95)	No					
NYHA Class	Yes			4		

Table 6.03. Summary of variables included in each study

KEY:

C	=	Continuous variable
N	=	Number of categories
2(N1,N2)	=	Category subdivided into N1 and N2 subcategories
N1,N2	=	Factor classified into two seperate categories
	=	$p < 0.05$ (in univariate analysis)
	=	$p < 0.005$ (in univariate analysis)

LVEF = Left ventricular ejection fraction
LVEDP = Left ventricular end diastolic pressure
SVG = Number Spahenous Vein graft distal anastomoses

Table 6.04. Key to Table 6.03

6.1.3. Multivariate Analyses

Any prediction model of intensive care unit length of stay, or any other clinical outcome has at least one clinical or non-clinical variable as input variables. Table 6.05 below lists all of the same variables included in Table 6.03 above and summarises each variable's role in each of the prediction models developed in the studies. In addition, Table 6.05, in the first row, notes how many categories were used to define the outcome length of stay variable. In most cases, the variable was a binary category – either the patient stayed longer than a particular number of days or hours, or they didn't. In those cases where the outcome length of stay prediction is a continuous variable, it is represented in Table 6.05 as 'C'.

For each variable for each study, Table 6.05 denotes whether the variable was measured as significantly predictive in the prediction model by the study (if not, it is represented as 'N/P'), and if so, the rank ordering of the variable amongst the other variables included in the prediction model according to its predictive power. Thus, for example, a value of 2/6 for a variable V signifies that it was the second most predictive variable in the prediction model out of a total of 6 variables included. Predictive power is normally calculated by the regression co-efficient.

Factor	**Pre-Op**	**BEC95**	**KNA93**	**TUM92**	**KAT88**	**BAR96 #1**	**BAR96 #2**
No. LOS Categories		**C**	**C**	**7**	**2**	**2**	**2**
Demographics							
Age	Yes	2/6	4/10	2/11	1/9		
Family History	Yes				N/P		
Sex	Yes	5/6		8/11	7/9		
Smoking	Yes				N/P		
Hospitalisation Variables							
Non ICU LOS	Yes						
Previous Hospital Location	Yes		3/10				
Reason for ICU Admission	Yes		2/10				
ICU Readmission	No		6/10				
Bed Size			8/10				
Hospital Location			5/10				
Teaching Status			10/10				
Cardiology							
Angina Grade	Yes						
Cardiovascular Function	Yes						
Chronic Pulmonary Disease	Yes			N/P			
Congestive Heart Failure	Yes			10/11			

Exercise Tolerance	Yes						
Hypertension	Yes				N/P		
LVEF	Yes			11/11			
LVEDP	Yes						
Pre-op MI	Yes			3/11	6/9		
Pulmonary Hypertension	Yes			9/11			
Serious Arrythmias	Yes			N/P			
Unstable Angina	Yes			N/P			
Operative Variables							
Bypass Conduit	Yes						
Bypass Time	No						
Left Main Stem Stenosis	Yes						
Number Diseased Coronaries	Yes	6/6					
SVG	Yes						
Operative Category	Yes			4/11	2/9		
Previous Procedure	Yes	3/6		7/11	5/9		
Urgency/Operative Status	Yes	4/6	9/10	1/11	8/9	2/3	1/2
Non Cardiologic Variables							
Cerebrovascular Condition	Yes			6/11			
Diabetes	Yes			N/P	9/9		
Hematologic Function	Yes						
Hepatic Function	Yes						
Renal Function	Yes			7/11			
Respiratory Function	Yes						
Weight/BMI	Yes			N/C,N/P	4/9		
Component Scoring Systems							
APACHE II	No					N/P	N/P
APACHE III	No					3/3	N/P
APS (of APACHE III)	No	1/6	1/10				
Comorbidity	Yes		7/10				
Glasgow Coma Score	No						
MODS (Re: [MAR95])	No					1/3	1/2
NYHA Class	Yes				3/9		

Factor	**Pre-Op**	**MAR95**	**TUJ93**	**TUJ94**	**MOU95 #1**	**MOU95 #2**
No. LOS Categories		**6**	**2**	**4**	**2**	**2**
Demographics						
Age	Yes		N/C	2/5	N/P	N/C
Family History	Yes		N/C			
Sex	Yes		N/C	5/5	N/P	N/P
Smoking	Yes				N/P	N/P
Hospitalisation Variables						
Non ICU LOS	Yes					
Previous Hospital Location	Yes					
Reason for ICU Admission	Yes					
ICU Readmission	No					
Bed Size						
Hospital Location						
Teaching Status						

Cardiologic Variables						
Angina Grade	Yes				N/P	N/P
Cardiovascular Function	Yes	5/5				
Chronic Pulmonary Disease	Yes				N/P	N/P
Congestive Heart Failure	Yes					
Exercise Tolerance	Yes				N/P	N/P
Hypertension	Yes				N/P	N/P
LVEF	Yes				N/C	N/P
LVEDP	Yes		N/C	4/5	N/P	N/P
Pre-op MI	Yes				N/P	N/P
Pulmonary Hypertension	Yes					
Serious Arrythmias	Yes					
Unstable Angina	Yes				N/P	N/P
Operative Variables						
Bypass Conduit	Yes				N/P	N/P
Bypass Time	No				N/U	N/C
Left Main Stem Stenosis	Yes				N/P	N/P
Number Diseased Coronaries	Yes				N/C	N/P
SVG	Yes				N/P	N/P
Operative Category	Yes		N/C	3/5	N/P	N/P
Previous Procedure	Yes			N/P	N/P	N/P
Urgency/Operative Status	Yes			1/5	N/P	N/P
Non Cardiologic Variables						
Cerebrovascular Condition	Yes					
Diabetes	Yes				N/P	N/P
Hematologic Function	Yes	4/5				
Hepatic Function	Yes	N/P				
Renal Function	Yes	2/5			N/P	N/P
Respiratory Function	Yes	3/5				
Weight/BMI	Yes		N/C		N/P	N/P
Component Scoring Systems						
APACHE II	No					
APACHE III	No					
APS (of APACHE III)	No					
Comorbidity	Yes		N/C	N/P		
Glasgow Coma Score	No	1/5				
MODS (Re: [MAR95])	No					
NYHA Class	Yes			N/P		

Table 6.05. Summary of the results for each variable for each study

KEY:

X/Y = xth most predictive factor out of y total predictive factors of secondary analysis
N/P = factor has significance in primary (i.e univariate) analysis, but not in second (multivariate) analysis
N/C = factor has predictive power, but no ranking of factors/signifance levels were calculated/provided in secondary analysis
N/U = factor used in primary analysis but not in secondary analysis
LVEF = Left ventricular ejection fraction
LVEDP = Left ventricular end diastolic pressure
SVG = Number Spahenous Vein graft distal anastomoses

Table 6.06. Key to Table 6.05

6.1.4. Conclusions

Because many of the prediction models developed in the studies either use different experimental methods, different case-mixes or different outcome variables, the comparison of each model's predictive power is largely meaningless.

The following lists summary comments from each of the studies, noting in each case the authors' interpretation of the utility and accuracy of the prediction model they have developed.

[BAR96A]

Correlating APACHE scores on admission with the progressive MODS score during treatment in the intensive care unit allowed prolonged length of stay to be identified more accurately and earlier than by use of APACHE scoring alone.

The results suggested that prolonged intensive care unit length of stay be considered as anything longer than 21 days as after that point the APACHE score on admission and MODS tend to level off. The R-squared value, which measures the predictive power of a regression equation in terms of the proportion of variance in the dataset it explains, for hospital mortality was 0.85, although for intensive care unit length of stay was only 0.08, and 0.13 for the first three days TISS score. The authors note that, while these values of R-squared is low across individuals, "there is a fairly strong and substantial impact across groups".

[BUC94]

The results showed that the neural network models outperformed the logistic regression models of length of stay prediction. The authors conclude that neural network modelling "surrenders insight for greater predictive power. A neural network implicitly constitutes a model, but only the predictions - not the model itself - remain accessible to the investigator."

[KAT88]

The results showed that, with all of the predictive factors present in a particular patient, the mean length of stay was only two days longer than if all of those factors were absent. A more significant difference was found amongst patients at extreme age groups, with prolonged lengths of stay in the age groups 20-30 and 80-90 years being 6-7 days more (equivalent to an increase of 60% above the mean).

[KNA93]

The authors classified each variable into variable groups and the combined predictive power of each group was calculated. It was found that a patient's underlying physiology is the most important (48.7% of total predictive power) with the second most important group being the characteristics of the disease and the disease process (34.1%).

The authors concluded that the same variables which predicted mortality were also involved in predicting length of stay, although "Physiology, Age, and Chronic Health are all less important in predicting LOS than mortality. Disease and Other variables are more important".

The resulting R-squared value for the length of stay regression equation was 0.15. However, the authors suggest that much of this relatively low value is due to those patients with prolonged lengths of stay, with R-squared increasing to 0.23 if length of stay is truncated at 15 days rather than the 40 days used in the study.

The authors argue that the reason for the R-squared value for the length of stay prediction being substantially lower than that for mortality is "greater random variations in whether a patient is discharged on a particular day, measurement of LOS in days rather than hours, and the complex relationship between LOS and severity (i.e., for mortalities, LOS will be less the greater the severity since they will die sooner)".

[MAR95]

The authors conclude in their comparison between MODS and APACHE that organ dysfunction, as measured by MODS, is a more important determinant of outcome than APACHE for both intensive care unit mortality and length of stay. When compared against a modified version of MODS, APACHE had a beta coefficient of 0.13, compared with 0.51 for the modified MODS score.

[MOU95]

The authors note that prolonged lengths of intensive care unit stay are often the result of "uncommon but severe complications such as perioperative myocardial infarction, strokes, and wound infections" and that these complications are not typically predictable in advance through any known clinical screening variables, and that as such an element of unpredictability of outcome and length of stay is unavoidable.

The study concluded that length of intensive care unit stay was particularly associated with the cardiologic variables of left ventricular ejection fraction and left ventricular end diastolic pressure (measured as percentage correctly classified): "The factors with the highest predictive accuracy were low left ventricular end diastolic pressure (90%) and one or two vessel disease (89%). The most sensitive factors were good left ventricular function (80%) and good renal function (87%)."

When comparing two prediction models for whether or not the patient would be suitable for fast track or not, the models achieved high rates of percentage correctly classified (89% and 80%), although the first model resulted in low sensitivity. The most significant predictors of fast track status were left ventricular ejection fraction, left ventricular end diastolic pressure, good renal function, the number of diseased vessels and bypass time (i.e. the amount of time the patient was placed on a cardiopulmonary bypass machine during the operation)

The authors found that only the patient's age and bypass time were significant predictive variables of length of stay.

[NIC87]

The authors concluded that age is less important than the severity of illness in predicting outcome. It was noted, however, that there could have been an overcompensating effect on the patient's age with older patients receiving 16% greater resource inputs (as measured by TISS) than the younger patients.

The study found that length of intensive care unit stay was shorter for older patients than for younger patients, although there was no significant difference amongst nonsurvivors.

[TUJ93]

The authors measured the predictive power of the neural network model by varying the threshold between short and long length of intensive care unit stay and calculating the area under the ROC. For the training set, ROC was 0.7094 and 0.6960 in the test set.

The authors stratified the output value of the neural network into three intervals. The first interval was assumed to predict a low probability of prolonged length of intensive care unit stay; the second an intermediate probability, and so on. In the low risk group, prolonged length of intensive care unit stay was only 16.3%, and in the high risk group 60.8%. The authors thus concluded that "the network was able to stratify patients into three fairly distinct risk categories for prolonged ICU length of stay following cardiac surgery."

[TUJ94]

The study concluded that five were independent predictors of prolonged length of intensive care unit stay (age, female sex, left ventricular function, urgency of surgery and type of surgery). Moreover, these variables retained their predictive power both when mortalities were included as well as when they were excluded from the study population.

The authors created a stratified risk index with five distinct categories. They concluded that the index "was found to be useful not only for predicting a stay longer than 2 days but also for predicting stays longer than 4, 7 and 10 days.

[TUM92]

The study results indicated that there is a significant relationship between multiple morbidity and length of stay in the intensive care unit. In scoring morbidity, the score correlated well with length of stay, although the authors noted that the variance in lengths of stay was large and increased with length of stay, thus limiting the utility of such a score.

ICU LOS and score correlated well, although SD was large for each interval and constantly increased with LOS

The authors concluded that the use of such scoring systems "should be in application to adjusting severity levels when reporting outcome statistics as well as prospective planning of resource allocations".

6.1.5. Discussion

The objective of this literature review was to argue for the viability of a computerised model of the process 2.07: predictResources as it was presented in the RBH POP model of the previous chapter.

As specified in RBH POP the purpose of 2.07: predictResources in the wider context of the operational model is to generate output predictions to be able to evaluate proposed admission scheduled in the subsequent process 2.08: evaluateSchedule. And the purpose of 2.08: evaluateSchedule is to enhance the operational cost-effectiveness of the RBH high-dependency environment through increasing the amount of control operational manager have over the allocation of resources and patient admissions. The criteria by which 2.07: predictResources is to be judged, therefore, is in terms of its ability to generate the kind of output which can be used by 2.08: evaluateSchedule to enhance operational cost-effectiveness. The hypothesis of this thesis is that this may be done by predicting the amount of resources that will be consumed by patients both queuing for admission to the high-dependency environment as well as those already admitted to the high-dependency environment and queuing for admission to another unit within the environment or discharge to either a lower-dependency unit or the mortuary.

From the studies described above, it is not immediately clear whether this is in principle possible. The study [MOU95] suggests that it is certainly possible to screen patients for suitability for being fast-tracked. This is significant as it goes a long way to satisfying the attribute POP Projected Admission Time Unit [U] Time = [T], since implicit in this attribute is that it is possible to predict the order in which the patients will pass through the component units of the high-dependency environment (in the description of the attribute in the previous chapter, if a patient is not scheduled to be admitted to a particular unit, the time of admission is defined as 0). In fact, in determining a patient's flow through the high-dependency environment, the decision as to whether they will be able to be recovered in the recovery room or will require intensive care unit admission (i.e. whether or not the patient may be fast-tracked or not) is the most important, since this is the only point at which a patient's flow may bifurcate, with all other transitions between unit being either procedural or occasionally the result of complications (such as a patient requiring re-admission to the operating theatre).

However, the critical question as to whether or not the process 2.07: predictResources can be effectively computerised is not whether a patient's flow can be predicted, but rather whether the rate of the flow can be predicted, and the rate of flow will depend on patient's length of stay in each unit. It is relatively easy to predict the length of stay in the operating theatre, since surgical procedures are relatively standardised and the distribution of lengths of stay follows a relatively normal distribution with not much variance. Similarly, it is relatively easy to predict lengths of stay in the recovery room, since patients admitted to the recovery room tend not to be suffering from complicating physiological or co-present disease processes, and are thus less susceptible to experiencing complications during the recovery process which would prolong their length of stay. Moreover, in those cases where complications do occur, or where the patient otherwise fails to recover within an acceptable period of time, the patient is transferred to the intensive care unit.

Thus, the most important unit for which the computerised model of the process 2.07: predictResources need to predict length of stay is the intensive care unit. In this regard, the above studies provide no clear answer as to whether or not this is possible. Many of the studies note that, due to the nature of the length of stay variable, and unlike the variable for mortality, it is a continuous variable rather than a binary one. This inevitably makes the process of prediction more difficult, simply because any prediction model needs to correlate the set of values for the predictive variables with a large set of values for the outcome variable, rather than simply two values – alive or dead.

Many of the studies have attempted to overcome this problem by defining a threshold between short lengths of stay and long or prolonged lengths of stay. Of course, this is not a real solution to the problem, but it does at least allow more direct comparison with the closely related prediction models for mortality. In this regard, the prediction models for length of stay do not compare well, if only for the simple reason that, whereas greater severity of illness tends to correlate linearly with increased probability of mortality, it does not correlate linearly with length of stay because beyond a certain level of severity of illness, the patient will be so critically ill that their length of stay will be reduced since they will die sooner than those who are less critically ill.

In those studies which did compare a mortality prediction model with a length of stay prediction model using the same study population, the R-squared value is substantially lower for the length of stay prediction model (for example, [BAR96a]). However, these models based themselves on statistical methods that were originally designed for prediction of mortality rather than length of stay – in particular the study method for validating the APACHE severity of illness scoring systems.

The most successful length of stay prediction model developed in the above studies appears to be [TUJ93]. In this study, rather than using a regression-based method, the authors used a neural network model. Moreover, the length of stay outcome variable was defined as a categorical variable based on a stratification of outcomes into short, intermediate and long lengths of stay. The ROC for this study was around 0.7, and its ability to predict short, intermediate and prolonged lengths of stay was good. The question, therefore, is Is it good enough?

In order to answer the question it would be helpful to consider the context in which the question is posed. The hypothesis of this thesis is that a computerised system of patient scheduling can increase operational cost-effectiveness than the current non-computerised system. If the non-computerised system of predicting resource requirements is therefore taken as control, the computerised model of 2.07: predictResources must be able to allow a greater level of operational cost-effectiveness than its non-computerised equivalent for it to be 'good enough'.

The most obvious way for the computerised model of 2.07: predictResources to increase operational cost-effectiveness is for it to be able to generate more accurate predictions than its non-computerised counterpart. This, however, is not the only way, since a computerised model allows for a much more efficient and quicker transmission of information to other processes – in particular 2.08: evaluateSchedule. In the case of human predictors, for example, it is of no consequence whether someone is able to predict the length of stay, for example, of patients if that information is not able to be used in subsequent processes, either because the information is not delivered in a timely manner, or in an incompatible format.

It is reasonable to assume that a computerised implementation of 2.07: predictResources is more effective in being able to integrate with subsequent processes and to deliver information in a format and in good time than a human could achieve.

Regarding the issue of whether a human predictor is able to generate more accurate predictions of lengths of stay, or other indicators of resource requirements, most of the studies included a comparison with human predictors as controls. The only study which did include a comparison with human controls was [MOU95]. In this study, the predicted outcome was suitability for fast-track status. The result was that the computer model has slightly better accuracy than human predictors. The ability to extrapolate this result to other predicted outcomes indicating resource requirements, such as intensive care unit length of stay, is of course very limited. However, with no other basis of drawing any conclusion, it seems reasonable to conclude that computer models are, potentially at least, no worse than human predictors, and may in some cases be better.

In summary, while no firm conclusions can be made, it is nonetheless reasonable to tentatively claim that the process 2.07: predictResources may be computerised successfully, and the most obviously form of this model – for length of stay prediction at least – is an artificial neural network. Thus, we arrive at the outline of a design model which satisfies the requirements of 2.07: predictResources as specified in the previous chapter.

The actual design of the computerised implementation of 2.07: predictResources would need to be a two-tier system comprising two prediction models. The first model would predict the patient's flow through the RBH high-dependency environment (which, as mentioned above, amounts in most cases as to whether or not the patient is suitable for recovery in the recovery room rather than the intensive care units). The second model would predict the length of stay in each unit to which the patient is predicted to require admission. It is probable that the second model would be a complex model consisting of an individual model for each unit, since, for example, the case-mix and severity-mix of an intensive care unit are very different from that of a recovery room.

As specified in the previous chapter, 2.07: predictResources is a pre-processing stage of 2.08: evaluateSchedule. Therefore, on the assumption claim that the process 2.07: predictResources may be computerised successfully, a similar conclusion drawn for 2.08: evaluateSchedule will demonstrate the viability of the computerised operational model RBH POP of the previous chapter.

6.1.6. 2.08: evaluateSchedule

The purpose of 2.08: evaluateSchedule is to provide the processor Operational Manager with the information to be able to control patient admissions and resource allocation which is performed in the process 2.06: admitPatient. This is done by evaluating the schedule as it is proposed in the resource allocation information generated by the process 2.04: allocateBedSlot. In this regard, 2.08: evaluateSchedule thus acts as a decision-support system, providing Operational Manager with information about projected levels of resource utilisation and census data on which he is able to base resource allocation decisions.

Just as the design of the process 2.07: predictResources was as a prediction model, so too 2.08: evaluateSchedule must be a prediction model if it is to make predictions of resource utilisation and census that results from making resource allocation decisions. However, just as the design models examined as possible candidates for the implementation of 2.07: predictResources were suitable for the task of predicting variables such as patient length of stay, so the types of model to be considered as possible candidates for implementing 2.08: evaluateSchedule must be suitable for the task. The purpose of this section is therefore to review work that has been done in this area and evaluate each type of model to select a preferred model. Before this is done, however, it is worth considering again the problem domain and those aspects which determine the final design of 2.08: evaluateSchedule.

The RBH high-dependency system is an example of a progressive-care system where patients are transferred between a series of interdependent healthcare units (or different locations within the same unit), depending on the different stages in the patients' treatment. The transfer of patients between units depends on the clinicians' perception of the progression of the underlying pathology and its treatment. In the RBH high-dependency environment, for example, a patient will usually begin in a pre-operative unit, where the patient is monitored and a treatment plan developed. If appropriate, the patient will then be admitted to the operating room and then to an intensive-care unit or the post-operative recovery room. Upon discharge from the intensive-care unit or post-operative recovery room, the patient may either be admitted to a so-called step-down unit or a general post-operative ward, depending on the facilities available.

Although this kind of patient management system is generally recognised as allowing for improved levels of treatment and monitoring through the division and specialisation of labour and other resources, it also has the potential to have an adverse impact on rates of resource utilisation: The larger the healthcare unit, the smaller is the proportion of beds which need to be set aside as slack in the system to cater for unpredictable admissions to the unit, or unexpectedly long lengths of stay of patients already admitted, since both these events will tend towards relatively constant quantities as number of admissions increases. Thus, in a progressive care system where a large healthcare facility is divided into small specialist units, a higher proportion of the total number of beds in the whole system would be required to be used as slack than would be the case in another facility with all the beds in just one unit..

The problem of low occupancy rates is especially important in the high-dependency environment, where overheads are much higher and the delaying or cancellation of admissions is more likely to result in an adverse outcome.

As argued in Chapter 2, the solution to the problem of managing a progressive care system is improved control over patient admissions and resource allocation. The main aim of scheduling is to control the admission of patients in such a way that the variance in utilisation rates is reduced, allowing for a reduction in system slack and a concomitant increase in utilisation rates. For systems involving a large proportion of emergency – and, by implication, unpredicted – admissions, this is not possible. But for systems such as the RBH high-dependency environment, where the majority of admissions are scheduled in advance, it is not such a great problem.

The (adult) high-dependency environment at the RBH consists of the following units:

1. One 20-bed adult intensive-care unit (AICU).
2. One 5-bed post-operative recovery room (RR).
3. One operating room suite, consisting of 5 fully-equipped operating rooms (OR).
4. Four general hospital wards, each with a 4-bed step-down unit (SDU).

RBH has no emergency room, with most admissions to OR being elective and scheduled and originating from within internal pre-operative wards. Most admissions to AICU originate from OR or tertiary referral from other hospitals.

Most patients at RBH start with admission to OR. Upon discharge from OR, the patient is transferred either to RR or AICU, depending on the severity of illness and the expected duration of the recovery period (patients who are more severely ill and requiring longer recovery periods go to AICU). Upon discharge from RR or AICU, the patient is usually admitted to a post-operative SDU before being transferred to a general ward area. Although this trajectory through the system is typical for operative patients, there are many possible variations on the theme. For example, surgeons may misjudge the post-operative recovery progress, and the patient may require admission to AICU from RR. There is also a significant number of non-operative patients whose initial admission to RBH will be to AICU, with subsequent admission to SDU, or possibly to OR if a surgical intervention becomes necessary. All possible unit-unit patient flows are shown in Figure 6.01below.

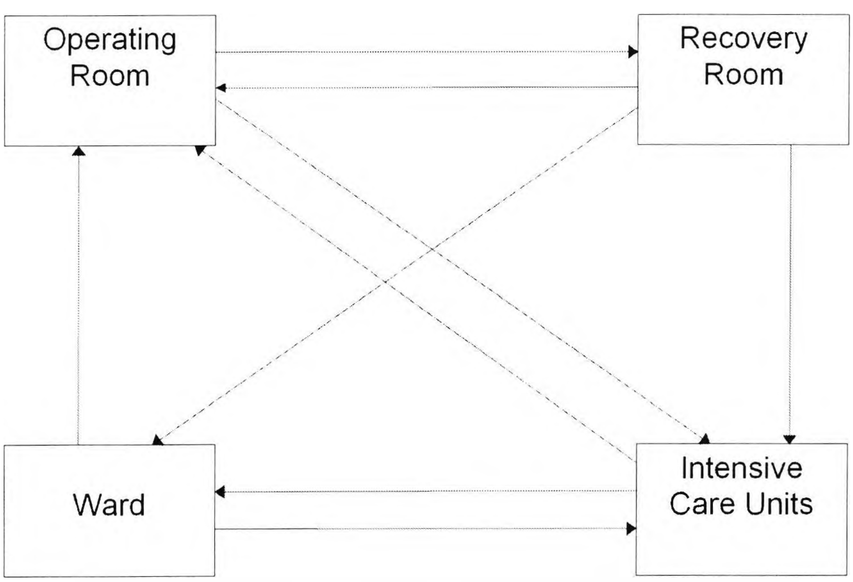

Figure 6.01. Unit-unit patient flows within the RBH High-Dependency Environment

To simplify this discussion and the models resulting from it, cases of mortality will not be considered. Nevertheless, it is important to note that mortality rates are relatively high at RBH due to the critical nature of the patients treated, and that the model could be easily extended by including the hospital mortuary as a component unit of the progressive-care system.

Currently at RBH, the system of scheduling is human-based and is centered around the scheduling of patients into OR. The schedule itself is based on the block-scheduling system [BRE91], [HAM94], [HAN92]. In this system a template is used, which ascribes blocks of OR time for each individual operating room to a particular surgical team, identified by the leading surgeon. This imposes a constraint on scheduling since each patient is under the care of a particular surgeon and must therefore be scheduled for an OR time-slot which is given over to that surgeon. A further constraint is included in the template by indicating the post-operative destination (i.e., RR or AICU) of patients admitted to OR during each time-slot.

Both of the above constraints are flexible in cases of emergency OR admissions, in which case the first available operating room will be used, and if the patient's surgeon is unavailable, an alternative surgeon will operate. Also, if the surgeon's assessment as to the appropriate post-operative destination changes peri-operatively, then the patient will be admitted to AICU instead of RR, or vice-versa.

The aim of scheduling at RBH, as with any healthcare facility, involves the optimisation of many different parameters, while remaining wherever possible within the constraints implied by limited resources. Looking at the process from the patient perspective, the aim is to provide the required resources and to provide them at the right time. From the perspective of the management team, the aim is to maximise the occupancy rates of beds within the progressive care system, while ensuring that the probability of having to deny treatment remains within acceptable limits. There is also the broader political dimension to consider, which requires healthcare facilities to fulfil their contractual obligations without diminishing the quality of care.

The objective of this section is to present the design of a model which may be used to implement the process 2.08: evaluateSchedule whose requirements were specified in RBH POP of Chapter 5. The model to be presented will assist a human operator to evaluate an admissions schedule by simulating the flow of patients around a progressive care system. It will achieve this by using the projections of resource utilisation for individual patients made by the process 2.07: predictResources, in conjunction with a proposed schedule of patient admissions and resource availability information, to predict the values of certain performance variables that would result if the admissions schedule were implemented.

The performance variables specified in Chapter 5 which are required of 2.08: evaluateSchedule are as follows:

POP ALLOCATED-OCCUPIED BED SLOTS DISCREPANCY TIME = [T]

POP Allocated-Occupied Bed Slots Discrepancy Time = [T] is a component attribute of the Unit object class and measures the difference between the projected total number of allocated and occupied bed slots at time T and the actual number of allocated and occupied bed slots at time T in the unit. It is derived from other attributes' values, where those attributes may be components of object classes other than Unit and is edited by the processor CAPSS. It may not be subsequently updated or modified in any way by CAPSS in normal operation of the system, except in the context of policy development or error-correction

POP MEAN OCCUPANCY RATE DISCREPANCY PERIOD = [P]

POP Mean Occupancy Rate Discrepancy Period = [P] is a component attribute of the Unit object class and measures the difference between the average (mean) projected occupancy rate and the average (mean) actual occupancy rate for bed slots within the unit during a period of time P. It is derived from other attributes' values, where those attributes may be components of object classes other than Unit and is edited by the processor CAPSS. It may not be subsequently updated or modified in any way by CAPSS in normal operation of the system, except in the context of policy development or error-correction

POP OCCUPANCY RATE DISCREPANCY TIME = [T]

POP Occupancy Rate Discrepancy Time = [T] is a component attribute of the Unit object class and measures the difference between the projected and actual proportion of bed slots in the unit which are occupied at time T. It is derived from other attributes' values, where those attributes may be components of object classes other than Unit and is edited by the processor CAPSS It. may not be subsequently updated or modified in any way by CAPSS in normal operation of the system, except in the context of policy development or error-correction

POP PROJECTED NUMBER ALLOCATED-OCCUPIED BED SLOTS TIME = [T]

POP Projected Number Allocated-Occupied Bed Slots Time = [T] is a component attribute of the Unit object class and measures the total projected number of bed slots within the unit which are occupied by a patient at time T. It is derived from other attributes' values, where those attributes may be components of object classes other than Unit and is edited by the processor CAPSS. It may be subsequently updated or modified by CAPSS as the underlying values from which it was derived change over time

POP PROJECTED NUMBER ALLOCATED-UNOCCUPIED BED SLOTS TIME = [T]

POP Projected Number Allocated-Unoccupied Bed Slots Time = [T] is a component attribute of the Unit object class and measures the total projected number of bed slots within the unit which are allocated to a patient, but not occupied by any patient at time T. It is derived from other attributes' values, where those attributes may be components of object classes other than Unit and is edited by the processor CAPSS. It may be subsequently updated or modified by CAPSS as the underlying values from which it was derived change over time

POP PROJECTED NUMBER UNALLOCATED-UNOCCUPIED BED SLOTS TIME = [T]

POP Projected Number Unallocated-Unoccupied Bed Slots Time = [T] is a component attribute of the Unit object class and measures the projected total number of unallocated and-unoccupied bed slots within the unit at time T. It is derived from other attributes' values, where those attributes may be components of object classes other than Unit and is edited by the processor CAPSS. It may be subsequently updated or modified by CAPSS as the underlying values from which it was derived change over time

POP PROJECTED OCCUPANCY RATE TIME = [T]

POP Projected Occupancy Rate Time = [T] is a component attribute of the Unit object class and measures the projected proportion of allocated and occupied bed slots within the unit at time T. It is derived from other attributes' values, where those attributes may be components of object classes other than Unit and is edited by the processor CAPSS. It may be subsequently updated or modified by CAPSS as the underlying values from which it was derived change over time

POP PROJECTED TISS COMPONENT [N] PER BED SLOT TIME = [T]

POP Projected TISS Component [N] per Bed Slot Time = [T] is a component attribute of the Unit object class and measures the total projected amount of the TISS component N consumed per bed slot within the unit per unit of time at time T, where N could be any TISS component which denotes the consumption of a particular healthcare resource or group of resources. It is derived from other attributes' values, where those attributes may be components of object classes other than Unit and is edited by the processor CAPSS. It may be subsequently updated or modified by CAPSS as the underlying values from which it was derived change over time

POP TISS COMPONENT [N] PER BED SLOT DISCREPANCY TIME = [T]

POP TISS Component [N] per Bed Slot Discrepancy Time = [T] is a component attribute of the Unit object class and measures the difference between the projected and actual average consumption of the TISS component N per bed slot within the unit per unit of time starting at time T. It is derived from other attributes' values, where those attributes may be components of object classes other than Unit and is edited by the processor CAPSS. It may not be subsequently updated or modified in any way by CAPSS in normal operation of the system, except in the context of policy development or error-correction

POP ALLOCATED-UNOCCUPIED BED SLOTS DISCREPANCY TIME = [T]

POP Allocated-Unoccupied Bed Slots Discrepancy Time = [T] is a component attribute of the Unit object class and measures the difference between the projected number of allocated and unoccupied bed slots within the unit at time T and the actual number of allocated and unoccupied bed slits within the unit at time T. It is derived from other attributes' values, where those attributes are other attributed of the object class Unit and is edited by the processor CAPSS. It may not be subsequently updated or modified in any way by CAPSS in normal operation of the system, except in the context of policy development or error-correction

POP LABOUR COMPONENT [N] PER BED SLOT DISCREPANCY TIME = [T]

POP Labour Component [N] per Bed Slot Discrepancy Time = [T] is a component attribute of the Unit object class and measures the difference between the projected and actual average amount of the labour component N that is consumed per bed slot per unit of time at time T within the unit. It is derived from other attributes' values, where those attributes are other attributed of the object class Unit and is edited by the processor CAPSS. It may not be subsequently updated or modified in any way by CAPSS in normal operation of the system, except in the context of policy development or error-correction

POP MEAN PROJECTED OCCUPANCY RATE PERIOD = [P]

POP Mean Projected Occupancy Rate Period = [P] is a component attribute of the Unit object class and measures the average (mean) projected occupancy rate for the bed slots within the unit during a period of time T. It is derived from other attributes' values, where those attributes are other attributed of the object class Unit and is edited by the processor CAPSS. It may not be subsequently updated or modified in any way by CAPSS in normal operation of the system, except in the context of policy development or error-correction

POP UNALLOCATED-UNOCCUPIED BED SLOTS DISCREPANCY TIME = [T]

POP Unallocated-Unoccupied Bed Slots Discrepancy Time = [T] is a component attribute of the Unit object class and measures the difference between the actual and projected number of unallocated and unoccupied bed slots within the unit at time T. It is derived from other attributes' values, where those attributes are other attributed of the object class Unit and is edited by the processor CAPSS. It may not be subsequently updated or modified in any way by CAPSS in normal operation of the system, except in the context of policy development or error-correction

For many of these variables such as POP Unallocated-Unoccupied Bed Slots Discrepancy Time = [T], the value will actually be calculated from a comparison of the value predicted by the implementation model of 2.08: evaluateResources and the actual value. Other variables, such as POP Projected Number Unallocated-Unoccupied Bed Slots Time = [T] are calculated by considering the proposed admissions schedule, the proposed availability of staff and other resources, and most importantly – the predictions of lengths of stay and other measures of projected resource consumption made by the component prediction models of 2.07: predictResources.

The discrimination between these two types of variable is important due to the role they play in the effective control of patients admissions and resource allocation. The former type of variable play a role in an adaptive control loop, which has an indirect on the effectiveness of control by allowing the control and optimisation of the model itself, while the latter type of variable presents the domain-level control loop which has a direct effect on the effectiveness of control by providing the necessary information to evaluate proposed admissions schedules prior to implementation, and thus allowing the Operational Manager to make amendments as necessary.

The modelling methodology that will be used in the design of the model for 2.08: evaluateSchedule is coloured-timed Petri nets (CTPNs) [COS92], [DAS91], [LIN98], [MOO95], [TAM97]. It will be argued that the CTPN formalism is particularly suited to the problem since it allows the dynamics of the system to be sensitive to different instantiations of system variables such as processing time, that is, sensitive to different patients having different lengths of stays and admissions requirements.

Before the model itself will be presented, alternatives to the CTPN-based model will be considered in the form of a literature review. The formalism of Petri-nets and then coloured-timed Petri-nets will be presented. Following the development of the CTPN model that demonstrates the in-principle feasibility of the computerisation of 2.08: evaluateSchedule as it is specified in RBH POP, there will be a discussion of the main features of the model and how it relates to the wider context of controlling patient admissions and resource allocation at RBH high-dependency environment.

6.1.7. Literature Review

Any model which aims to evaluate the effectiveness of a schedule must be able to model the flow of patients in to and out of, as well as within, the system. That is, it must be able to predict with reasonable accuracy the census of each unit within the system at each point in time that would result from implementing the proposed schedule. In the next section various approaches which have been proposed for modelling patient flows within a healthcare facility will be briefly discussed. It will be argued that of all the different approaches, CTPNs are best able to model the complex patient flows within a progressive care system.

Various formalisms have been deployed to model patient flows in a healthcare facility. Proposed models may be broadly classified as being either analytic or simulation models. Of the analytic models proposed, most came out of the operational research community in the '60s and '70s and are usually based on either Queuing Theory or Markov/semi-Markov processes. The queuing models which have been proposed are generally able to easily model the performance of individual units corresponding to different mixes of patient types and admissions policies [ESO76], [LAM95], [YOU65], [YOU66]. The application of Queuing Theory to the performance analysis of progressive care systems is very difficult to model mathematically, however, given the complex feedback relationships that can arise during certain periods of operation.

An alternative to Queuing Theory is the use of Markov [NAV70], [STA71] or semi-Markov [HER81], [KAO73], [KAO74], [KAO72] processes to model patient flow. In the closed versions of these models, patient-flows between units are constructed from the individual transition probabilities P_{ij} of a patient being transferred from unit i to unit j during one transition time period. Hershey et al ([HER81]) have shown how a semi-Markov model can be used to derive some important performance measures regarding utilisation rates and the probability of full capacity. However, it is more cumbersome to model different types of patient or admissions policies in Markov-based models, since a new matrix is required for each new patient type or policy.

Despite the importance of simulation in healthcare being well recognised (Mahachek, 1992), most simulation models proposed have been used only for the validation of analytic models (see, for example, Vassilacopoulos, 1985 [VAS85]). We believe that the best way to approach the problem of modeling patient flows is through the use of a computer simulation model. Simulation models have the potential flexibility to model many different patient types and admissions policies, as well as the many different patient flows within a progressive care system.

The advantages and disadvantages of the three approaches are summarised Table 6.07 below, where a simple binary classification is made as to whether or not the formalism may EASILY incorporate those evaluation criteria mentioned above.

Formalism	**Queuing**	**Markov**	**Simulation**
Can model progressive-care system?	No	Yes	Yes
Can model different patient types?	Yes	No	Yes
Can model different admissions policies?	Yes	No	Yes

Table 6.07. Comparison of different modelling approaches.

6.1.8. Petri Nets

A basic Petri net may be defined by the 4-tuple (P, T, A, M_0), where

- P is a set of places $P = (p_1, p_2, \ldots, p_n)$ representing states
- T is a set of transitions $T = (t_1, t_2, \ldots, t_m)$, representing functions
- A is a set of directed arcs $A = (P \times T)$, connecting places and transitions
- M_0 is the initial marking of the system. A marking M_i is defined over places such that $M_i = \{\mu_i(p_1), \mu_i(p_2), \ldots, \mu_i(p_n)\}$. The marking M_0 defines the system state at time t = 0. That is, $M_0 = \{\mu_0(p_1), \mu_0(p_2), \ldots, \mu_0(p_n)\}$.

The marking of a basic Petri net is represented in terms of tokens and determines whether a transition will be enabled or not. A transition, t_j is enabled when

Eq.6.01 $$\mu(p_i) \geq I(p_i, t_j)$$

where $I(p_i, t_j)$ is the input mapping function for t_j and p_i such that:

Eq.6.02 $I(p_i, t_j): p_i \times t_j \rightarrow N$, where N is an integer

When the transition is enabled, it fires, taking tokens from the input place and placing tokens into the output place resulting in the marking

Eq.6.03 $\mu'(p_i) = \mu(p_i) + O(p_i, t_j) - I(p_i, t_j)$

where $O(p_i, t_j)$ is the output mapping function of transition t_j such that:

Eq.6.04 $O(p_i, t_j): p_i \times t_j \rightarrow N$, where N is an integer

Time may be introduced into a basic Petri net by either defining a firing delay for transitions, or an enabling delay for places. Defining time over places results in the 5-tuple (P, T, A, M_0, Γ), where

Eq.6.05 $\Gamma = (\gamma_1, \gamma_2, \ldots, \gamma_n),\ \gamma_i \geq 0$

In which case a transition t_j only becomes enabled after a delay γ, and the inequality 1) holds.

A more in-depth description of Petri nets may be found in Peterson (1981). In the next section the Petri net formalism introduced here will be used as the basis of an introduction to CTPNs.

6.1.9. Coloured-Timed Petri Nets (CTPNs)

In coloured Petri nets, the tokens are divided into types, called 'colours' for historical reasons. This allows for the basic Petri net as presented above to be abbreviated through duplicated structural features being subsumed into a single net with no redundancy, the basic idea being that enabling conditions, markings and other extensions to the basic Petri net formalism may be made sensitive to the colour of the tokens involved. A more in-depth introduction to coloured Petri nets can be found in [JEN89]. A coloured Petri net without timing is thus the 5-tuple (P, T, A, M_0, C), where P, T and A are as before, and C is the colour sets of transitions and places, $C(p)$ and $C(t)$ where,

Eq.6.06 $C(p_i) = \{a_{i1}, a_{i2}, \ldots, a_{iui}\},\ u_i = |C(p_i)|,\ i = (1, 2, \ldots, n)$; and

Eq.6.07 $C(t_j) = \{b_{j1}, b_{j2}, \ldots, b_{jvj}),\ v_j = |C(t_j)|,\ j = (1, 2, \ldots, m)$

Where a and b are colours of places and transitions, and |...| denotes cardinality.

M_0 is reinterpreted in a coloured Petri net such that $\mu_0(p_i)$ denotes the number of tokens of each colour in place p_i. Thus, $\mu_0(p_i)$ is an nx1 vector such that:

Eq.6.08 $$\mu_0(p_i) = \sum_{h=1}^{\mu_i} n_{ih} a_{ih}$$

where n_{ih} is the total number of tokens of colour a_{ih} at time t=0. Similar modifications apply to the enabling and firing of transitions, as follows:

A transition, t_j is enabled for tokens of colour b_{jk} when

Eq.6.09 $$\mu(p_i)(a_{ih}) \geq I(p_i, t_j)(a_{ih}, b_{jk})$$

where $I(p_i, t_j)(a_{ih}, b_{jk})$ is the input function for place p_i with colour a and transition t_i with colour b such that:

Eq.6.10 $$I(p_i, t_j)(a_{ih}, b_{jk}): C(p_i) \times C(t_j) \rightarrow N$$

When the transition fires, it takes tokens from the input place and placing tokens into the output place according to the colour of the tokens. resulting in the marking

Eq.6.11 $$\mu'(p_i)(a_{ih}) = \mu(p_i)(a_{ih}) + O(p_i, t_j)(a_{ih}, b_{jk}) - I(p_i, t_j)(a_{ih}, b_{jk})$$

Where $O(p_i, t_j)(a_{ih}, b_{jk})$ is the output mapping function of transition t_j with colour b to the place p_i with colour a such that:

Eq.6.12 $$O(p_i, t_j): p_i \times t_j \rightarrow N, \text{ where N is an integer}$$

Introducing time into a CTPN results in the 5-tuple (P, T, A, M_0, C, Γ), where

Eq.6.13 $$\Gamma(p_i) = (\gamma_{i1}, \gamma_{i2},\ldots,\gamma_{iui}),\ u_i = |C(p_i)|,\ i = (1, 2, \ldots, n)$$

In which case a transition t_j with colour b_{jk} only becomes enabled with respect to place p_i with tokens of colour a_{ih} after a delay γ_{ih} and the inequality Eq.6.09 holds.

The graphical formalism used to represent CTPNs is shown in Figure 6.02 below.

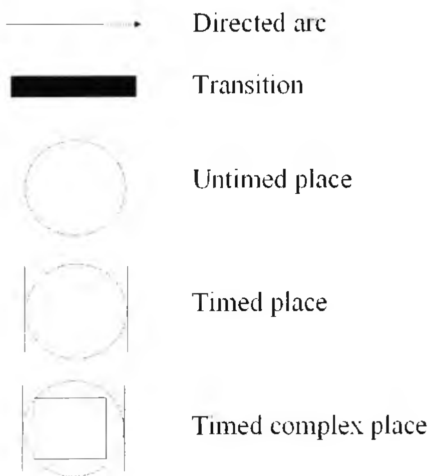

Figure 6.02. Graphical CTPN formalism

6.1.10. The Model

The CTPN model of the patient flows at RBH has a modular structure, each module corresponding to a unit within the progressive care system. Two colour sets are defined as follows:

- *Bed Slot Set*, B: The healthcare resources of the progressive care system are modeled as bed slots. A bed slot for a unit X is assumed to be the capacity to admit a patient for therapy to X for a period of time dependent on the type of patient admitted. B therefore consists of the token colours *or* for OR bed slots, *rr* for RR bed slots, and *icu* for ICU bed slots. Ward bed slots are not modelled explicitly for reasons to be explained below.

- *Patient Set*, Π: Patients are classified in terms of their admissions requirements. That is, the order in which they pass through the units in the system. Because of the large number of possible routes a patient can take through the system, we have simplified the model by considering only the three most common routes, as follows:

  1. $\pi_{1.1}$ = Ward → OR → RR → Ward

  2. $\pi_{1.2}$ = Ward → OR → ICU → Ward

  3. π_2 = Ward → ICU → Ward

Figure 6.03 shows the OR CTPN module. In the graph the arcs have been labelled with the condition states of the transitions. For instance the arc *–or→* means *"if the colour of token = or, then pass"*. Readability of the CTPN model has been simplified by showing arcs labelled with tokens of colour set B as dashed, and arcs labelled with tokens of set Π as solid.

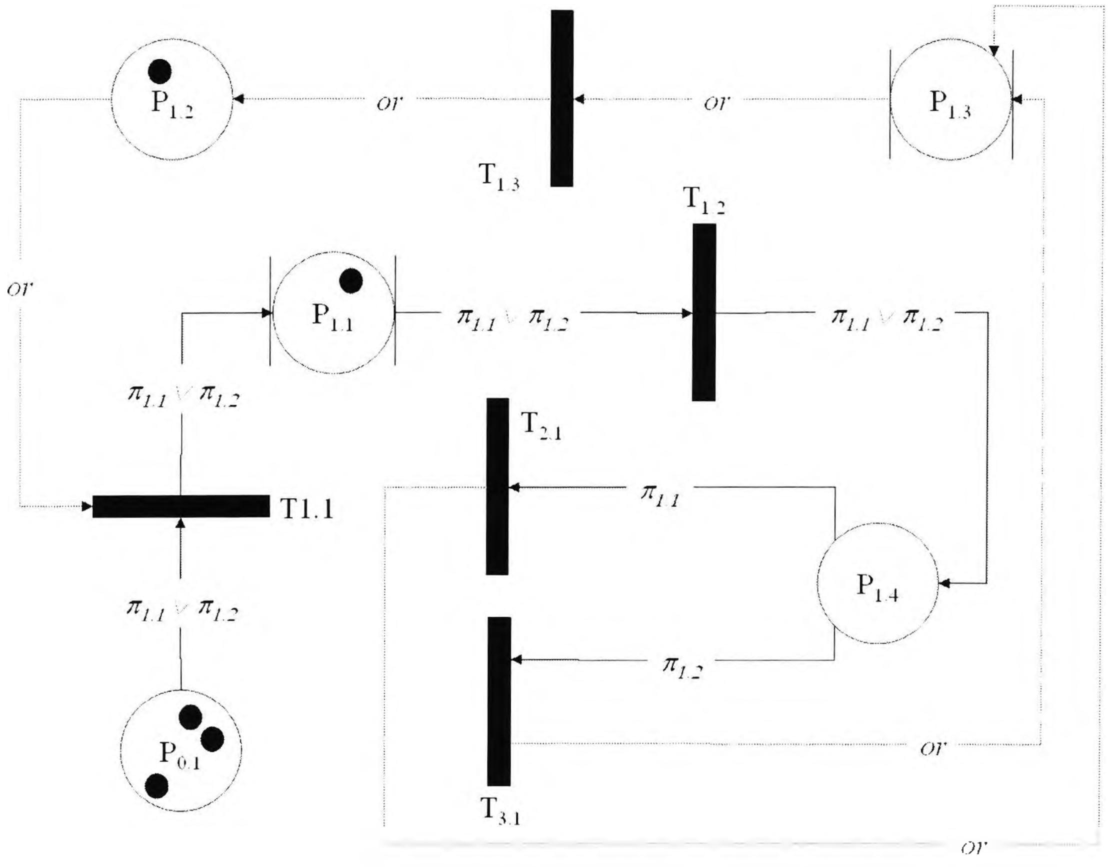

Figure 6.03: CTPN Model of Operating Room

With the initial marking shown in the diagram, there are 2 x $\pi_{1.1}$ patients and 1 x $\pi_{1.2}$ patient queuing for admission to OR (place $P_{0.1}$), there is one available *or* bed slot ($P_{1.2}$), and there is 1 patient admitted to OR ($P_{1.1}$). Transition $T_{1.2}$ represents the process of clinical discharge from OR. Clinical discharge is the process of placing the patient in a queue for admission to another unit. Thus, patients which are clinically discharged from OR ($P_{1.4}$) will still be consuming an *or* bed slot. Transition $T_{2.1}$ represents a complex process comprised of the administrative discharge of a patient from OR, the admission of a patient to RR, and the release of an *or* bed slot. Similarly, $T_{3.1}$ represents the administrative discharge from OR, admission to ICU and the release of an *or* bed slot, and so on.

For the sake of simplicity it is assumed that all patients of the same type will all have the same time delays with respect to the timed places in each unit . In reality, however, there will be many different lengths of stay in each unit for the same type of patient, which requires the further classification of patients into subtypes. This point will be discussed further in the next section. Note that the term 'length of stay' does not necessarily equate to the total length of stay within the unit, since a patient may be required to queue for admission to a subsequent unit. Thus, only when a patient is admitted to a subsequent unit is the bed-slot resource freed, represented in the case of OR by placing a token in $P_{1.3}$.

The transition $T_{1.3}$, represents the process of preparing an *or* bed slot for a subsequent admission to OR. $T_{1.3}$ is only enabled when there is a freed *or* bed slot in the timed place $P_{1.3}$, and a subsequent patient may only be admitted to OR once $T_{1.3}$ has fired and an available *or* bed slot is placed in $P_{1.2}$. The time delay for $P_{1.3}$ may be assumed to be the same for each type of patient, although the delay will vary between units, it being longer, for example, to prepare for a new patient in OR than it does in RR.

Since the processes of clinical discharge, admission/bed slot release/administrative discharge, and bed slot preparation are generic across all units, the structure depicted in Figure 6.03 is duplicated for RR and ICU. The exception is Ward, which represents the general hospital wards, which does not represent the processes of clinical disharge and bed slot preparation, since in the model Ward only ever acts as either an origin unit or sink unit for patients entering or leaving the system.

Since the same basic structure is duplicated for each component unit of the RBH high-dependency environment, and to simplify the model, each CTPN model of each unit may be represented as a set of complex places as an abbreviation of the model to enhance readability, according to the formalism depicted in Figure 6.02 above. In this case, we may combine the two bed-slot places and the two patient places and their associated transitions into complex places. The resulting model for the operating room is depicted in Figure 6.04 below, where the place $OR_{(B)}$ represents the places $P_{1.2}$ and $P_{1.3}$ of Figure 6.03 above, and the place $OR_{(P)}$ represents the places $P_{1.1}$ and $P_{1.4}$. A similar labelling scheme for the complex places of other units will be used as appropriate.

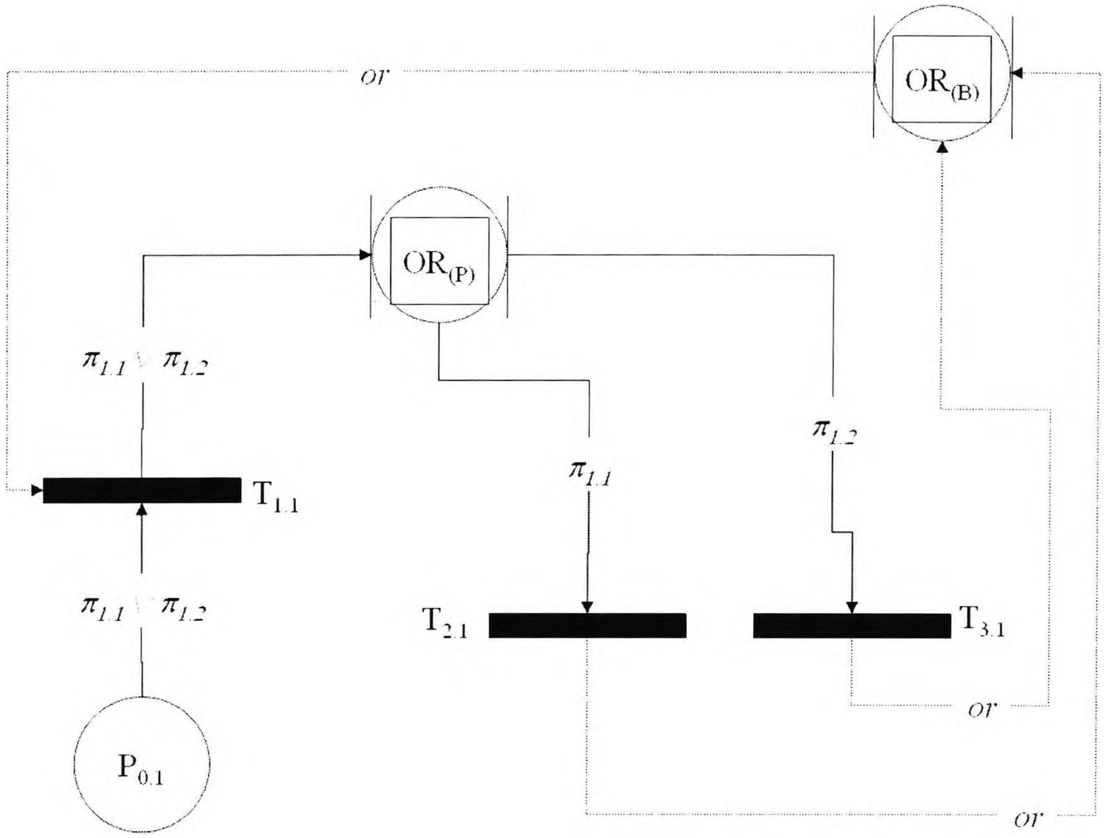

Figure 6.04: Abbreviated CTPN model of the operating room using complex places.

The semantics for each module and the place/transition numbering system are provided in Table 6.08.

Node	Description
P0.i (T0.i)	Ward place (Ward transition).
P1.i (T1.i)	OR place (OR transition).
P2.i (T2.i)	RR place (RR transition).
P3.i (T3.i)	ICU place (ICU transition).
Pi.1	Patient awaiting clinical discharge from unit i (i ≠ 0).
Pi.2	Available bed-slot in unit i (i ≠ 0).
Pi.3	Unit i bed-slot awaiting turnaround (i ≠ 0).
Pi.4	Patient awaiting administrative discharge from unit i (i ≠ 0).
P0.1	Ward patients awaiting admission to OR or ICU
P0.2	Ward patients discharged from ICU or RR.
Ti.1	Patient admission to unit i.
Ti.2	Clinical discharge from unit i (i ≠ 0).
Ti.3	Unit i bed slot preparation (i ≠ 0).

Table 6.08. Module semantics and place/transition numbering system.

The global model of patient flows is constructed by connecting the individual abbreviated unit modules together according to the three possible patient flows defined by the different patient types. The resulting CTPN is shown in Figure 6.05below, where the following abbreviations have been used:

1. Places $P_{i.1}$ and $P_{i.4}$, and transition $T_{i.2}$ are reduced to the complex places $OR_{(P)}$ (i=1), $RR_{(P)}$ (i=2), and $ICU_{(P)}$ (i=3).
2. Places $P_{i.2}$ and $P_{i.3}$, and transition $T_{i.3}$ are reduced to the complex places $OR_{(B)}$ (i=1), $RR_{(B)}$ (i=2), and $ICU_{(B)}$ (i=3).

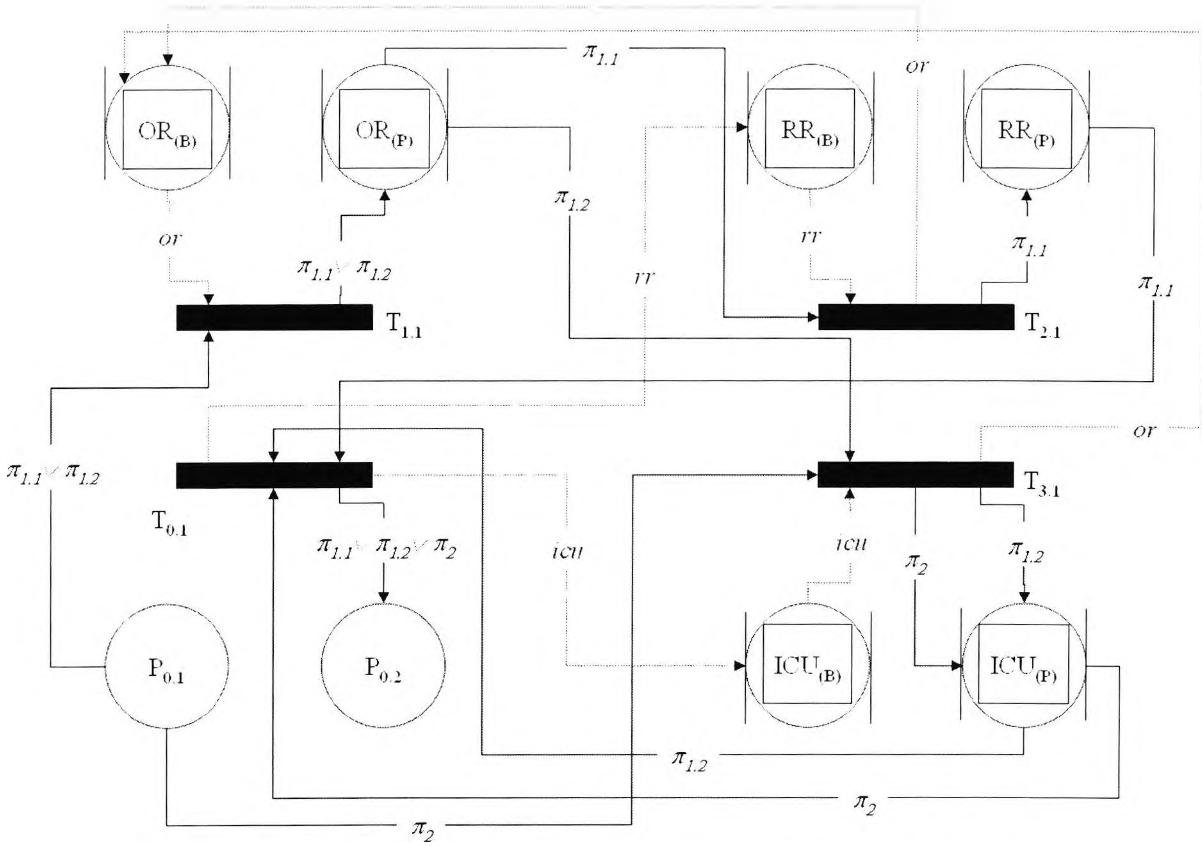

Figure 6.05: CTPN Model of the RBH high-dependency environment.

6.1.11. Discussion

The CTPN model which has been developed in this section has the objective of demonstrating the in-principle feasibility of using the formalism of CTPN to model patients flows within the RBH high-dependency environment for the purpose of schedule evaluation as defined by the process 2.08: evaluateSchedule of the preceding chapter.

The use of CTPNs to model patient flows is particularly suited to the problem of scheduling evaluation since the formalism easily allows for the ad hoc construction of different types of patients by defining a new colour for each length of stay prediction and admissions requirement prediction. Further, the CTPN model we have proposed has the advantage of being modular in structure, which allows it to be more easily modified according to a change in the underlying structure of the clinical setting, or the application of the model to a different setting.

Before the above model may be implemented as a computerised process, various extensions are required due to the simplifying assumptions regarding the problems domain which were made prior to the construction of the model. Among these simplifying assumptions were the absence of mortalities and a simple case-mix of patients. Fortunately, both of these extensions may be included in the model simply and quickly due to the modular structure and the schematic definition of patients by different coloured tokens.

The model presented above would be the central component of any software application designed to encode the process 2.08: evaluateSchedule. Other components to the software application would, of course, include an extensive reporting tool to display the results of each simulation of each proposed admissions schedule and associated resource availability data. It would also include another component for the adaptive control aspect of the system, where model projections and actual values are compared against each other and displayed in a report.

The model would also required to be integrated with the CAPSS database as specified in RBH POP of the preceding chapter. In particular, the model needs to be closely integrated with the data generated by 2.07: predictResources, since the projections of patients' lengths of stay and their projected flow around the system will define the colour set and token properties of the CTPN model. Further, the data regarding the current census and resource availability will define the initial marking, M_0 of the CTPN model.

6.2. Summary

The objectives of this chapter have been to demonstrate the feasibility of the computerisation of processes 2.07: predictResources and 2.08: evaluateSchedule as they were specified in RBH POP.

The feasibility of the computerisation of 2.07: predictResources was demonstrated through a review of the literature on prediction models of intensive-care unit length of stay and related variables that quantify projected resource requirements of individual patients.

The conclusion of the review of the literature was that, while the evidence currently available is minimal on the viability of using models to predict length of stay and other similar variables, the best possibility was to make use of a neural network model which seems to generate more accurate results with a stratification of outcome variables into three categories of length of stay. The models designed for the purpose of predicting the flow of a patient around a progressive-care system showed good results that could be applied to individual cases and incorporated into a computerised model.

The feasibility of the computerisation of 2.08: evaluateSchedule was demonstrated through the development of a coloured-timed Petri-net model of the RBH high-dependency environment capable of simulating patient flows and thereby generating quantifiable results for resource utilisation rates and other performance parameters.

The formalism of coloured-timed Petri nets was preferred over other simulation formalism because of its ability to be easily configured to incorporate the output of the prediction models that comprise 2.07: predictResources as system parameters. A simulation-based approach to schedule evaluation was preferred because of its ability to easily model the complexities of progressive-care systems and the network of multiple and interdependent queuing systems that they represent.

Considered as a single software system, the computerised models of both 2.07: predictResources and 2.08: evaluateSchedule, would be closely integrated with both the CAPSS database, as well as the main patient record database at RBH, CareVue. It would also require a distributed architecture where data could be entered and displayed from various locations around RBH and the model used as a decision-support tool for both operational and strategic members of the RBH management team.

7. Concluding Remarks

7.1. Modelling Healthcare Resource Allocation

The healthcare industry throughout the developed world is experiencing rapid technological and pharmaceutical developments not seen before. Healthcare consumers are also expecting more from their healthcare service providers – particularly publicly funded providers. The result is escalating costs for treatment and a resulting need on the behalf of healthcare providers to cut costs while enhancing the quality of the service provided.

The hypothesis which has been argued here is that healthcare managers can increase the cost-effectiveness of healthcare provision by increasing the level of control they have over the admission scheduling and resource allocation process.

To demonstrate the link between the degree of control that healthcare managers have over the admission of patients and the allocation of healthcare resources, Chapter 2 was devoted to the exploration of some issues in healthcare resource allocation and the development of a mathematical model of healthcare resource allocation.

In Chapter 2 an important classification was made of healthcare resources. Unlike conventional classification systems, that proposed in Chapter 2 was constructed of two concepts: the degree of generality of the resource, and its degree of variability. The concept of generality measured whether or not a resource was generic for all patients admitted to a healthcare unit. That is, whether or not all patients admitted to a healthcare unit consumed the resource or not. The concept of variability measured whether or not the supply of a resource may be changed quickly and easily according to demand.

Unlike other resource classification systems, the generic-fixed taxonomy corresponded well to important properties of resources that are required for modelling healthcare resources. Most particularly, it provides the basis for deriving the concept of a bed-slot. It was argued that, since generic resources – whether fixed or variable in their supply – were consumed by all patients, and hence that there was no healthcare resource allocation decisions to be made regarding those resources by healthcare managers when scheduling patients for admission, the notion of enhancing the degree of control over the consumption of those resources to particular patients is meaningless.

For those resources which were not generic, it was argued that only those which were both non generic and fixed were important components in any model of resource allocation. The reasoning for this was based on the concept of opportunity cost, since non generic and non fixed resources – which broadly correlated with consumable resources such as pharmaceuticals – had little or no opportunity cost if they were not consumed at a particular time by a particular patient other than the cost of storage. Non generic and fixed resources, however – which broadly correlated with capital resources such as beds, surgical equipment, mechanical ventilators and other large expenditure items of capital equipment – do have significant opportunity cost if left idle, simply because of their large cost.

It was thus concluded that a model of patient admission scheduling and resource allocation may be based on the concept of a bed-slot defined as a package of non-generic and fixed healthcare resources.

The concept of a bed-slot was first assumed as a generic entity in a model of resource allocation in a theoretical healthcare unit. In this scenario, the rate of patient admissions was assumed to be normally distributed with the variance in admission rates measuring whether or not those admissions were controlled to any extent or not. In the case where there is no control over admissions and a limited supply of bed-slots at any given time, it was noted that in the resulting distribution of bed-slot consumption rates was heavily skewed to the right, since at those times when demand exceeds supply only that number of patients equal to the number of available bed-slots may be admitted. This important consideration has been ignored in other similar models of healthcare resource allocation.

Given this foundation, the variance in admission rates was argued as a possible measure of the control over healthcare resource allocation. With complete control over admissions – and thereby resource allocation – the resulting occupancy rate of each bed-slot comprising a healthcare unit should, in theory, be 100%. With no control over admissions, however, the resulting occupancy rate of the bed-slots will not be the mean level of demand for bed-slots as argued by other models of healthcare resource allocation because of the skewed nature of the distribution of the admission rates in the case of limited healthcare resources. Rather, the mean occupancy rate will be greater than the mean demand rate. As the level of control over patient admissions increases, so mean occupancy rate increases. Moreover, with increased control, the level of demand for bed-slots can be safely increased while remaining within acceptable denied admission rates.

The relationship between the level of control over resource allocation and the cost-effectiveness of healthcare delivery was demonstrated by the model through the derivation of a simple equation for measuring cost-effectiveness. The equation measured cost-effectiveness in terms of the mean bed-slot utilisation rate and the amount of unsatisfied demand for bed-slots. This equation represents a significant advance in modelling healthcare resource allocation with no other similar measure being developed elsewhere. Only with a well-defined measure of cost-effectiveness of healthcare resource allocation does it make sense to speak of models of healthcare resource allocation.

In the real world, many of the assumptions made in the model simply do not hold. In particular:

1. Each patient has different levels of urgency of admission;
2. Each patient has different lengths of stay upon admission;
3. Each patient's length of stay as judged by clinicians upon admission can change from hour-to-hour;
4. The level of resources available in a healthcare unit, as measured by bed-slots, is not constant either throughout the day, throughout the week or throughout the year.

The theoretical model developed was contrasted with the real world by considering the patient scheduling and resource allocation process as a tiling problem. But, unlike other tiling problems typical in mathematics, the tiling problem representing patient scheduling and resource allocation is one where all tiles are not only of different shapes, but where those shapes and dimensions may change once the tile is placed on the plane. Moreover, the tiling really occurs in a hyper plane rather than a two-dimensional space.

There were also other problems in the real-life scenario. For example, the distribution in the lengths of stay of patients, at least in the high-dependency environment, is not a normal distribution. Rather, it typically has an extremely elongated tail, with a significant proportion of patients staying for many times longer in a particular healthcare unit that the majority of other patients. This implies serious problems for any system aimed at increasing the control over resource allocation with an increase in 'emergency' bed-slots set-aside at any time to safeguard against the situation where one or two patients stay an unexpectedly long period of time which would otherwise block bed-slots for other patients.

With all of these considerations, the theoretical model of patient admission scheduling and resource allocation developed in Chapter 2, while demonstrating the in-principle connection between the level of control and the cost-effectiveness of healthcare delivery, is not sufficient to form the basis of a model to actually enhance the level of control. The development of a model capable of achieving this task was the main objective of the remaining chapters of this thesis.

7.2. The Empirical Domain

The model was developed using the scenario and data from the Royal Brompton NHS Trust's High-Dependency Environment (RBH). As a healthcare facility with a large throughput of critically ill patients, a mix of both surgical and non-surgical admissions, and a suite of interdependent high-dependent units each with their own clearly defined admission and discharge policies, RBH is an ideal healthcare facility to use in the basis of the model.

To attempt to answer the question of whether or not a model could be developed of RBH to satisfy the objective of improving control, a series of studies were devised. The first of these studies looked at a dataset of patients admitted to the Adult Intensive Care Unit (AICU) and the factors which influenced their length of stay in AICU.

All patients consume different resources during their stay in the intensive-care unit. This is implied by each patient requiring different treatment, and therefore different resource input. Even in those cases where two patients have the same diagnosis and undergo the same surgical procedure, there will inevitably be differences in, for example, the amount of a particular drug they consume during their post-operative recovery and convalescence. Nonetheless, all of these differences are relatively trivial in the progressive-care system of healthcare compared to differences in length of stay. After all, the whole point of a progressive-care system is, in essence, to group patients together that have similar therapeutic requirements. Therefore, the greatest determiner of differences in the quantity of resources consumed by different patients within the same unit of a progressive-care system – such as the AICU – will be the length of time that patients stay within the unit.

The study which examined the relationship between various demographic and clinical factors and AICU length of stay was therefore of fundamental importance in the demonstration of the feasibility of a model of RBH capable of improving control over patient scheduling and resource allocation. If a strong relationship could be shown to exist between these factors and length of stay, then this would demonstrate a capability to control admissions based on patients' expected length of stay – and, by implication, their resourcing requirements - in the intensive-care unit. Moreover, since the intensive-care units typically have the most diverse case and severity-mixes of any units, it would be reasonable to assume that if such projections could be made for the intensive-care units, then they could equally be made for other units also.

In measuring AICU lengths of stay, the study considered only a very blunt measure of chronicity that defined any length of stay in excess of 48 hours as chronic, and any length of stay less than 48 hours as non-chronic. It was considered that a finer calibration of chronicity was not required due to the very general objectives of the study.

The study hypothesised that patients who died within the intensive-care unit were also more likely to have longer lengths of stay. This hypothesis was confirmed, with only 3.3% of patients with a non-chronic length of stay dying within the AICU, compared with 20.6% of patients with a chronic length of stay dying within the AICU.

The study also examined the relationship between chronicity and operative category. It was found that there were substantial differences in the likelihood that a patient would require a chronic stay in the AICU depending on whether their therapy involved a surgical procedure or not, and if so whether the procedure was cardiac, thoracic, vascular or general surgery. In particular, it was found that those patients whose therapy did not involve a surgical procedure were almost 5 times more likely to require a chronic length of stay in the AICU than those patients whose therapy involved a cardiac surgical procedure.

Finally, the study examined the relationship between AICU chronicity and the source unit from which patients were admitted to the AICU. Perhaps unsurprisingly, patients who were admitted to the AICU from the operating theatres had the lowest risk of a chronic AICU length of stay, whereas patients from the wards had the highest risk. This difference most probably reflects the fact that many patients originating from wards will either be re-admissions to the AICU following complications, or fall within the group of non-surgical patients which have a higher risk of a chronic AICU length of stay.

All of these findings are important for the development of CAPSS. Not only does it provide a descriptive statistical picture of the patients and their treatment, but more importantly it suggests that any model which categorises patients according to their expected length of stay in the intensive-care units and other component units of RBH allows managers to increase their control over the patient scheduling process.

Demonstrating a relationship between patient characteristics and length of stay is necessary for the development of a model of patient scheduling, but it is not in itself sufficient, for it assumes that length of stay is a viable measure of consumed resources. This assumption can be recast in terms of bed slots – where it is assumed that one unit of time in a healthcare unit per patient (or, in those cases where some resources are left idle, per package of healthcare resource which would otherwise be used in the treatment of a patient) is equivalent to a bed slot. The notion of a bed slot is therefore defined as a package of healthcare resource per unit time.

According to the classification of resources developed in Chapter 2, the notion of a bed slot can be restricted to considering only those resources relevant to any model of patient scheduling where the explicit objective of that model is to improve cost effectiveness through enhancing operational control.

Given this restricted conception of the bed slot, the question remains whether or not the bed slot is a viable measure of resources to be used in a model of healthcare patient scheduling. If it is not a viable measure, then the enterprise of developing a model of resource allocation and patients scheduling is stopped in its tracks given the enormous simplification of the problem domain that the use of the bed slot concept affords.

In an attempt to answer the question of whether or not the bed slot concept is viable in the case of RBH, a study was devised using data from the AICU. This study based itself on four different hypotheses which together imply the viability of the bed slot concept. The objective of the study was therefore to confirm each hypothesis and thereby the viability of the bed slot. These four hypotheses were as follows:

1. *The Predictability Hypothesis*. Individual non-generic resources or different types of bed-slot whose individual allocation sizes are not predictable in advance of consumption within an acceptable degree of accuracy for different types of patient should be modelled as generic resources, either individually or as components of a bed-slot type in the case of individual non-generic resources, or in the case of types of bed-slot as instances of a more generic type of bed-slot; and

2. *The Non-Generic and Fixed Resources (–G&F) Proportion Hypothesis*. The –G&F resources whose individual allocation sizes are predictable in advance of consumption within an acceptable degree of accuracy do not constitute an excessively large proportion of the total resource consumption; and

3. *The Variance Hypothesis*. The overall level at which those resources which are modelled as components of the bed-slot are consumed at the population level is within an acceptable degree of variance within different categories of patients, or that such variance may not be reduced through the categorisation of patients into different categories where the categorisation of patients into those categories may be made in advance of consumption within an acceptable degree of accuracy; and

4. *The Difference Hypothesis*. There is no significant difference in the overall level of consumption of those resources which are modelled as components of the bed-slot between different categories of patients, or that where there is a difference between different categories, the categorisation of patients into those categories may be made in advance of consumption within an acceptable degree of accuracy.

The first of these hypotheses was proved formally, although it was noted that the utility of categorising packages of resources into bed slots could only be determined if two empirical questions were answered, namely, in predicting the bed slot requirements of patients, what is the cost (measured either economically or clinically) of making a mistake in the prediction, and further, is there any extra utility to be gained by making a further classification of existing bed slots?

The other hypotheses were tested empirically by using the Therapeutic Intervention Scoring System, TISS, which measures the amount of resource utilisation of intensive-care patients by measuring the amount of resources consumed for a variety of different resources.

The confirmation of the second hypothesis consisted of categorising each component resource as measured by the TISS score into whether it represented a –G&F – that is, a non-generic but fixed resource – or not. It was concluded on an intuitive level that the proportion of –G&F resources which were predictable in advance of consumption was not sufficiently high to disconfirm the hypothesis.

In the case of the variance and difference hypotheses, patient's resource consumption data were correlated with the outcome of regression-based prediction models. An outcome variable that measured the patient's resource consumption data was constructed from the data by using a cluster analysis of TISS scores. Two clusters were derived – one representing high-resource patients; the other low-resource patients. The hypotheses were therefore considered as being confirmed if either these two TISS clusters were not significantly separable in terms of either mean TISS score or variance in TISS scores, or that if they were that patients could accurately be predicted to belong to either cluster.

The hypotheses were not confirmed on an intuitive level when only pre-admission data was considered, with a large degree of separability between the two clusters, and an accuracy rate of 62% for bed slot classification. However, a new model of prediction was introduced which combined pre-admission data with the previous day's TISS score. This rolling system of prediction resulted in predictive accuracy of 80% which on an intuitive level was considered as confirmation of both hypotheses, although of course it also indicates the viability of creating two different types of bed slots within the AICU.

This study is the first of its kind and may be considered as a very useful tool in evaluating the level of targeting of healthcare resources. Its applicability to other types of healthcare units other than intensive-care units may be assured by the typically large degree of variance and range of therapeutic requirements found in the resource consumption profiles of intensive-care units.

The final study examined patterns of resource utilisation in the operating theatres. The purpose of this study had two main objectives. The first was to provide a set of statistically descriptive data for the operating theatres to inform the modelling process; the second was to evaluate the level of control over allocation and scheduling of operating theatre resources and to identify certain factors in the management of operating theatre resources that resulted in reduced level of operating efficiency.

There were two data sets used in this study. Both sets of data comprised records of patients scheduled for admission to the operating theatres at RBH over the same period of time. The difference between the two sets of data lay in the epistemological status of the resource allocation decisions resulting from each record. In the first set of data, each record corresponded to the intended schedule of admissions of patients to the operating theatres. That is, the schedule as it is first derived by the admissions manager. The second set of data corresponded to the actual schedule of admission of patients to the operating theatres. That is, the schedule of patients as it actually occurred after the inclusion of urgent cases and the subsequent modification of the original schedule of admissions that such cases imply.

From each set of data various performance and workload measures were derived by counting the number of admissions and in some cases weighting each patient according to the projected workload that their surgical procedure implied, not only for the operating theatres, but also for the recovery room and intensive-care units.

Because the first data set – the so-called non-monotonic data set, since the resource allocation decisions implied by the data were non-monotonic in nature – represented the optimal operating scenario given certain control limiting factors which were beyond the control of the admissions manager, a comparison of this data set with the second data set – the so-called monotonic data set, since the resource allocation decisions implied by the data were monotonic in nature – should indicate weaknesses in the allocation of operating theatre resources. Once such weaknesses are identified, this not only informs the subsequent development of a model of patient scheduling, but should also provide a benchmark for evaluating the performance of such a model in improving operating efficiency.

The two data sets were compared first for significant differences in mean values as well as variances of various performance and workload measures for the operating theatres, adult and paediatric intensive care units and the recovery room.

In the comparison of means it was found that there were fewer numbers of fast-tracked patients per day, fewer AICU bed slots allocated per day and less operating theatre workload per procedure in the monotonic allocation of resources than the non-monotonic allocation.

In the comparison of variances it was found that there was greater variance in operating workload per procedure and operating workload per day in the monotonic case than in the non-monotonic. Surprisingly, however, there was found to be significantly less variance in recovery room workload per procedure in the monotonic than the non-monotonic allocation of resources.

Apart from the result for the variance in recovery room workload per procedure, the results indicate that the actual allocation of resources shows evidence of less control than the non-monotonic. This result is unsurprising given that the non-monotonic allocation does not include urgent cases and other required changes to the schedule of patient admissions and allocation of resources. Unsurprising as this result is, however, it is of use not only in quantifying the level of control over resource allocation, but also in providing a benchmark measure of performance in evaluating a model of patient scheduling.

The second comparison of the two data sets had the objective of identifying the causes of suboptimal control over patient scheduling. It was assumed that the optimal operating scenario would be as described in Chapter 2, with 100% utilisation rates of all resources and no cancellations or delayed admissions for urgent cases. Naturally, this assumption is unrealistic, but it nonetheless gives the control scenario for any subsequent study comparing the two data sets.

A distinction was made in the study between what was called epistemological and non-epistemological control limiting factors. The former prevented the fulfilment of the optimal operating scenario through a lack of knowledge of the future supply and demand of healthcare resources; the second through system constraints implicit in the design of the system and beyond the control of the admissions manager. An example of the former type of control limiting factor would include those cases where a patient's length of stay in the intensive-care unit is longer than expected. An example of the latter type of control limiting factor would be characteristics of the nursing schedule which means that at certain times of the day or at certain days of the week there are fewer nurses available to treat patients than at other times of the week.

A method was devised which was able to identify and classify control limiting factors according to whether they were epistemological or non-epistemological based on a comparison of the longitudinal variances in performance and workload measures within each data set. This resulted in four different possibilities:

1. There was significant longitudinal variance in the non-monotonic data set for a variable V but not in the monotonic data set;
2. There was no significant longitudinal variance in the non-monotonic data set for a variable V but there was in the monotonic data set;
3. There was significant longitudinal variance in the non-monotonic data set for a variable V and in the monotonic data set;
4. There was no significant longitudinal variance in the non-monotonic data set for a variable V nor in the monotonic data set.

Based on these four possible outcomes, a set of Boolean expressions was derived to identify, for each variable, whether the results indicated the presence of either an epistemological or non-epistemological control limiting factor or not. The assumption from which these expressions were based was that non-epistemological control-limiting factors would be considered by the admissions manager in the construction of the admissions schedule. For example, patients would be scheduled for admission to the operating theatre according to the availability of surgeons and their surgical teams. In this case, the non-epistemological artefact of surgeon availability would be reflected in variance in patient admission rates to the operating theatre that could not be accounted for by random variation.

Various measures of unit performance were considered in the study, and for the identification of classification of control limiting factors, longitudinal variance was considered for both on a monthly basis (i.e. the independent variable being the month of the study period), as well as per day of the week (i.e. the independent variable being the day of the week, Monday thru Friday).

The results indicated that there were present various epistemological and non-epistemological control-limiting factors in the allocation of resources and the scheduling of patient admission at RBH when considered longitudinally by month, as well as by day of the week.

The result that there should be non-epistemological artefacts when considering comparisons between different days of the week is unsurprising, since the availability of specialist surgical teams is not constant throughout the week, and there is often a deliberate effort to reduce workload towards the end of the week to release nursing resources over the weekend period.

The results also indicated, however, the presence of epistemological artefacts for both the monthly and day-of-the-week perspectives. Of particular note is the presence of epistemological artefacts by month for the fast-track related variables of the proportion of patients classified as fast-track and the total number of fast-track patients per day. From the day-of-the-week perspective, there was an epistemological control limiting factor indicated for the variable measuring the number of PICU (Paediatric Intensive-Care Unit) bed slots allocated per day, suggesting a particular lack of control over paediatric admissions throughout the week.

Apart from affording an insight into the operational strengths and weaknesses of the control system in place to allocate resources and schedule patients at RBH, the method of comparing the non-monotonic and monotonic patient admission schedules also represents a new way of looking at control systems and their evaluation.

Traditionally, the study of control systems has been limited to Newtonian-style conceptions of controllers and controlled systems. Still, many of the control systems which most directly affect our quality of life are human-based systems which permit no analysis by any known form of equation – at least not to the extent necessary for precision engineering or re-engineering. It thus remains of great importance to develop methods that allow for the evaluation of such control systems and the identification of control-limiting factors. For, in the absence of any formal and precise description of the system and the controller, there is no other means of systematically redesigning the system and/or its control mechanism. The statistical method presented in the study is perhaps one possible way of achieving this. Moreover, it has a wider arena of application than healthcare alone, with the problem of cost-effectively scheduling resources between different processor units being a common problem in many other human-based systems, both at the micro scale of individual businesses and organisations, and at the macro scale of economies.

7.3. Modelling Approach and Formalism

The results of the various studies indicate various weaknesses in the process of patient scheduling at RBH. These results, in conjunction with the model of resource allocation developed in Chapter 2, indicate the need for a better control system to manage the patient scheduling. To be able to develop this control system, it was argued in Chapter 4 that a modelling formalism needs to be derived which is able to be used in both the process of redesigning the healthcare system to optimise its performance and to allow for the introduction of the enhanced control system, as well as to redesign the control system itself.

The approach that was proposed was interpreted as a hybrid between modelling formalisms used in business process re-engineering and formalisms used in software engineering. This hybrid approach was formulated on the basis of a five-staged development process, represented below as Figure 7.03

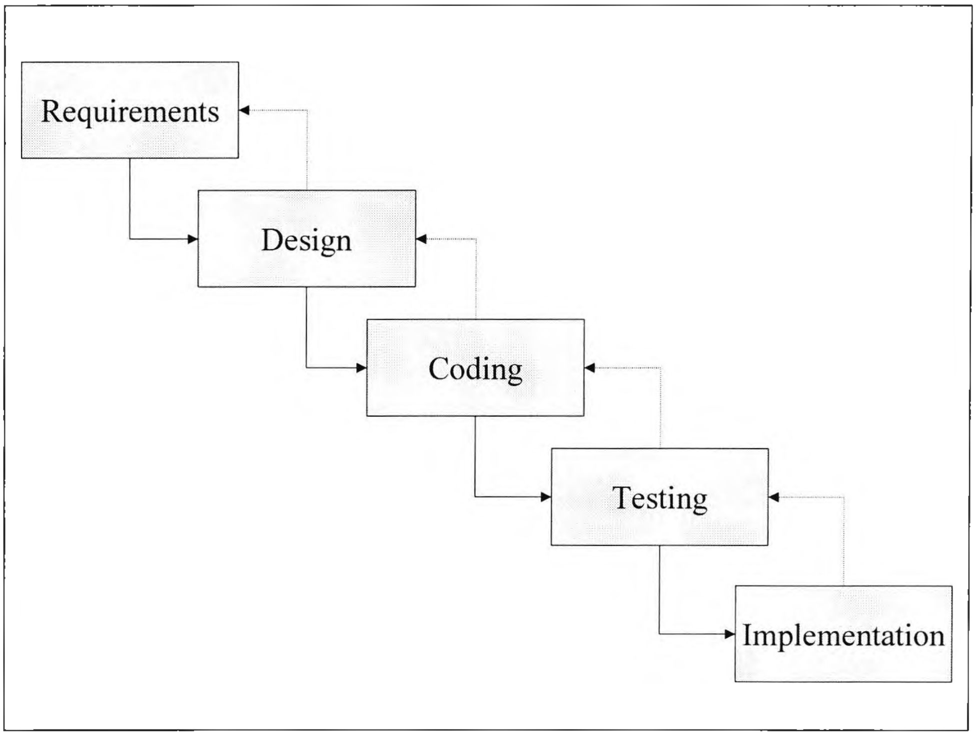

Figure 7.01. The software engineering process.

Conventional software engineering formalisms were not considered able to be used in all of these developmental phases because of their restricted scope – they consider only the software components of the system, and ignore other system components, by, for example, making the distinction between user and software. These kinds of distinctions and scope limitations are problematic when it is the entire system which is being redesigned, and certainly ignores the thesis that if any piece of software is to be anything more than a mere replication of the processes currently performed manually, then the relationship between user and software and between the software and other components of the system, and possibly even the architecture and capabilities of the system itself, will inevitably change also.

Modelling formalisms used in business process re-engineering (BPR) were equally inappropriate given their lack of formal rigour as well as their tendency to ignore issues of implementation and hardware. There was also the problem that BPR often leaves issues of structuring and representing the system's data resources and processes unaddressed, meaning that many of the complexities of specifying informational processes in the system and any corresponding data repositories are unable to be represented or managed effectively.

The hybrid approach to modelling the RBH high-dependency environment takes the strengths of both the BRP and software engineering approaches by being a systemic modelling approach, including both users and software and other system components, as well as including representations for the hardware implementations of different processes and having a comprehensive representation for system data resources and informational processes.

The hybrid approach's modelling formalism was based on Petri-nets, where system processes are represented by transitions, and the input/output data from those processes represented by places. The system was represented using the object-oriented modelling paradigm rather than the function-oriented paradigm.

The choice of using object-orientation was based on a information-theoretic comparison of object-oriented versus function-oriented system decompositions. To make this comparison, it was assumed that the hybrid modelling approach needed to choose between function-orientation and object-orientation on the basis of the ability to represent and model the RBH high-dependency environment in a manner which is easily comprehensible and manageable. That is, issues of resulting code efficiency of any software deliverables were not considered, along with other conventional software metrics. The result of the comparison indicated that object-orientation as a modelling paradigm was more efficient in representing systems beyond a minimal level of complexity.

Apart from grouping data into object classes according to the object-oriented paradigm, a grouping of processes was also introduced akin to the notion of a use case in the Unified Modelling Language (UML). As with UML, these process groups represented related processes which conjointly satisfied a high-level function of the control system. The resulting data flow was represented using the Petri-net formalism where each processing thread was considered as a loop within the Petri-net.

One important aspect of the modelling approach adopted was to include the implementations of the system processes within the modelling formalism. This resulted in two distinct perspectives on process groups – the object perspective and the processor view. The object view represented the data flow of the process group between the different component object classes of the system; the processor view represented the data flow of the process group between the different component processors of the system.

Data repositories were represented as a subclass of processes and were thus incorporated into the modelling formalism as transitions within the Petri-net formalism.

The resulting hybrid modelling approach and its corresponding formalism represents a potential resolution of the conflicting objectives of BPR-style modelling formalisms and software engineering modelling formalism such as UML. On the one hand, it affords much greater capabilities in representing data structures and software requirements than BPR-style modelling formalisms, and yet unlike software engineering formalisms takes a systemic approach encompassing both users and hardware as integral components in the model derivation.

Using the modelling formalism it is possible to calculate simple but very useful system metrics from resulting models. These system metrics are able to measure important aspects of system architecture and performance, and can thus provide a useful tool in the re-engineering and software design exercise. This ability to derive system metrics may be considered akin to the derivation of software metrics from source code and serves a similar purpose. One system metric used in the modelling of the RBH high-dependency environment was the Object Class Fragmentation Ratio, and measured the degree to which the component data of object classes was distributed amongst different data repositories. Ultimately, it is metrics such as these which indicate the efficiency and effectiveness of any operational model.

7.4. Operational Models

According to the hybrid modelling approach, it is necessary to derive two models. The first of these represents the current system; the second the proposed system. The derivation of these models was the objective of Chapter 5.

In both the current (COP) and proposed (POP) operational models of the RBH high-dependency environment, there were three object classes: Unit, Patient and Bed-Slot. This may appear to be a relatively small number of object classes given the complexity of the problem domain, although it is worth noting that many considerations that would require an increase in the number of object classes would only come into play during the coding of any resultant software deliverable, and would comprise object classes subsidiary to the main three above, such as those to support lookup functionality and other systems of data grouping.

The processes were grouped as follows:

1. Create Policy: The function of allocating resources between the different units of the RBH high-dependency environment and the derivation of admissions and treatment policies.

2. Schedule Patient: The function of allocating bed-slots within the different units of the RBH high-dependency environment and updating the relevant databases as appropriate.

3. Treat Patient: The function of administering treatment to patients and updating clinical databases as appropriate.

4. Transfer Patient: The function of discharging patients or transferring them to another unit within the RBH high-dependency environment.

These four process groups occurred in both the COP and POP models of the RBH high-dependency environment, as did the three object classes. In the process groups of both the current and proposed operational models, the processes were categorised into two types, depending on whether the process involved a read or write message to or from a data repository or whether the process involved the manipulation or updating of the data being read from or written to a data repository. The former of these process types were very different in number from the proposed operational model than in the current operational model.

In the current operational model there were many more data repository read or write processes (38 compared with 32 in the proposed operational model). This reduction in the number of processes was achieved through the replacement of many data repositories with a computerised database, the Computer-assisted patient scheduling system Database (CAPSS Database). In the current operational model, there were many distinct and paper-based data repositories which caused many instances of data duplication and time-consuming manual reading and writing of data in those data repositories.

Each of the read or write database processes in each model were classified according to whether or not they were manual processes or automatic (i.e. computerised) processes, and whether or not they were read or write processes. The count of processes in each category is as follows:

Process Type	**RBH COP**	**RBH POP**
Manual Read	23	10
Automatic Read	0	6
Manual Write	14	10
Automatic Write	1	6
	38	**32**

Table 7.01. Comparison of different database process types between the RBH COP and RBH POP models

Thus, as can be seen from the above table, not only does the proposed operational model have a reduced number of data repository processes, but also represents a significant increase in the number of automated processes, and thereby an increased level of computerisation of the whole system.

This decrease in data repository processes and the resulting increase in system integration can be measured by the Object Class Fragmentation Ratio. The Object Class Fragmentation Ratio for COP was 0.20, compared with only 0.06 for POP.

Model	Objects	Relations	Databases	Ratio
RBH COP	3	7	7	0.06
RBH POP	3	5	3	0.20

Table 7.02. Comparison of Object Class Fragmentation Rates between the RBH COP and RBH POP models

Metrics such as the count of data repositories and the Object Class Fragmentation Ratio are good indicators of the efficiency of the system operation, and in particular the amount of resources needed to communicate information amongst the different processing units of the system and record that information in an appropriate format. Still, it does not imply any necessary change in the effectiveness of control over system operation – all that such metrics represent is a duplication of existing system processes using different media and communication systems and different types of data repositories, such as computer databases rather than paper records.

To imply any change in the effectiveness of operational control, the system itself needs to change in terms of its component processes and the flow of work and information amongst those processes.

7.5. Design Models

All of the processes in the current operational model had corresponding processes in the proposed operational model. However, there were two additional processes in the proposed operational model that did have equivalents in the current operational model. These two additional processes were:

- **POP2.07: predictResourceConsumption(Patient)**: POP2.07: PredictResource Consumption(Patient) is a component process of the Schedule Patient process group. It makes projections of the resource requirements of the patient prior to admission to the high-dependency environment, including to which units the patient will require admission, and when. POP2.07: predictResourceConsumption(Patient) is a component process of the Patient object class and is performed by the processor CAPSS.

- **POP2.08: createScheduleEvaluation(Unit)**: POP2.08: createScheduleEvaluation (Unit) is a component process of the Schedule Patient process group. It evaluates the proposed admission schedule by generating economic and clinical performance statistics of each unit that are projected to result from admitting the patients according to the proposed schedule. POP2.08: createScheduleEvaluation(Unit) is a component process of the Unit object class and is performed by the processor CAPSS.

These two processes represent a paradigm shift in the way patient scheduling is managed and controlled in a high-dependency environment such as RBH. The process POP2.07: predictResourceConsumption(Patient) takes patient data as input and generates predictions such as length of stay and urgency based on that data. The process POP2.08: createScheduleEvaluation(Unit) takes the output of POP2.07: predictResourceConsumption(Patient) and uses it in combination with the data from staff rosters, surgeon availability schedules, and the proposed patient admission schedule derived by the admissions manager, to evaluate the proposed admission schedule by generating projected performance figures that would result given the combined factors of the proposed admission schedule, the availability of resources, and the current profile of patients already being treated within the high-dependency environment.

The incorporation of processes such as POP2.07: predictResourceConsumption(Patient) into healthcare delivery systems is not a new idea. For example, clinicians regularly use clinical scoring in making treatment, diagnosis and prognosis decisions. Such clinical scores represent prediction systems for various outcomes, such as probability of mortality, responsiveness to certain drugs, and so on. Amongst these types of scores, the probability of mortality scores are the most complex and most developed. However, their use, as with other types of scores, has been almost exclusively in prioritising treatment and clinical auditing and evaluation. Such scores have not been previously used in patient scheduling processes.

When integrated with P2.08: createScheduleEvaluation(Unit), P2.07: predictResourceConsumption(Patient) the result is a model-driven computerised decision-support system where .08: createScheduleEvaluation(Unit) provides information to scheduling managers in the evaluation of their admission schedules and staff rosters and their subsequent scheduling decisions. It is this use of prediction scores in combination with a scheduling evaluation process that represents a novel use of such scores in healthcare, reflecting their ability to be used directly in patient management and the improvement of operating cost-effectiveness.

The objective of Chapter 5 was to describe the processes P07: predictResourceConsumption(Patient) and P2.08: createScheduleEvaluation(Unit) only in terms of their relationships with other processes, object classes and processors and the data fields which define the processes themselves. The objective of Chapter 6 was to consider P07: predictResourceConsumption(Patient) and P2.08: createScheduleEvaluation(Unit) as actual design models which satisfy the data requirements specified in Chapter 5.

The feasibility of P07: predictResourceConsumption(Patient) as a computerised model was demonstrated with a literature review and meta-analysis of comparable prediction models. In all ten studies were included in the study, all of which considered the problem of either predicting patient's length of stay in an intensive care unit or the prediction of whether or not cardiovascular surgical patients can be recovered in the post-operative recovery room or require admission to the intensive-care unit. Most of the studies used various forms of linear regression to derive their prediction models, although an example of the use of an artificial neural network prediction model was also used. Both the regression models and the neural network models of patient resource consumption requirements may be implemented as a software routine.

The conclusion of the study is that amongst all of the prediction models, the artificial neural network generated the best results, although a direct comparison with other models could not be made due to differing outcome measures being used. However, the only study which compared the performance of the derived prediction model with that of human predictors suggested that the derived prediction model generated more accurate results.

Despite these conclusions, for P07: predictResourceConsumption(Patient) to be deployed in the context of a decision-support system such as the Computer-assisted patient scheduling system (CAPSS) proposed, the generated predictions of resource consumption need to be sufficiently accurate to be applied to individual patients, rather than patient populations. Moreover, a prediction model needs to consider the case of deviations from the initial prediction made pre-admission for P2.08: createScheduleEvaluation(Unit) to provide useful results. Fortunately, the bed-slot validation study of Chapter 4 showed that using a rolling system of prediction made every 24 hours can improve predictions and achieve accuracy of 80% as to whether the patient will be discharged within the subsequent 24 hours. Thus, while it is unlikely that this level of accuracy can be duplicated prospectively, this use of prediction models along with pre-admission predictions can make the use of a computerised implementation of P07: predictResourceConsumption(Patient) in combination with P2.08: createScheduleEvaluation(Unit), and thus of CAPSS as a whole, a feasible system.

With regards to P08: createScheduleEvaluation(Unit), Chapter 6 considered a simulation model using the formalism of coloured Petri-nets. The choice of using a simulation model was made based on the comparison of simulation models with markov and queuing theory based models. The choice of coloured Petri-nets was made due to their extensive use in comparable problem domains such as the design of manufacturing systems. Moreover, the use of the colour abbreviation of the standard Petri-net formalism to create coloured Petri-nets allows different patients with different resource requirements to be modelled as different colours.

In the Petri-net model, each component healthcare unit of the high-dependency environment was considered as a subnet. This allowed the component processes of each unit to be aggregated into a set of complex transitions in a system-level model of the high-dependency environment. Apart from providing a mechanism for managing model complexity, this also enabled patient flows between units to be represented more transparently and thus aiding in comprehension.

The relationship between the Petri-net model of P08: createScheduleEvaluation(Unit) and a software implementation – such as a suite of artificial neural networks, as suggested above – of P07: predictResourceConsumption(Patient) is the most innovative aspect of the proposed operational model of the RBH high-dependency environment of Chapter 5. With patients being represented by coloured tokens in the Petri-net simulation model of the high-dependency environment, the output of the prediction models may be used to specify the properties of each token colour. In particular, each token colour may be ascribed properties for:

1. The sequence of admissions for the patients represented by the tokens
2. The length of stay in each unit to which patients represented by the tokens are admitted.

Therefore, upon the creation of a proposed schedule by the admissions manager, the data for each patient on the admissions schedule, as well as data for all patients already admitted to the high-dependency environment, along with data regarding the availability of resources in the form of staff rosters and so on, the resource consumption requirements predictions may be generated by the software implementation of P07: predictResourceConsumption(Patient). These predictions are then used to populate the Petri-net model which then outputs an evaluation of the schedule in terms of its projected impact on the performance parameters of each component unit of the high-dependency environment. This output, which may be presented in the form of a report, can then be used by the admissions manager to decide whether to make any changes to the admissions schedule before distributing it to members of the clinical team to implement. In this sense, the proposed computer-assisted patient scheduling system involves a decision-support component whereby the process of patient scheduling involves a human as well a computer input into the decision-making process.

The use of decision-support systems in healthcare is not a new idea, but the use of a decision-support system that is based on clinical prediction models and a simulation model for use in patient scheduling and resource allocation is a novel use of the decision-support system paradigm.

7.6. Summary

The aim of this thesis has been to show how healthcare managers can increase the cost-effectiveness of healthcare delivery through the development of a computer-assisted patient scheduling system (CAPSS) that enables increased levels of control over the consumption of healthcare resources.

To achieve the aim of demonstrating how CAPSS allows healthcare managers to increase the cost-effectiveness of healthcare delivery the thesis was divided into five main chapters. Each chapter satisfied specific objectives as laid out in Chapter 1 and summarised as follows.

7.6.1. Theoretical Analysis

Chapter 2 provided a theoretical analysis of the process of patient scheduling. A mathematical model was developed showing that patient scheduling may be modelled as a control process. This was the first logical step in supporting the hypothesis that CAPSS may improve operational cost-effectiveness by increasing the level of control that operational managers have over the patient scheduling process.

It was argued that the extent to which operational managers can control the patient scheduling process may be measured by the variance in the distribution of different consumption rates of a healthcare facility's resources. This represented an important assumption in developing a method for evaluating system performance.

An equation was developed based on the distribution in consumption rates of healthcare resources that measured the cost-effectiveness of healthcare delivery. This equation provided the basis for a far more formal approach to healthcare facility performance evaluation than has been used thus far in the literature.

Demonstrating a positive correlation between the effectiveness of control and the operational cost-effectiveness of a healthcare facility was necessary in supporting the hypothesis that CAPSS can improve cost-effectiveness, with this improvement in cost-effectiveness being derived from an increase in the effectiveness of control over patient scheduling.

The difficulty in optimising the operational performance of a healthcare facility was demonstrated by equating the patient scheduling process to that of a tiling problem. This equivalence was presented as a conceptual aid in the modelling of the patient scheduling process.

7.6.2. Empirical Analysis

Chapter 3 provided an empirical analysis of the process of patient scheduling as it occurs in the Royal Brompton and Harefield NHS Trust (RBH) which was the empirical domain used throughout this thesis.

A fundamental assumption introduced Chapter 2 stated that under certain conditions healthcare resources may be modelled as a single resource, referred to as a the bed-slot. Chapter 3 included the conclusions of a study which validated this assumption as it applies to RBH, with the main results of the study included as Appendix 10.

The requirement that any system with the objective of increasing the level of control that managers have over patient scheduling must be to be able to predict the amount of resources that each patient is likely to consume was shown was tested using RBH data. The study showed that a rolling system of updating the predictions of resource consumption every 24 hours would satisfy the requirement. The full results of the study were included within Appendix 10.

Chapter 3 argued that the success of CAPSS in increasing the cost-effectiveness of healthcare delivery at RBH is dependent upon the current process of patient scheduling being sub-optimal. Further, that any sub-optimality must be caused by control-limiting factors that result from a lack of knowledge regarding patients' projected resource consumption or the projected availability of healthcare resources if the introduction of CAPSS is to support the hypothesis that it is able to improve cost-effectiveness. This is opposed to control-limiting factors that are known by managers and are inherent in the system design.

The distinction between those factors whose effects are known to managers and those that are not was defined as one between epistemological and non-epistemological control-limiting factors and a statistical method was developed to identify these two types of control-limiting factor and quantify their effects on cost-effectiveness. The full method and its application to RBH were included as Appendix 11.

Chapter 3 presented the conclusions of a study using the statistical method to identify different control-limiting factors applied to data from the RBH high-dependency environment. The study concluded that there were various control-limiting factors present in the patient scheduling process in operation at RBH. Moreover, the study showed that at least some of these control-limiting factors were epistemological in nature, and thus supported the hypothesis that the introduction of CAPSS would improve the level of control over patient scheduling, and thereby improve cost-effectiveness.

7.6.3. Modelling Approach and Formalism

Chapter 4 developed a modelling approach and formalism suitable for the development of operational models of RBH and the requirements of CAPSS. The utility of modelling approaches and formalisms adopted in Business Process Re-engineering (BPR) and software engineering were considered in providing a suitable basis for the modelling of RBH and the requirements of CAPSS. It was concluded however that neither formalism by itself was adequate for this task. This conclusion was made on the basis of considering various properties necessary for the modelling of RBH and the requirements of CAPSS and whether or not those properties are included in conventional BPR and software engineering approaches or formalisms.

As a result of the evaluation of modelling approaches and formalisms deployed for the activities of software engineering and business process re-engineering an argument was made for a hybrid modelling approach particularly suited to the development of operational models of the RBH patient scheduling system and other similar systems.

In evaluating the appropriate modelling formalism to adopt for the modelling of the RBH patient scheduling system the object-oriented modelling paradigm was compared against its function-oriented equivalent in terms of the efficiency with which each could represent and model systems. The comparison method used was based on information theory and the results indicated that the most efficient modelling paradigm to use was dependent on the degree of system complexity. In the case of modelling the RBH patient scheduling system, it was concluded that due to the level of system complexity an object-oriented approach was appropriate. The full method and results of the study were included as Appendix: Information-Theoretic Evaluation of Object-Oriented System Representations.

Chapter 4 combined the properties of a basic definition of Petri-nets with the static modelling properties of a generic object-oriented modelling formalism to create a comprehensive modelling formalism. It was argued that this formalism may be used in the modelling of RBH and the requirements of CAPSS.

In order to provide a formal comparison of the current process of patient scheduling at RBH and the proposed process with the introduction of CAPSS, Chapter 4 defined a set of system metrics that may be used to evaluate the processing efficiency and degree of system integration in both the current and proposed operational models of the patient scheduling process. These system metrics were used in the comparative evaluation of the operational models of the RBH system of patient scheduling developed in Chapter 5.

7.6.4. Operational Modelling

Chapter 5 was the first modelling chapter whose purpose was the development of two models to represent the operational processes and data involved in the process of patient scheduling at RBH. The first model developed was the current operational model (COP) representing the current patient scheduling process. The second model was the proposed operational model (POP), representing the patient scheduling process with the inclusion of CAPSS. The main points of the two models were as follows.

Three object-classes were defined in both the current and proposed operational models:

4. **Patient** contains all of the data and processes defining patients within the RBH high-dependency environment
5. **Bed-Slot** contains all of the data and processes defining bed-slots within the RBH high-dependency environment

6. **Unit** contains all of the data and processes defining the component healthcare units of the RBH high-dependency environment

The main difference between the current and proposed operational models was the introduction of two new processes into the proposed operational model that were absent from the current operational model. Both of these new processes were represented as computerised processes and as such represented the main software components of CAPSS. The two new processes defined the prediction of patients' resource consumption and the subsequent prediction of system performance.

In the proposed operational model predictions were made of each patient requiring admission to the RBH high-dependency environment. These predictions consisted of the patient's projected bed-slot consumption and which of the component units of the RBH high-dependency environment where the bed-slots would be consumed and when. The manager responsible for scheduling patient admission could then propose a schedule of admissions. This proposed schedule of admissions would be evaluated by the system by predicting the performance of each component unit of the RBH high-dependency environment that would result given each patient's predicted resource consumption and the availability of resources within the high-dependency environment. The operational manager would then be able to modify the schedule of admissions and repeat the evaluation procedure, or to implement the schedule unmodified.

Because there is no closed-loop involved in the process of patient scheduling in so far that the decision-making is still undertaken by a human agent, CAPSS represents a decision-support system. Unlike other decision-support systems deployed in healthcare settings, however, CAPSS is not concerned with the diagnosis or treatment processes. Rather, it is concerned with the optimisation of operational control that managers have over healthcare patient scheduling. Any benefits to be achieved by the introduction of CAPSS are therefore economic as much as they are clinical.

For CAPSS to improve control system performance it must not only result in sufficiently accurate predictions of patient resource consumption and system performance evaluation, but it must also accomplish these objectives in a cost-effective manner. For example, it must generate the information necessary to improve control system performance using as few resources as possible, while not compromising on the accuracy of this information. In practical terms this implies a need for increased levels of system automation.

One of the system metrics developed in Chapter 4 was designed specifically to measure the degree of system automation. This metric was applied to both the current and proposed operational models of the RBH patient scheduling system to compare the degree of system automation implied by each model. According to this metric, the proposed operational model had a much larger degree of automation than the current operational model.

While system automation measures the extent to which system processes are performed by non-human agents, system integration measures the extent to which system processes are distributed between a number of different processing agents – human or non-human.

A second metric developed in Chapter 4 was designed to measure an important aspect of system integration by measuring the degree of data distribution between different data stores. Applying this metric to the models developed in Chapter 5 showed that the proposed operational model had a much greater degree of system integration than the current operational model.

In summary, the proposed operational model introduced two new computerised processes into the RBH model of healthcare patient scheduling, one of which was designed to predict individual patients' resource consumption profiles; the other the projected performance of the component units of the RBH high-dependency environment. Together, these two new processes represent the core functionality of a computer-assisted patient scheduling system (CAPSS). Comparing the current and proposed operational models in terms of system metrics demonstrated that CAPSS is capable of supporting the hypothesis that computerising the process of patient scheduling allows healthcare managers greater control and thereby greater potential to increase the cost-effectiveness of healthcare delivery.

7.6.5. Design Modelling

Chapter 6 was the second modelling chapter whose purpose was the development of two models proposed as implementations of the two new processes introduced in the proposed operational model of the RBH patient scheduling process of Chapter 5. The first of these models was for the prediction of patient resource consumption requirements; the second was for the evaluation of system performance based on the predictions generated by the first model. The main points of the two models were as follows.

The purpose of the first model was to demonstrate that a suite of computerised prediction models is capable of predicting patients' resource consumption – within an acceptable degree of accuracy – for patients scheduled for admission to the RBH high-dependency environment.

The method used was to evaluate and compare prediction models that have been developed and described in the literature through a systematic literature review. Both the methods used in the derivation of each model, as well as the accuracy of the predictions made by each model, were considered in evaluating and comparing the models.

If any of the models were capable of demonstrating a sufficient level of accuracy in its predictions, then it was assumed to have the capability of being implemented as a software routine in the context of CAPSS for the prediction of patients' resource consumption requirements and the integration with a model for the evaluation of system performance.

The objective of the second model was to evaluate the operational performance of healthcare systems such as the RBH high-dependency environment. Various methods were compared for the design of the model, based on the methods that have been successfully used for comparable models and have been described in the literature. Although none of the models reviewed from the literature had precisely the same objectives as the model to be developed, a set of evaluation parameters were extrapolated on the basis of the requirements of the model to be developed, and applied to the models that have been presented in the literature. On the basis of this evaluation, it was concluded that the model would be best developed using a simulation-based approach. For the modelling formalism to be used for the design of the simulation model, it was argued that an abbreviation of the basic Petri-net formalism used would be appropriate. According to this formalism, different resource and patient types were categorised according to colours of tokens used to simulate the system dynamics. This conclusion was based on the successful use of coloured Petri-nets in other comparable modelling domains

7.7. Contributions to Knowledge

This thesis has made contributions in various fields of research, including healthcare management, system modelling and simulation, critical care medicine, business process modelling and software engineering.

1. **Healthcare Management**. The use of a business simulation model in tandem with a suite of prediction models as proposed in Chapter 5 and Chapter 6 is a novel use of both clinical prediction models, as well as of coloured Petri-nets. Typically, prediction models have only been used for the normalisation of different patient groups to make valid comparisons between different healthcare facilities, as well as in the use of aiding in treatment decisions. However, they have yet to be used in any extensive way to enhance the operational performance of healthcare units.

2. **System Modelling and Simulation**. Coloured Petri-nets have been used extensively in the simulation modelling of systems. However, their use has typically been in non-human systems such as flexible manufacturing systems and assembly lines. The use of coloured Petri-nets to model healthcare systems represents a transfer of a modelling formalism to a new problem domain.

3. **Critical Care Medicine**. The use of outcome prediction models has been a mainstay of research into critical care medicine for many years. However, the methods used in the development of such models, as well as the breadth of their application, have to date been very limited. This thesis has shown the viability of not only using such models in other domains, but also provided scope for the development of new types of models as well as new methods for the development of such models.

4. **Business Process Modelling**. To date the field of business process modelling has suffered from a lack of rigour and a dearth of advocates non schooled in the principles of system engineering and design. This thesis has developed both an overarching control-theoretic approach to business process modelling, as well as a well-suited formalism that is capable of serving the dual needs of demonstrating how business processing may be changed given the introduction of new technologies or new processes, and how this would reflect on the specific requirements of new software and hardware components of the proposed design. Moreover, objective system metrics were derived that measured important characteristics of particular system designs, meaning that an objective comparison may be made between current and proposed system designs prior to implementation.

5. **Software Engineering**. The modelling approach and formalism derived in this thesis is an advance on current modelling approaches and formalisms in that it takes a much more

systemic and panoramic perspective on the software engineering process. More systemic since it involves the modelling of all system components, including users and processors, rather than either assuming such components are mere interfaces with the system and the outside world, or simply ignoring them completely; more panoramic since in including all system components and how they interact, the models developed are able to be used at every stage of the software engineering process.

7.8. Future Research

The CAPSS models presented in Chapter 5 and Chapter 6 are deficient in one important respect: they have not been implemented. The hypothesis that CAPSS is capable of improving the cost-effectiveness of RBH through enhancing the level of control that operational managers have over the process of patient scheduling is supported from the mathematical models developed in Chapter 2 and the empirical studies summarised in Chapter 3. However, it would be unreasonable to make the claim that the hypothesis is confirmed without the models of Chapter 5 and Chapter 6 being implemented as a software application. This confirmation thus represents the next logical step in the validation of CAPSS, and thereby the development of a new breed of medical information systems.

8. References

[AUD95] Audry, N; Prunet, F (1995), *Controlled batches Petri nets*. IEEE Systems, Man & Cybernetics Soc. Conf. Proc. 2 pp.1849-1854

[BAR96a] Barie, P S; Hydo, L J (1996), *Influence of multiple organ dysfunction syndrome on duration of critical illness and hospitalization*. Arch. Surg. 131 pp.1318-1324

[BAR96b] Barie, P S; Hydo, L J; Fischer, E (1996), *Utility of illness severity scoring for prediction of prolonged surgical critical care*. J. Trauma 40(4) pp.513-519

[BEC95] Becker, RB, Zimmerman, JE, Knaus, WA, Wagner, DP, Seneff, MG, Draper, EA, Higgins, TL, Estafanous, FG, Loop, FD (1995), *The use of APACHE III to evaluate ICU length of stay, resource use, and mortality after coronary artery by-pass surgery*. J. Cardiovasc Surg 36(1)

[BOO91] Booch, G (1991), *Object-Oriented Design*. Benjamin-Cummings

[BRE91] Breslawski, S, Hamilton, D (1991), *Operating room scheduling*. AORN Journal 53(5) pp.1229-1237

[BRU86] Bruno, G; Marchetto, G (1986), *Process-translatable Petri Nets for the rapid prototyping of process control systems*. IEEE Trans., Software Engineering SE-12(2) pp.346-357

[BUC94] Buchman, T G; Kubos, K L; Seidler, A J et al (1994), *A comparison of statistical and connectionist models for the prediction of chronicity in a surgical intensive care unit*. Critical Care Medicine 22(5) pp.750-762

[BUT89] Butler, P; Duce, G; Morgan, C J (1989), *National Heart and Chest Hospitals Phase I: Monitoring and patient information management strategy*. Internal Report.

[CAR95] Carevue (1995), *HP Carevue Users International Conference 1995: Summary of Proceedings*.

[COS92] Cossins, R; Ferreira, P (1992), *Celeritas: a coloured Petri net approach to simulation and control of flexible manufacturing systems*. Int. J. Prod. Research 30 pp.1925-1956

[CRA95] Cramp, D G; Carson, E R (1995), Assessing health policy strategies: A model-based approach to decision support. Unpublished Manuscript N/A pp.N/A

[CUL74] Cullen, DJ; Civetta, J M; Briggs, B A (1974), Therapeutic Intervention Scoring System: A method for quantitative comparison of patient care. Critical Care Medicine 2(2) pp.57-60

[CUL94] Cullen, DJ, Nemeskal, AR, Zaslavsky, AM (1994), Intermediate TISS: a new therapeutic intervention scoring system for non-ICU patients. Critical Care Medicine 22(9) pp.1406-1411

[DAS91] Das, S; Sarkar, D; et al (1991), Extended colored Petri net: An efficient tool for analyzing concurrent systems. Information Sciences 54 pp.191-218

[DAV94] David, R (1994), Petri Nets for the Modeling of Dynamic Systems - A Survey. Automatica 30(2) pp.

[DER93] Der Jeng, M (1993), A Review of Synthesis Techniques for Petri Nets with Applications to Automated Manufacturing Systems. IEEE Trans, Systems, Man, Cybernetics 23(1) pp.

[DIT95] Dittrich, G (1995), Modeling of complex systems using hierarchically represented Petri nets. IEEE Systems, Man & Cybernetics Soc. Conf. Proc. 3 pp.2695-2699

[DOH01] Department of Health (2001), Building the information core implementing the NHS Plan. London: Department of Health.

[DOH02] Department of Health (2002), Delivering IT in the NHS a summary of the national programme for IT in the NHS. London: Department of Health.

[DOH98] Department of Health (1998), Hospital waiting list statistics England (resident based) quarter ended 31 December 1998. London: Department of Health

[DVI90] Davis, A (1990), Software Requirements. Prentice-Hall: Englewood Cliffs, NJ.

[EMB92] Embley, W D (1992), Object-Oriented Systems Analysis. PRENTICE-HALL (Englewood-Cliffs, NJ).

[ESO76] Esogbue, A O; Singh, A J (1976), A stochastic model for optimal priority bed distribution problem in a hospital ward. Operations Res. 24 pp.884-889

[EZP97] Ezpeleta, J; Colom, J M (1997), Automatic synthesis of colored Petri nets for the control of FMS. IEEE Trans, Robotics and Automation 13(3) pp.327-337

[FRE91] Freedman, P (1991), Time, Petri Nets, and Robotics. IEEE Trans, Robotics and Automation 7(4).

[GRA93] Graham, A (1993), *Object-Oriented Methods (2nd ed.)*. Addison-Wesley.

[HAM94] Hamilton, D, Breslawski, S (1994), *Operating room scheduling: factors to consider*. AORN Journal 59(3) pp.665-680

[HAN92] Hancock, W M, Isken, M W. (1992), *Patient-scheduling methodologies*. J. Soc. Health Systems 3(4) pp.83-94

[HAS93] Hasibeder W; Haisjackl M; Germann R; Vergeiner R; Seyr M; Sparr H (1993), Physiological scoring systems and audit. Lancet 342(8866) pp.307

[HER81] Hershey, J C; Weiss, E N; Cohen, M A (1981), A stochastic network model with application to hospital facilities. Operations Res. 29 pp.1-22

[HIL89] Hillion, H P (1989), Timed Petri nets and application to multi-stage production systems. in Rozenberg, G (ed.), Advances in Petri Nets, Berlin: Springer-Verlag N/A pp.281-305

[HMM90] Hammer, M (1990), Reengineering Work: Don't Automate, Obliterate. Harvard Business Review 7 pp.104-112

[HOL89] Holliday, M A (1989), Some issues in timed Petri nets. The Third Intl. Wkshp on Petri Nets and Performance Models: Pre-Workshop Tutorials N/A pp.117-150

[HOL93] Holt, A W; Bersten, A D; Worthley, L I; Vedig, A E (1993), Physiological scoring systems and audit. Lancet 342(8866) pp.307-308

[HOR95] Horvath, L; Rudas, I J (1995), Modeling of manufacturing processes using object-oriented extended Petri nets and advanced knowledge representations. IEEE Systems, Man & Cybernetics Soc. Conf. Proc. 3 pp.2576-2581

[JEN88] Jensen, K (1988), Coloured Petri-Nets: A way to describe and analyze real-world systems - without drowning un unnecessary detail. Meta Software Corporation Design/Simulator Project Papers N/A pp.N/A

[JEN89] Jensen, K (1989), Introduction to coloured Petri-nets. The Third Intl. Wkshp on Petri Nets and Performance Models: Pre-Workshop Tutorials N/A pp.77-116

[KAO72] Kao, E P C (1972), A semi-Markov model to predict recovery progress of coronary patients. Health Serv. Res. 7 pp.191-208

[KAO73] Kao, E P C (1973), A semi-Markovian population model with application to hospital planning. IEEE Trans. System Man Cybernetics 3 pp.327-336

[KAO74] Kao, E P C (1974), Modelling the movement of coronary patients within a hospital by semi-Markov processes. Operations Res. 22 pp.683-699

[KAT88] Katz, N M; Ahmed, S W; Clark, B K et al (1988), Predictors of length of hospitalization after cardiac surgery. Ann. Thorac. Surg. 45 pp.656-660

[KEE83] Keene, AR, Cullen, DJ (1983), Therapeutic intervention scoring system (TISS): update 1983. Critical Care Medicine 11(1) pp.35490

[KLE90] Klem, S A; Pollack, M M; Getson, P R (1990), Cost, resource utilisation, and severity of illness in intensive care. J. Paediatrics 116(2) pp.231-237

[KNA81] Knaus, WA, Zimmerman, JE, Wagner, DP, Draper, EA, Lawrence, DE (1981), APACHE - acute physiology and chronic health evaluation: a physiologically based classification system. Critical Care Medicine 9(8) pp.591-597

[KNA85a] Knaus, WA, Draper, EA, Wagner, DP, Zimmerman, JE (1985), APACHE II: a severity of disease classification system. Critical Care Medicine 13(10) pp.818-829

[KNA91] Knaus, WA, Wagner, DP, Draper, EA, Zimmerman, JE, Bergner, M, Bastos, PG, Sirio, CA, Murphy, DJ, Lotring, E, Damiano, A, Harrell, FE, Jr (1991), The APACHE III prognostic system: risk prediction of hospital mortality for critically ill hospitalized adults. Chest 100(6) pp.1619-1635

[KNA93] Knaus, W A; Wagner, D P; Zimmerman, J E; Draper, E A (1993), Variations in mortality and length of stay in intensive care units. Annals Internal Med. 118(10) pp.753-761

[LAM95] Lamiell, J M (1995), Modeling intensive care unit census. Military Medicine 160 pp.227-232

[LEM94b] Lemshow, S, Le, Gall, J-R (1994), Modeling the severity of illness of ICU patients: a systems update. JAMA 272(13) pp.1049-1055

[LEP94] Lee, P J, Chen, D J, Chung, C G (1994), An object-oriented modelling approach to system software design. Information and Software Technology 36(11) pp.683-694

[LIN98] Lin, S; Huang, H (1998), Modeling and emulation of a furnace in IC Fab based on colored-timed Petri nets. IEEE Trans. Semiconductor Manufacturing 11 pp.410-420

[LIT88] Litt R (1988), Therapeutic intervention scoring. Critical Care Medicine 8(6) pp.79-83

[LOP95] Lopez-Mellado, E (1995), Simulation of timed Petri net models. IEEE Systems, Man & Cybernetics Soc. Conf. Proc. 3 pp.2270-2274

[LOS94] Losavio, F, Matteo, A, & Schlienger, F (1994), Object-oriented methodologies of Coad Yourdon and Booch: comparision and graphical notation. Information and Software Technology 36(8) pp.503-514

[MAK94] Makhlouf, M; James, J H; Wigfield, E C; Wheeler, T M; Modeen, R A; Schaffner, S C (1994), A system engineering approach for the migration of computer-based systems. The MITRE Corporation N/A pp.N/A

[MAL92] Malstam, J; Lind, L (1992), Therapeutic intervention scoring system (TISS)--a method for measuring workload and calculating costs in the ICU. Acta Anaesthesiol Scand 36(8) pp.758-763

[MAR95] Marshall, J C; Sibbald, W J; Cook, D A et al (1995), The multiple organ dysfunction (MOD) score: A reliable indicator of a complex clinical outcome. Critical Care Medicine 23(10) pp.1638-1652

[MET95] Metnitz, P; Lenz, K (1995), Patient data management systems in intensive care - the situation in Europe. Intensive Care Medicine 21 pp.

[MIR91] Miranda, D R (1991), Management of resources in intensive care. Intensive Care Med 17(2) pp.127-128

[MOO95] Moore, K; Gupta, S (1995), Stochastic colored Petri net models of flexible manufacturaing systems: Material handling systems and machining. Computers in Industrial Engineering 29 pp.333-337

[MOR97] Moreno R; Morais P (1997), Validation of the simplified therapeutic intervention scoring system on an independent database. Intensive Care Medicine 23(6) pp.640-644

[MOU95] Mounsey, J P; Griffith, M J; Heaviside, D W; Hedley Brown, A; Reid, D S (1995), Determinants of the length of stay in intensive care and in hospital after coronary artery surgery. Br. Heart J. 73 pp.92-98

[MRG73] Morgan, D M (1973), Improved scheduling through block booking. Canadian Hospital Feb pp.41,45

[MUR89] Murata, T (1989), High-level Petri nets for logic programming and AI applications. The Third Intl. Wkshp on Petri Nets and Performance Models: Pre-Workshop Tutorials N/A pp.209-

[NAV70] Navarro, V; Parker, R; White, K L (1970), A stochastic and deterministic model of medical care utilization. Health Serv. Res. 5 pp.342-356

[NHS98] NHS Executive (1998), Information for Health: An Information Strategy for the Modern NHS 1998 -2005. NHS Executive pp.

[NHS98a] NHS Executive (1998), The Modern NHS Modern and Dependablae: A National Framework for Assessing Performance. NHS Executive pp.

[OZK95] Ozkarahan, I (1995), Allocation of surgical procedures to operating rooms. J. Medical Systems 19(4) pp.333-352

[PED94] Pedrycz, W (1994), A Generalized Fuzzy Petri Net Model. IEEE Trans, Fuzzy Systems 2(4) pp.

[PET81] Peterson, J L (1981), Petri Net Theory and the Modeling of Systems. PRENTICE-HALL pp.

[POL88] Pollack, MM, Ruttimann, UE, Getson, PR (1988), Pediatric risk of mortality (PRISM) score. Critical Care Medicine 16(11) pp.1110-1116

[PTR62] Petri, C A (1962), Kommunikation mit Automaten, PhD Dissertation. University of Bonn.

[RAT97a] Rational Software Corporation (1997) Unified Modeling Language (UML): Notation Guide. Rational Software Corporation: Santa Clara, CA N/A pp.N/A

[RAT97b] Rational Software Corporation (1997) Unified Modeling Language (UML): UML Semantics, Appendix M1 - UML Glossary. Rational Software Corporation: Santa Clara, CA

[REI97] Reis M D (1997), *The Therapeutic Intervention Scoring System: one single tool for the evaluation of workload, the work process and management?* Intensive Care Medicine 23(6) pp.615-617

[ROG94] Rogers, J; Fuller, H D (1994), Use of daily Acute Physiology and Chronic Health Evaluation (APACHE) II scores to predict individual patient survival rate. Critical Care Medicine 22(9) pp.1402-1404

[ROW94] Rowan, K M; Kerr, J H; Major, E; McPherson, K; Short, A; Vessey, M P (1994), Intensive Care Society's Acute Physiology and Chronic Health Evaluation (APACHE II) study in Britain and Ireland. Critical Care Medicine 22(9) pp.1392-1401

[SIL89] Silva, M; Valette, R (1989), Petri nets and flexible manufacturing. in Rozenberg, G (ed.), Advances in Petri Nets, Berlin: Springer-Verlag N/A pp.374-417

[SQU92a] Squire, S, Elia, A (1992), Implemening Therapeutic Intervention Scoring System on CareVue 9000. Unpublished Manuscript.

[SQU92b] Squire, S, Elia, A (1992), Implementing Therapeutic Intervention Scoring System Within the CareVue Flowsheet. Unpublished Manuscript

Within the CareVue Flowsheet. Unpublished Manuscript.

[STA71] Staff, P J; Vagholkar, M K (1971), Stationary distributions of open Markov processes in discrete time with application to hospital planning. J. Applied Probab. 8 pp.668-680

[STA94] Stansifer, R; Beven, M; Marinescu, D C (1994), Modeling concurrent programs with colored Petri nets. J. Systems and Software 26 pp.129-148

[TAM97] Tam, H; Law, H; Au, A (1997), Integrated modeling of electronics assembly systems by colored Petri nets. J. Electronics Manufacturing 7 pp.171-169

[TER94a] Teres, D (1994), Comment on "The case for using objective scoring systems to predict intensive care unit outcome". Crit Care Clin 10(1) pp.91-92

[TER94b] Teres, D, Lemeshow, S (1994), Using severity measures to describe high performance intensive care units. Crit Care Clin 10(1) pp.543-554

[TER94c] Teres, D, Lemeshow, S (1994), Why severity models should be used with caution. Crit Care Clin 10(1) pp.93-110

[TUJ92] Tu, J V, Guerriere, M R J (1993), Use of a neural network as a predictive instrument for length of stay in the intensive care unit following cardiac surgery. Proc. Ann. Symp. Computer Apps in Med. Care pp.666-672

[TUJ94] Tu, J V; Mazer, C D; Levinton, C; Armstrong, P W; Naylor, C D (1994), A predictive index for length of stay in the intensive care unit following cardiac surgery. Can. Med. Assoc. J. 151(2) pp.177-185

[TUJ96] Tu, J V; Mazer, C D (1996), Can clinicians predict ICU length of stay following cardiac surgery? Can. J. Anaesthesia 43(8) pp.789-794

[TUM92] Tuman, K J; McCarthy, R J; March, R J et al (1992), Morbidity and duration of ICU stay after cardiac surgery. A model for preoperative risk assessment. Chest 102 pp.36-44

[VAS85] Vassilacopoulos, G (1985), A simulation model for bed allocation to hospital inpatient departments. Simulation 45(5) pp.233-241

[WAT94a] Watts, CM, Knaus, WA (1994), Comment on "Why severity models should be used with caution". Crit Care Clin 10(1) pp.111-115

[YIM94] Yim, D-S (1994), A Petri-Net-Based simulation Tool for the Design and Analysis of Flexible Manufacturing Systems. J. Manufacturing Systems 13(4) pp.251-261

[YOU65] Young, J P (1965), Stabilization of inpatient bed occupancy through control of admissions. Hospitals 39 pp.41-48

[YOU66] Young, J P (1966), Administrative control of multiple-channel queuing systems with parallel input streams. Operations Res. 14 pp.145-156

[ZUR94] Zurawski, R (1994), Petri Nets and Industrial Applications: A Tutorial. IEEE Trans, Industrial Electronics 41(6)

9. Appendix: Intensive-Care Units

9.1.1. The Adult Intensive Care Unit (AICU)

AICU has a total of 20 beds, divided into 4 rooms each with 4 beds, plus 4 isolation rooms, each with one bed. A proportion of the intensive-care area is designated for MRSA-infected[1] patients, the rest is for non-infected patients. Operative and non-operative patients tend not to be distinguished geographically within the unit, although common practice is to designate one of the 4-bed rooms each day for the intake of surgical patients.

As would be expected in a healthcare facility with a large operative workload, most AICU bed-slot allocations are for patients coming direct from the operating room suite. However, there are also a proportion of bed-slots which are allocated to patients coming from the RR room (via the operating room suite) as well as from other wards located either within RBH or other hospitals.

Table 9.01 below show the percentage of patients allocated AICU bed-slots for each previous bed-slot allocation. It can be seen that 81.7 percent of allocated AICU bed-slots were for patients coming direct from the operating room suite, with 12.5% coming direct from an adult ward. The proportion of those patients coming direct from an adult ward that are operative patients and those who are non-operative is not recorded, although it is reasonable to assume that the majority would be non-operative. The majority of the 5.8% of patients coming direct from the RR room represent failed fast-track or overnight recovery cases.

Previous bed-slot allocation	Frequency	Percent
OR	2690	81.7
Ward	410	12.5
RR	191	5.8
Total	3291	100.0

Table 9.01. Percentage of patients allocated AICU bed-slots for each previous bed-slot allocation

As well as recording the previous bed-slot allocation, the initial bed-slot allocation within RBH is also recorded. In all cases this will either be another ward within RBH or another hospital. The dataset for the initial bed-slot allocation is as follows:

[1] Methicillin-resistant staphylococcus aureus, a non-fatal but antibiotic-resistant infection commonly found amongst high-dependency patients.

- **Alex Ward**. This is a ward whose bed-slots are primarily designated for consumption by operative cardiac patients.
- **Elizabeth Ward**. This is a ward whose bed-slots are primarily designated for consumption by operative thoracic patients.
- **Other Hospital**. This could be a bed-slot allocation in any type of unit in any other hospital or healthcare facility other than RBH.
- **Paul Wood Ward**. This is a ward whose bed-slots are primarily designated for consumption by operative cardiac patients.
- **Reginald Wilson Ward**. This is a ward whose bed-slots are primarily designated for consumption by either cardiac or thoracic operative patients who are self-funded.
- **South Block**. This is a ward whose bed-slots are primarily designated for consumption by operative cardiac patients.
- **York Ward**. This is a ward whose bed-slots are primarily designated for consumption by non-operative cardiac patients.

Table 9.02 below show the percentage of patients allocated AICU bed-slots for each initial bed-slot allocation.

Initial bed-slot allocation	Frequency	Percent
Alex	1261	38.3
York	546	16.6
Elizabeth	505	15.3
Reginald Wilson	465	14.1
Paul Wood	342	10.4
Other Hospital	139	4.2
South Block	32	1.0
Total	**3290**	**100.0**

Table 9.02. Percentage of patients allocated AICU bed-slots for each initial bed-slot location

With regards to the subsequent bed-slot allocation after a patient is discharged from AICU, Table 9.03 below shows the percentage of subsequent bed-slot allocations for the different possible units to which the patient may be allocated a bed-slot. The dataset used is the same as that used for the previous bed-slot allocation above with the additional possibility that a patient may die within AICU and therefore have the hospital morgue as the subsequent location.

Destination	Frequency	Valid Percent
Alex	1638	49.9
Elizabeth	578	17.6
Reginald Wilson	542	16.5
Morgue	248	7.6

Paul Wood	102	3.1
Other Hospital	77	2.3
York	54	1.6
South Block	45	1.4
Total	**3284**	**100.0**

Table 9.03. Percentage of patients allocated bed-slots for each destination ward from AICU

The case-mix of patients admitted to AICU is, as mentioned earlier, constituted primarily of operative patients. The majority of the surgical patients are cardiac patients (78.2%, see Table 9.04 below), which is more a reflection of the increased level of criticality of post-operative cardiac patients as opposed to thoracic, general, or vascular[2,3] patients than of the larger absolute number of cardiac procedures undertaken.

Operative Category	**Frequency**	**Percent**
Cardiac	2574	78.2
Non-operative	361	11.0
Thoracic	232	7.0
General	84	2.6
Vascular	40	1.2
Total	**3291**	**100.0**

Table 9.04 Percentage of AICU admissions from each Operative Category

With regards to the gender-mix of the patients allocated AICU bed resources, the majority of patients are male with 70.3% as shown in Table 9.05 below. This result is a reflection of the elevated risk of heart disease amongst men.

Sex	**Frequency**	**Percent**
Male	2312	70.3
Female	979	29.7
Total	**3291**	**100.0**

Table 3.05. Gender-mix of AICU admissions

With regards to the age distribution of patients allocated AICU bed-slots, the mean age 60 (SD=14.61). As would be expected, the age distribution is not normal as is shown in Figure 9.01 below, with a significant skew to the left.

---

[2] Vascular surgery was ceased at RBH effective from April 1997.

[3] In fact, because of the need to use a bypass machine during major cardiac surgery, it is true to say that cardiac surgery involves the greatest level of criticality of all types of surgery.

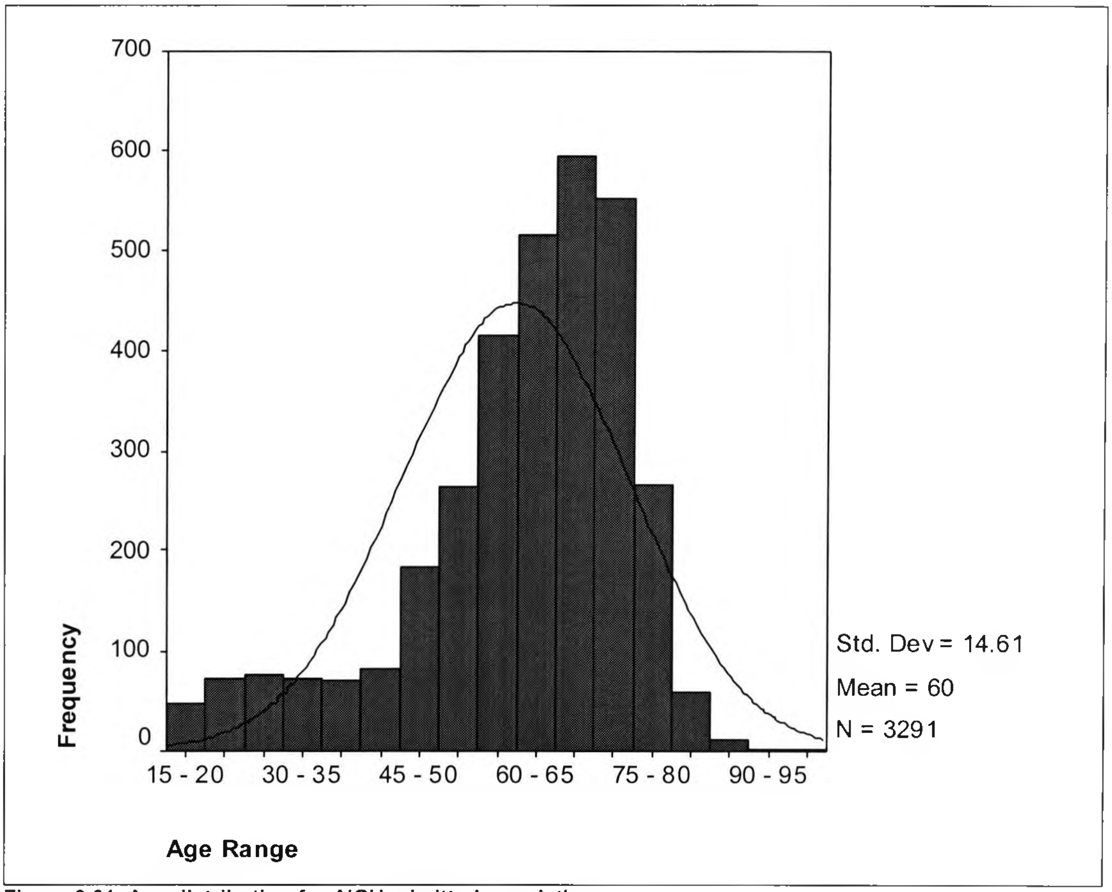

Figure 9.01. Age distribution for AICU admitted population.

There is a relatively high mortality rate in AICU, as would be expected given the critical state of the patients admitted. As can be seen in Table 3.06 below, the mortality rate in AICU is 7.7%.

AICU Outcome	**Frequency**	**Percent**
Alive	3039	92.3
Dead	252	7.7
Total	**3291**	**100.0**

Table 3.06. Outcome distribution of AICU admissions

From the perspective of resource allocation, one of the most important distributions is that of the bed-slot allocation size which shows the distribution of bed-slots amongst the patient population. Figure 9.02 below shows the AICU bed-slot allocation size distribution expressed as the length of stay (LOS) measured in days of AICU patients.

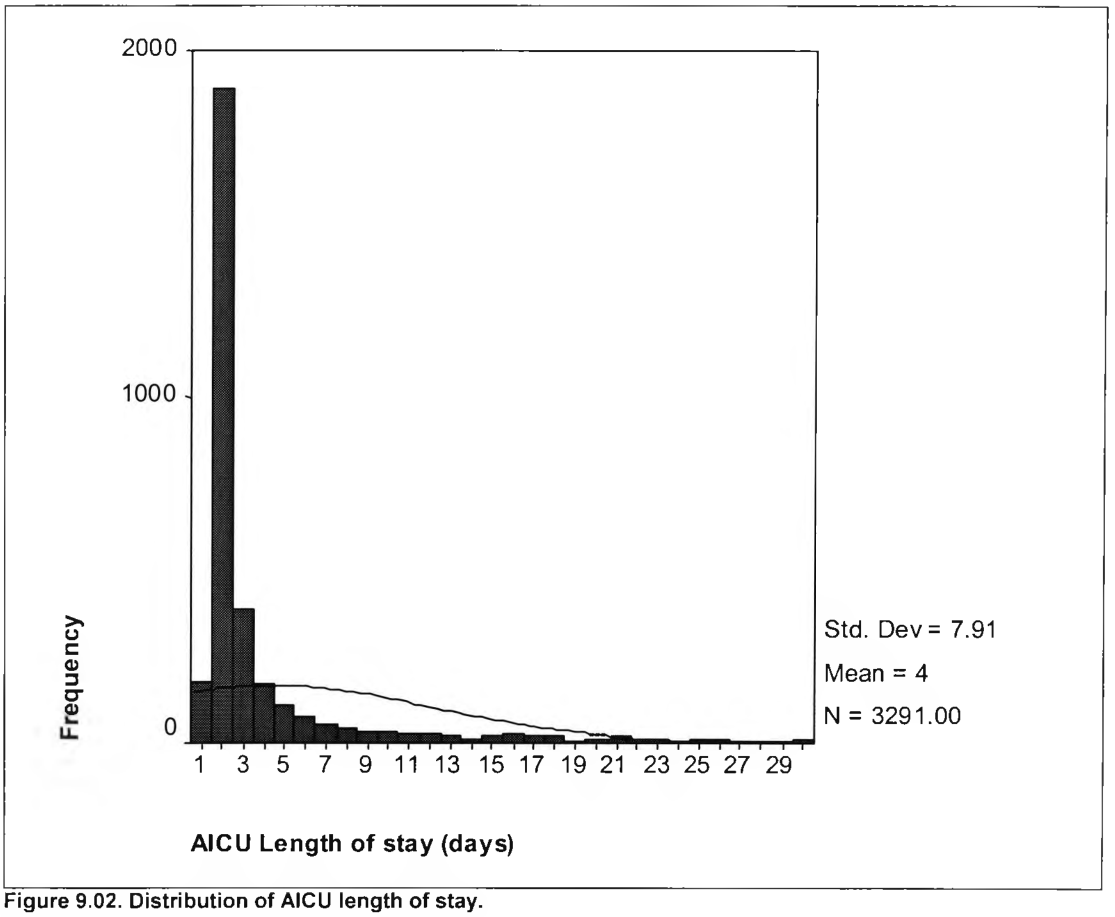

Figure 9.02. Distribution of AICU length of stay.

As can be seen from Figure 9.02, the distribution is not normal with a skew to the right and a long tail. This is the characteristic type of distribution for resource allocation in healthcare. It is significant since, although the proportion of bed-slot allocations which make up the right-hand tail is small, the number of consumed bed-slots which they represent is disproportionately large. This is clearly demonstrated in Figure 9.03 below which compares on a log scale the percentage of the total number of bed-slot allocations for each allocation size and the percentage of the total consumption for each allocation size.

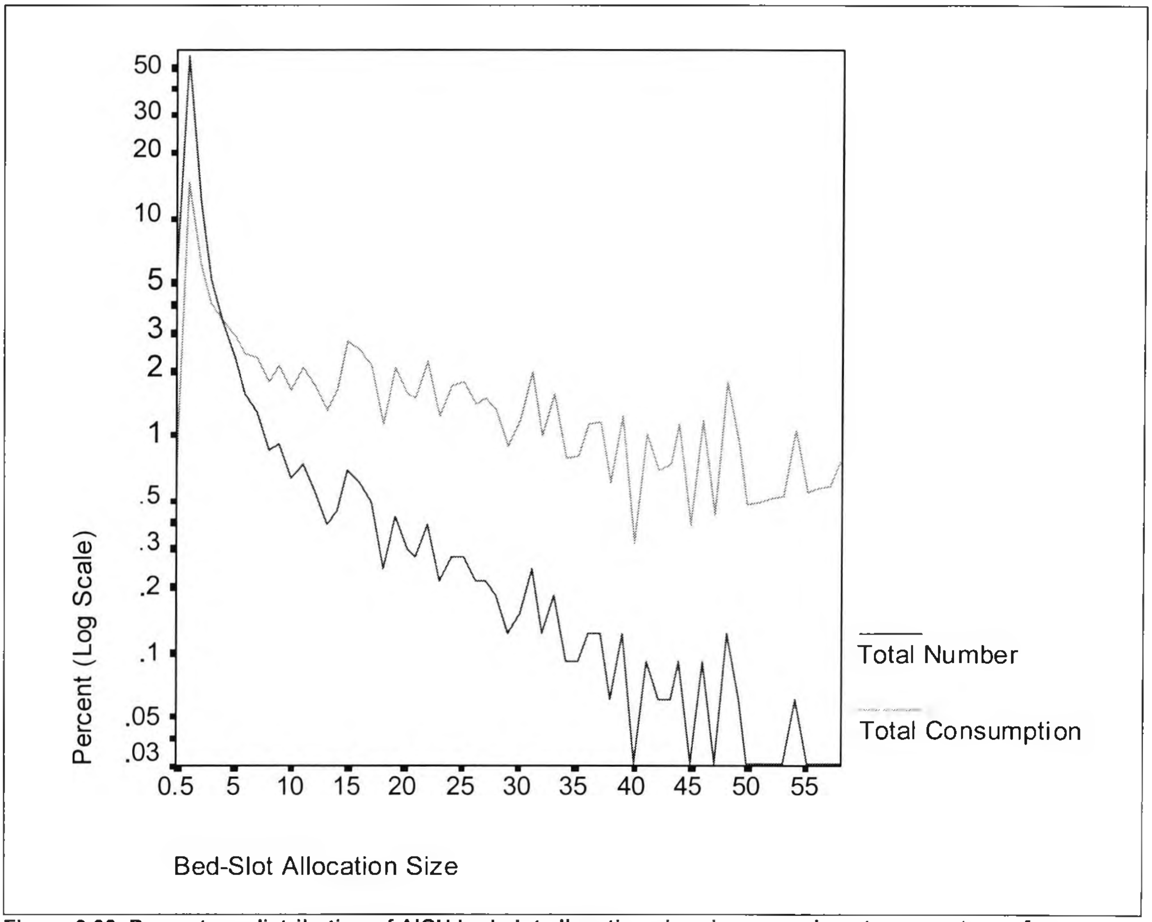

Figure 9.03. Percentage distribution of AICU bed-slot allocation sizes in comparison to percentage of resources consumed by each allocation size

It can be seen from Figure 9.03 that although the larger bed-slot allocation sizes represent only a fraction of a percent of the total number of bed-slot allocations, they can represent around ten times as much of the percentage of the total consumption of bed-slots.

9.1.2. The Paediatric Intensive Care Unit (PICU)

For the paediatric intensive-care unit, most of the comments made above concerning the AICU apply in much the same measure to PICU. There are some differences, however. In particular, surgical paediatric cases will tend to have less predictable and larger bed-slot allocation sizes than adult operative patients. This is primarily due to most paediatric cases being congenital, which often involves less standardised and less frequently undertaken surgical procedures, often requiring much longer intra-operative durations than in adult cases. All of these factors contribute to a longer post-operative recovery period.

For the admission and discharge profiles of PICU, in terms of location, the situation is similar to that of AICU in terms of admissions; most admissions will come direct from theatre, or direct from a hospital ward, whether external or internal. The only patients being admitted direct from RR will be those for whom RR was unable to recover - there are no paediatric fast-track cases. The discharges from PICU will be very different from AICU in that the majority will be discharged to the paediatric intermediate-care unit, or Rose ward, which is designated specifically for paediatric cases.

With regards to the resourcing of PICU, the level and profile is proportionately much the same as with AICU. There are, however, much fewer beds in PICU - 11 in total, 7 of which are publicly financed - which, when combined with the increased levels of unpredictability in LOSs for paediatric cases, tends to magnify the problems involved in maximising resource utilisation.

10. Appendix: Evaluating Resource Targeting in an Adult Intensive Care Unit

The aim of this study is to examine the extent to which the Bed-slot Assumption referred to in the previous chapter is validated by the available data. The Bed-slot Assumption states that the adoption of the notion of a bed-slot is a valid simplification for use in resource allocation. As it stands, the notion of validity being used here is in need of further clarification. In order to determine the extent to which the assumption is validated by the data it needs to be re-cast in terms of testable hypotheses. Four such hypotheses may be identified, as follows:

5. **The Predictability Hypothesis**. Individual non-generic resources or different types of bed-slot whose individual allocation sizes are not predictable in advance of consumption within an acceptable degree of accuracy for different types of patient should be modelled as generic resources, either individually or as components of a bed-slot type in the case of individual non-generic resources, or in the case of types of bed-slot as instances of a more generic type of bed-slot; and
6. **The –G&F Proportion Hypothesis**. The –G&F resources whose individual allocation sizes are predictable in advance of consumption within an acceptable degree of accuracy do not constitute an excessively large proportion of the total resource consumption; and
7. **The Variance Hypothesis**. The overall level at which those resources which are modelled as components of the bed-slot are consumed at the population level is within an acceptable degree of variance within different categories of patients, or that such variance may not be reduced through the categorisation of patients into different categories where the categorisation of patients into those categories may be made in advance of consumption within an acceptable degree of accuracy; and
8. **The Difference Hypothesis**. There is no significant difference in the overall level of consumption of those resources which are modelled as components of the bed-slot between different categories of patients, or that where there is a difference between different categories, the categorisation of patients into those categories may be made in advance of consumption within an acceptable degree of accuracy.

The first two of these hypotheses refer to ability to identify the existence of a bed-slot type. That is, for any given healthcare facility, if both of these hypotheses may be shown to be supported, then there is a type of bed-slot that can be validated by the data. It does not show, however, that this bed-slot type should be used as the basis of resource allocation.

The last two of these hypotheses refer to a specific application of the notion of the bed-slot. That is, for any given healthcare facility, if both of these hypotheses may be shown to be supported by the available data, then any bed-slot(s) which have already been identified may not effectively be categorised further into subtypes of bed-slot. Conversely, if the hypotheses are not supported by the available data, then there are subtypes of bed-slots which can be identified and consequently result in improved healthcare resource allocation.

To see how these hypotheses may be applied to a real example, consider the case of the AICU. There currently exists one type of AICU bed-slot so that all patients who are allocated to AICU are allocated the same bed-slot. If it can be shown for AICU that both the Predictability Hypothesis and the -G&F Hypothesis are supported by the available data, then this implies that an AICU bed-slot may be identified. It does not show, however, that this bed-slot type should be used as the basis of resource allocation. If it can be further shown that the Variance Hypothesis and Difference Hypothesis are supported by the data, then this implies that the use of the AICU bed-slot is economically or clinically justified as the basis of resource allocation. If at least one of these latter hypotheses are not supported, however, this would imply that, while the AICU bed-slot type exists, it should not be used as the basis of resource allocation since there also exists subtypes of bed-slot of the AICU bed-slot type which would result in improved resource allocation.

In the situation where the first two hypotheses are shown to be supported, and at least one of the latter two are not, this does not mean that the AICU should be split into different units according to the subtypes of AICU bed-slots which have been identified. Nor does it mean that the component resources of each subtype of bed-slot must be located within distinct areas within AICU (although this would probably make some sense and is the route often taken). Rather, it means only that when allocating bed-slots to patients requiring AICU admission, allocation should be in terms of subtypes of the general AICU bed-slot type - for example, AICU operative bed-slots or AICU non-operative bed-slots - instead of in terms of the general AICU bed-slot type itself.

It can be seen that each of these hypotheses is framed in somewhat vague terminology. For example, “excessively large proportion” or “acceptable degree of accuracy”. Unfortunately, this is inevitable since no clear criteria exist for determining, for example, what is an acceptable level of accuracy. In particular, with regards to The Difference Hypothesis, the notion of significance cannot be plausibly reduced to that of statistical difference, since statistical significance need not imply a difference which is either clinically or economically significant. For this reason, it is necessary to define sets of threshold values in each case which are intuitively plausible. Thus, whenever such threshold values are defined it will be alongside some argumentation for adopting that particular value as the most plausible.

The structure for this study is to consider each of the four hypotheses in turn after first providing a definition of the variables and a discussion of the data which was used in the study. This will be followed by a general discussion and summary. The method, results and discussion will be self-contained in each section which deals with a specific hypothesis. The hypotheses will be tested in the order in which they are presented above.

10.1. Data

The data which is used in this study comes from the Combined AICU Database. This database was constructed by merging the AICU Audit Database and the AICU TISS Database. There are a total of 5,867 records in the AICU TISS Database for the period 20/04/93 to 17/11/94, 5,485 of which were matched with records from the AICU Audit Database for the same period. Each record in the Combined AICU Database corresponds to an AICU bed-slot allocation, where a bed-slot is defined as being of a 24 hour duration. Thus, a patient who consumes 3 AICU bed-slots would generate three records within the Combined AICU Database.

The reason for choosing bed-slot from AICU rather than the operating room or the recovery room, for example, may be justified on both pragmatic and clinical grounds. Pragmatically, the TISS score is specifically designed to measure resource consumption in intensive-care units, with no equivalent measure existing for either the operating room or the recovery room, nor any alternative dataset which is recorded that may serve the same purpose. Clinically, the choice of the AICU is justified on the grounds that the cost-effective allocation of intensive-care resources is more dependent on the accurate categorisation of patients into the types of bed-slots implied by their resourcing requirements. This is because of the higher cost of intensive-care medicine and the economic advantages in being able to provide only the level of resourcing implied by the patient’s treatment, rather than that level of resourcing implied by the healthcare unit in which they are being treated (the idea of the bed-slot being that they should amount to the same thing).

With regards to the data itself, it should be noted that because of the constantly changing nature of intensive-care medicine, the range of dates used in the study means the data is now of less relevance than when it was collected. Nevertheless, it is still possible to evaluate the hypotheses 'in principle' using the data even though it may not be a true reflection of current practice.

The fields in the Combined AICU Database relevant to this study are as follows:

- *Consumed AICU Bed-slots*. The number of (contiguous) AICU bed-slots which have already been consumed from the total allocation by the same patient;
- *Remaining AICU Bed-slots*. The number of (contiguous) AICU be-slots which remain to be consumed from the total allocation by the same patient;
- *Chronicity*. A contrast variable indicating whether or not the patient to whom the bed-slot is allocated requires a chronic AICU bed-slot allocation, defined as chronic if Consumed AICU Bed-slots is greater than 2, otherwise defined as non-chronic;
- *Operative Status*. A contrast variable indicating whether the patient to whom the bed-slot is allocated is classified as operative or non-operative;
- *TISS Variables*. A set of 77 different contrast variables, each indicating whether or not the bed-slot is constituted by or implies various specific or groups of specific resource inputs. Examples of TISS variables are Naso-Gastric Feeding or Chest X-Ray;
- *Total TISS Interventions*. A quantitative variable indicating the number of TISS components which were consumed as part of the bed-slot;
- *Total TISS Score*. A quantitative variable similar to Total TISS Interventions, but with a weight attached to each TISS component according to the implied cost of the resources input;
- *TISS Ratio*. A quantitative variable derived by the division of Total TISS Score by Total TISS Interventions which represents the average weighting given to each consumed TISS component of the bed-slot.

Each individual TISS variable and the total TISS variables (items 6, 7 and 8 in the above list) for each previous contiguous bed-slot consumed by the same patient is also recorded for those records where

Consumed AICU Bed-slots is greater than 1. Thus, for example, if the Total TISS score for the first AICU bed-slot consumed by a particular patient is 35, then if the patient consumes a second AICU bed-slot this figure will also be recorded in the record for that bed-slot although under the field Previous Total TISS Score.

10.2. The Hypotheses

10.2.1. The Predictability Hypothesis

The Predicability Hypothesis states that individual non-generic resources or different types of bed-slot whose individual allocation sizes are not predictable in advance of consumption within an acceptable degree of accuracy for different types of patient should be modelled as generic resources, either individually or as components of a bed-slot type in the case of individual non-generic resources, or in the case of types of bed-slot as instances of a more generic type of bed-slot.

For example, if it is not possible to predict which patients will require mechanical ventilation and for how long, then the resources associated with mechanical ventilation should be modelled as a generic resource or be included as a component resource of a bed-slot type.

In the case of bed-slots, if it is not possible to accurately predict a patient's bed-slot requirements in advance of consumption of a particular type of bed-slot, then this is an indication that the notion of a more generic bed-slot should be introduced instead of particular types of bed-slots. For example, if two types of AICU bed-slots were defined according to whether or not a patient would require mechanical ventilation as a component of the bed-slot, then the use of these two types of bed-slot can only be effective if it is possible to predict in advance of consumption of those bed-slots which one a patient will require. If this cannot be predicted, then a generic AICU bed-slot should be defined which does not discriminate between the inclusion of mechanical ventilation resources.

The support for the Predictability Hypothesis comes from the following line of reasoning:

- Let $\boldsymbol{P}$ be a representative sample of a patient population whose resourcing requirements are known post hoc; and
- Let $\boldsymbol{P'}$ be a representative sample of a patient population whose resourcing requirements are not known; and

- Let two types of bed-slot, P_1 and P_2, be defined post-hoc for each patient within $\boldsymbol{P}$ such that each individual patient, $p_1,p_2,\ldots,p_m$, is assigned exactly one type of bed-slot as follows according to a set-membership rule R, where R is based on patients' resourcing requirements:

  $P_1 = \{p_1,p_2,\ldots,p_{n-k}\}$; and

  $P_2 = \{p_{n-(k-1)},p_{n-(k-2)},\ldots,p_n\}$

If it is now assumed that bed-slots are identified with sets, then not only are P_1 and P_2 bed-slots, but so are $\boldsymbol{P}$ and $\boldsymbol{P'}$. If it is further assumed that the members of $\boldsymbol{P}$ and those of $\boldsymbol{P'}$ are samples from the same population, then $\boldsymbol{P}$ bed-slots are the same as $\boldsymbol{P'}$ bed-slots. However, while the two sets P'_1 and P'_2 defined as

$P'_1 = \{p'_1,p'_2,\ldots,p'_{m-j}\}$; and

$P'_2 = \{p'_{m-(j-1)},p'_{m-(j-2)},\ldots,p'_m\}$

may be hypothesised to exist, each individual patient may in reality only be assigned exactly one of the above types of bed-slot according to the rule R if the patients' resourcing requirements are known. It was assumed, however, that $\boldsymbol{P'}$ is a sample of a patient population whose resourcing requirements are not known. Therefore, the patients $p'_1,p'_2,\ldots,p'_m$, cannot be assigned to the bed-slots P'_1 and P'_2. In this case, the only bed-slot which can with certainty be assigned to each patient in $\boldsymbol{P'}$ is the $\boldsymbol{P'} = \boldsymbol{P}$ bed-slot.

This conclusion may be expressed in terms of the distinction between generic and non-generic resources by saying that a resource is generic only if the knowledge regarding its consumption is generic, rather than the consumption itself; a resource is non-generic only if the knowledge regarding its consumption is non-generic. Thus, for example, the consumption of a mechanical ventilator may be non-generic insofar that some patients will consume the resource and others not consume it. However, for each patient, the knowledge of whether or not that patient will consume a mechanical ventilator may be generic insofar that for each patient it will not be known if the patient will consume a mechanical ventilator.

This line of reasoning assumes a binary classification of whether or not patients' specific resourcing requirements are known. If a patient has already consumed the resources then resourcing requirements are known, since the consumption of those resource has occurred in the past; if a patient is yet to consume the resources then resourcing requirements are not known, since consumption occurs in the future and is thus uncertain.

In reality, however, it may be more appropriate to use a fuzzy or probabilistic conception of patient categorisation according to bed-slots or specific resourcing requirements. In this case, the probability that a $\boldsymbol{P}$ ($\boldsymbol{P'}$) patient will consume a $\boldsymbol{P}$ ($\boldsymbol{P'}$) bed-slot is 1. Similarly, a patient whose resourcing requirements are identified post-hoc will have a probability of having consumed, for example, a P_1 bed-slot of 1, and a probability of having consumed a P_2 bed-slot of 0. But in those cases where the patient is yet to consume all resources, the probability that the patient will require a P_1 bed-slot is x, where $1>x>0$, and a probability of requiring a P_1 bed-slot is $1-x$.

Adopting this probabilistic notion, it is appropriate to pose the question as to which level of probability is necessary to justify the further classification of patients into more specific types of bed-slots? There are two main considerations in this respect:

- What is the cost of making a misclassification?
- What is the benefit of making a further classification of bed-slots?

Unfortunately it is beyond the scope of this discussion to examine these considerations in more detail, although, as shall be discussed below, they are important factors in interpreting the results of this study.

10.2.2. The –G&F Proportion Hypothesis

Method

The –G&F proportion hypothesis states that –G&F (non-generic and fixed) resources whose individual allocation sizes are predictable in advance of consumption within an acceptable degree of accuracy do not constitute an excessively large proportion of the total resource consumption.

The method adopted to evaluate the –G&F proportion hypothesis is to first identify those –G&F resources whose consumption is directly implicated in individual TISS components. To be classified as being a TISS component which implicates a –G&F resource, the TISS component must satisfy the following criteria:

- It directly implicates the consumption of a resource item which is both high-cost and fixed; and

- It has an incidence rate which is greater than 0.05 and less than 0.80.

Note that the first of these criteria determine the fixed or 'F' component of –G&F resources; the second determines the non-generic or '-G' component of –G&F resources. The method for determining whether or not a TISS component satisfies the first criterion is largely a matter of judgement. In this regard, there are TISS components which obviously implicate high-cost fixed resources, such as those which imply the consumption of mechanical ventilators or dialysis machines. There are also components, however, whose classification as high-cost and fixed or not may be thought of as borderline. In these borderline cases it is appropriate to err on the side of caution by classifying them as high-cost and fixed.

The second of the above criteria puts threshold values on whether or not a TISS component implicates resources which are both non-generic and whose consumption is a regular occurrence. The threshold for determining whether or not the resources should be considered as non-generic is an incidence rate of less than 0.80; that is, the resources are not considered as being generic if less than 80% of all AICU bed-slots contain those resource inputs. The threshold for determining whether or not the resource should be considered as being regularly consumed is an incidence rate of greater than 0.05; that is, the resources are considered as being regularly consumed if more than 5% of all AICU bed-slots contain those resource inputs.

The second stage is to determine if the total number of those TISS components which satisfy the above two criteria represents a sufficiently small enough proportion of the total number of TISS components to support the –G&F proportion hypothesis. The criteria used here will inevitably be to some extent arbitrary. The threshold adopted here is 25%. That is, if the proportion of TISS components which directly implicate –G&F resources according to the above criteria is less than 25% of the total, then this may be taken as being supporting evidence for the –G&F Proportion Hypothesis. 25% is appropriate in this context since the weighting of –G&F TISS components will tend to be higher than those of other resources. Therefore, assuming a liberal ratio of 3:1 between the average weightings of –G&F TISS components and other components, if the number of –G&F TISS components is less than 25% of the total, the proportion of the total TISS score represented by those components cannot be greater than 50%.

According to this method, if the proportion of –G&F TISS components is less than 25%, the –G&F Proportion Hypothesis is assumed to be consistent with the data. It does not follow, however, that if the proportion is greater than 25% then the data will be inconsistent with the hypothesis since the hypothesis also states that the consumption of those –G&F resources also be predictable in advance of consumption. Therefore, only if the total proportion of –G&F TISS components is greater than 25% and at least 25% of the TISS components are –G&F TISS components whose consumption cannot be predicted in advance of consumption within an acceptable degree of accuracy will the data be inconsistent with the –G&F Proportion Hypothesis. The reasoning behind this further condition will be given in support of the –G&F Prediction Hypothesis.

RESULTS

Table 10.01 below shows the classification of each TISS component according to the criteria given above.

TISS Component Variable	Mean	High-Cost AICU Capital	0.05<Mean<0.80	-G&F
Controlled Ventilation	0.42	Yes	Yes	Yes
Clinitron Bed	0.73	Yes	Yes	Yes
Blood Transfusion ...	0.12	Yes	Yes	Yes
Platelet Transfusion	0.09	Yes	Yes	Yes
SIMV and/or Pressure Supp.	0.51	Yes	Yes	Yes
CPAP	0.12	Yes	Yes	Yes
Hypothermia Blanket	0.23	Yes	Yes	Yes
Hourly Neuro Vital Signs	0.42	Yes	Yes	Yes
C. Arrest and/or Def.	0.02	No	No	No
Swan-Ganz/LA Line	0.38	No	Yes	No
Haemodyalysis in unstable	0.02	Yes	No	No
Peritoneal Dyalysis	0.00	Yes	No	No
CAVHD/CVVHD	0.07	No	Yes	No
Intracranial Pressure Monitored	0.00	Yes	No	No
Intra-Aortic Balloon Pump	0.06	No	Yes	No
Em. Op. in past 24 hrs	0.03	No	No	No
Sengstaken Tube	0.00	No	No	No
Em. Endoscopy/Broncoscopy	0.02	No	No	No
>1 Inotopic/CVS Drug	0.34	No	Yes	No
Active Pacing	0.07	No	Yes	No
TPN	0.12	No	Yes	No
Pacing wire in situ not ..	0.21	No	Yes	No
Chest Drains	0.49	No	Yes	No

Concentrated Infusions of ..	0.65	No	Yes	No
Active Treatment for Elec.	0.16	No	Yes	No
1-2 hour Suctioning	0.89	No	No	No
Triple Lumen Line	0.84	No	No	No
>4 Stat. B/Tests per s.	0.75	No	Yes	No
Continuous Antiarrythmic	0.10	No	Yes	No
Cardioversion for Arrythmia	0.02	No	No	No
Arterial Line	0.91	No	No	No
CO Measurement	0.34	No	Yes	No
Active Diuresis	0.28	No	Yes	No
Active Anticoagulation	0.20	No	Yes	No
>2 IV Antibiotics	0.25	No	Yes	No
Semi-Emergency IV Stat do.	0.05	No	Yes	No
Treatment of fits	0.01	No	No	No
Complicated Othopaedic	0.00	Yes	No	No
Acute Digitalisation	0.03	No	No	No
1 Inotropic/CVS Drug	0.26	No	Yes	No
Intubation in ICU	0.04	Yes	No	No
Tonomoeter	0.01	No	No	No
CVP	0.89	No	No	No
2 Peripheral IV Caths	0.13	No	Yes	No
Haemodialysis-Stable	0.02	Yes	No	No
Tracheostomy in past 48 hrs.	0.03	No	No	No
Spontaneous Breathing	0.20	No	Yes	No
NG Feeding	0.46	No	Yes	No
Regular IV Drugs not anti.	0.37	No	Yes	No
Multiple Dressing Chnages	0.10	No	Yes	No
Betadine Irrigation	0.01	No	No	No
Clear IV Fluids for Dehyd.	0.61	No	Yes	No
Renal Dose Dopamine	0.63	No	Yes	No
Haemofiltration Fluid Balance	0.09	No	Yes	No
ECG Monitoring	1.00	Yes	No	No
Hourly Vital Signs	0.99	Yes	No	No
1 Peripheral IV Cath.	0.36	No	Yes	No
Fluid Balance Chart	0.95	No	No	No
Tracheostomy Care	0.29	No	Yes	No
Pressure Score	0.12	No	Yes	No
Urinary Cath.	0.92	No	No	No

Oxygen via nasal specs/ma.	0.30	No	Yes	No
2 or less IV Antibiotics	0.54	No	Yes	No
Chest Physiotherapy	0.89	No	No	No
Colostomy	0.01	No	No	No
Enema/Glycerine Supp.	0.04	No	No	No
Routine Dressing Changes	0.84	No	No	No
Standard Orthopaedic Trac.	0.00	No	No	No
Chronic Anticoagulation	0.02	No	No	No
PRN IV Drugs	0.66	No	Yes	No
Chest X-Ray	0.97	No	No	No
Unconvent. Vent.	0.02	Yes	No	No
Unconvent. Mech.	0.01	Yes	No	No
Confusion	0.01	No	No	No
Cont. Diarrhoea	0.01	No	No	No
SG Plus LA	0.01	No	No	No

Table 10.01. Classification of TISS Variables according to resource-type

From Table 10.01 above it can be seen that there are 8 –G&F TISS components according to the criteria above. This represents a percentage of 11% of the total number of TISS components (74), which is less than the threshold of 25% necessary for the data to be considered inconsistent with the –G&F Proportion Hypothesis.

Because the proportion of –G&F TISS components is less than 25%, the –G&F Proportion Hypothesis is assumed to be consistent with the data. It is therefore not necessary to further show that the consumption of those –G&F resources be predictable in advance of consumption within an acceptable degree of accuracy. The extent to which the consumption of those -G&F TISS components will, however, be examined in a later chapter in relation to a slightly different hypothesis.

10.2.3. The Variance Hypothesis

The overall level at which those resources which are modelled as components of the bed-slot is within an acceptable degree of variance within different categories of patients. Alternatively, that if the variance is not within an acceptable range, it may not be substantially reduced through the further categorisation of patients into different categories where such categorisation may be made in advance of consumption within an acceptable degree of accuracy

METHOD

The measures of overall resource consumption to be used are the three overall TISS measures: Total TISS Interventions, Total TISS Score and TISS Ratio. The reason for using all three of these scores rather than just one is that each indicates subtle differences in how the overall level of resource consumption is constituted which would otherwise be masked if one single measure was adopted. For example, Total TISS Score could be quite high, but TISS Ratio may be quite low, indicating that overall consumption is primarily constituted by low scoring TISS components.

The measures of variance which will be used are standard deviation (SD) and mean 95% confidence intervals. The determination of whether or not the variance is within an acceptable range will be based on intuitive notions of what is acceptable. To support the use of an intuitive evaluation of variance measures, these will be compared to the resulting variance measures which result from the further classification of patients according to Age and Operative Status. Thus, if the reduction in variance achieved by such further classification is insubstantial, this will be considered as supportive evidence of the Variance Hypothesis. A further comparison will be made with an optimal classification of bed-slots according to two categories generated by a k-means cluster analysis using the three overall TISS variables as parameters. In each case the basis of the comparison will be made using the R^2 measure derived from a set of linear regression models using the proposed classification variables as independent variables (interpreted in the regression analysis as dummy quantitative variables) and each of the three overall TISS variables as dependents.

Each distribution for the overall TISS variables will be tested for normality for each category of each of the proposed bed-slot classifications using the Kolmogorov-Smirnov test for normality and a Q-Q normality plot.

RESULTS

Each distribution for the overall TISS variables was judged to be normal ($p<0.01$) for each category of each of the proposed bed-slot classifications using the Kolmogorov-Smirnov test for normality and a Q-Q normality plot.

The descriptive statistics for the total distributions for Total TISS Interventions, Total TISS Score and TISS Ratio is shown in Table 10.02 below:

TISS Variable	Mean	SD	95% CI (lower bound)	95% CI (upper bound)
Total TISS Interventions	22.7	4.84	22.60	22.86

Total TISS Score	48.6	14.11	48.25	49.01
TISS Ratio	2.10	0.234	2.096	2.108

Table 10.02. Descriptive statistics for aggregated TISS Variables

As may be seen from Table 10.02 above, the standard deviation for each TISS variable is relatively small in comparison to the mean, and the differences between the upper and lower bounds of the 95% confidence intervals are also within a relatively narrow range. This interpretation of the results lends support for the Variation Hypothesis.

The descriptive statistics for the distributions of each aggregate TISS variable for each category of the proposed bed-slot classifications are shown in Table 10.03, Table 10.04 and Table 10.05 below

Proposed bed-slot classification	**N**	**Mean**	**SD**	**95% CI (lower bound)**	**95% CI (upper bound)**
Chronicity:					
Non-chronic	1994	23.83	4.28	23.65	24.02
Chronic	3491	22.10	5.03	21.94	22.27
Operative Status:					
Non-operative	1648	21.60	5.19	21.35	21.85
Operative	3837	23.22	4.60	23.07	23.36
TISS Cluster:					
TISS Cluster 1	2702	19.57	4.30	19.40	19.73
TISS Cluster 2	2783	25.81	3.00	25.69	25.92

Table 10.03. Descriptive statistics for Total TISS Interventions for each proposed bed-slot category

Proposed bed-slot classification	**N**	**Mean**	**SD**	**95% CI (lower bound)**	**95% CI (upper bound)**
Chronicity:					
Non-chronic	1994	51.42	12.69	50.86	51.98
Chronic	3491	47.04	14.64	46.55	47.53
Operative Status:					
Non-operative	1648	45.66	14.64	44.95	46.36
Operative	3837	49.91	13.69	49.48	50.35
TISS Cluster:					
TISS Cluster 1	2702	39.01	11.24	38.59	39.44
TISS Cluster 2	2783	57.97	9.65	57.61	58.33

Table 10.04. Descriptive statistics for Total TISS Score for each proposed bed-slot category

Proposed bed-slot classification	**N**	**Mean**	**SD**	**95% CI (lower bound)**	**95% CI (upper bound)**
Chronicity:					
Non-chronic	1994	2.132	0.208	2.123	2.141
Chronic	3491	2.085	0.246	2.077	2.093
Operative Status:					

Non-operative	1648	2.066	0.250	2.054	2.078
Operative	3837	2.118	0.225	2.111	2.125
TISS Cluster:					
TISS Cluster 1	2702	1.964	0.219	1.956	1.972
TISS Cluster 2	2783	2.237	0.158	2.231	2.242

Table 10.05. Descriptive statistics for TISS Ratio for each proposed bed-slot category

As can be seen from Table 10.03, Table 10.04 and Table 10.05 above, there are significant differences between the means at the $p<0.05$ level for each proposed bed-slot classification for each TISS variable. With regards to the reduction in variance in comparison to the control situation where no classification of bed-slots is adopted, it can be seen that neither classification on the basis of Chronicity or on the basis of Operative Status results in a reduction in variance which would be able to justify making such classifications as the basis of a proposed further bed-slot categorisation on an intuitive level. This is in contrast, however, to the reduction in variance attained through the derived TISS Cluster classification which is far more significant.

This reasoning behind this conclusion may be seen more clearly in Table 10.06 below which shows for each classification the R^2 measure resulting from a linear regression analysis, the resulting change in the F statistic and the associated p value of the change in the F for the overall measure Total TISS Score.

Classification	R2	F Change	p
Chronicity	0.022	125	<0.001
Operative Status	0.019	107	<0.001
TISS Cluster	0.451	4502	<0.001

Table 10.06. R-squared, F change and significance level for each TISS classification

It can clearly be seen from Table 10.06 that the reduction in variance (as measured by R^2) is an order of magnitude greater for the derived TISS Cluster classification than for the others. According to the Variance Hypothesis, however, this will be interpreted as a disconfirmation of the hypothesis only if it is possible to predict within an acceptable degree of accuracy which TISS Cluster bed-slot category is implied by a patient's clinical and demographic characteristics in advance of consumption of either of those bed-slots.

To determine whether or not patients bed-slot requirements may be accurately predicted on the basis of clinical and demographic characteristics in advance of consumption, a binary logistic regression was used with TISS Cluster as the dependent variable and the variables listed below as independents (with contrast variables interpreted as dummy quantitative variables in the analysis). Reference categories are noted in parentheses except in the diagnosis categories where the reference category in each case is where the diagnosis is positive.

1. Patient Age
2. Sex (male)
3. Previous bed-slot allocation (OR)
4. Diagnosis - AAA
5. Diagnosis - Aortic valve disease
6. Diagnosis - ARDS
7. Diagnosis - ASD
8. Diagnosis - Athsma
9. Diagnosis - Atelectasis
10. Diagnosis - CA Lung
11. Diagnosis - CCF
12. Diagnosis - Coarctation
13. Diagnosis - Heart transplant
14. Diagnosis - Heart-lung transplant
15. Diagnosis - Ischaemic heart disease
16. Diagnosis - Lung transplant
17. Diagnosis - Mitral valve disease
18. Diagnosis - Other
19. Diagnosis - Other Cardiac
20. Diagnosis - Other Congenital
21. Diagnosis - Other Thoracic
22. Diagnosis - Peripheral vascular disease
23. Diagnosis - Pneumonia
24. Diagnosis - Renal failure
25. Diagnosis - Sepsis
26. Diagnosis - VSD
27. Re-admission (re-admission)
28. Operative Status (operative)

The summary classification table resulting from applying the equation to the records in the Combined AICU Database is shown in Table 10.07 below.

		Predicted (Pre-Admission)		
		TISS Cluster 1	*TISS Cluster 2*	*Percent Correct*
Observed	*TISS Cluster 1*	1,362	1,340	50.41%
	TISS Cluster 2	734	2,049	73.63%
			Overall Correct	62.19%

Table 10.07. Summary classification table for TISS Cluster prediction

As can be seen from Table 10.07above, the average percentage of bed-slot categorisations which are made correctly is 62.19%. Intuitively, this figure seems too low to be able to justify the further categorisation of bed-slots according to the TISS Cluster categories. To be able to be more confident in making this judgement it is necessary to weigh the costs of implementing the proposed classification against the costs involved in making such a high proportion of misclassifications.

If it is not considered costly to misclassify patients insofar that the re-allocation of the required bed-slot does not involve great cost either clinically or economically, then it becomes feasible to implement a 'rollling' system of bed-slot allocation and subsequent re-allocation, all within the same unit, as patients' clinical characteristics change over time. This is an important consideration since it changes the nature of the categorisation process. In this case, categorisation is based not only on the patients' clinical and demographic characteristics as collected pre-consumption, but also on those characteristics collected during the time the patients was consuming previous bed-slots.

Using this rolling system of categorisation a further binary logistic regression was performed on the records of the Combined AICU Database. This time, however, the previous day's values for each TISS Variable were used in addition to all of the pre-consumption characteristics used in the previous regression analysis. The resulting summary classification table is shown in Table 10.08 below.

		Predicted (Post-Admission)		
		TISS Cluster 1	*TISS Cluster 2*	*Percent Correct*
Observed	*TISS Cluster 1*	1,504	427	77.89%
	TISS Cluster 2	380	1,695	81.69%
			Overall Correct	79.86%

Table 10.08. Summary classification table for TISS Cluster prediction with rolling system of categorisation

As can be seen from Table 10.08 above, using the rolling system of categorisation, nearly 80% of patients are allocated the correct bed-slots. This figure seems very good, although before any firm conclusions may be made, two points need to be considered.

First, the rolling system of classification may, by definition, be applied only to those patients whose Consumed AICU Bed-slots measurement is greater than one, since classification requires the patient to have already consumed one bed-slot to be able to use the resulting consumption characteristics as the basis of classification. This consideration results in a significant proportion (27%) of bed-slots thus being excluded. For these bed-slots, classification can only be on the basis of pre-consumption data, which reduces the overall percentage correctly classified accordingly.

Second, and more importantly, the summary classification statistics needs to be weighted against the cost of misclassification. While it is not possible here to assign values to the cost of misclassifying patients' bed-slot requirements, it is reasonable to make the following claim: that the cost of misclassification will be proportionate to the reduction in the variance of overall resource consumption between the two bed-slot categories from the case where no classification of bed-slots is considered. The reasoning for this claim is based on the assumption that the advantages of bed-slot categorisation is, clinically, to be able to better identify particular treatment requirements and, economically, to be able to better target resources. It follows from this that if the variance in overall levels of resource consumption within each of the different bed-slots categories is smaller than it would be otherwise, these advantages would be greater than otherwise.

With regards, therefore, to the bed-slot classification in terms of TISS Cluster categories, because the reduction in variances of overall levels of resource consumption is relatively large (in comparison to the other proposed bed-slot classifications), it is reasonable to make the claim that the cost of misclassification in this case will similarly be relatively large.

Both of these considerations - the exclusion of patients with only pre-consumption characteristics available to be used as the basis of classification, and the relatively high cost of misclassification - together make the figures shown in Table 10.08 look less able to disconfirm the Variance Hypothesis than initial appearances may suggest. Equally, however, they may not reasonable be taken as supportive evidence of the hypothesis.

10.2.4. The Difference Hypothesis

The Difference Hypothesis states that there is no significant difference in the overall level of consumption of those resources which are modelled as components of the bed-slot between different categories of patients. Alternatively, that where there is a difference between different categories, it may not be substantially reduced through the further categorisation of patients into different categories where such categorisation may be made in advance of consumption within an acceptable degree of accuracy.

The Difference Hypothesis is closely related to the Variance Hypothesis in that the latter determines the validity of the Bed-slot Assumption on the basis of differences in the variance of overall resource consumption between different proposed bed-slot categories; the former on the basis of differences in the mean of overall resource consumption between different proposed bed-slot categories. Although closely related, however, the confirmation of either hypothesis by itself does not constitute necessary and sufficient validation of the Bed-slot Assumption. Consider the case where the Variance Hypothesis is disconfirmed and the Difference Hypothesis is confirmed. In this scenario, it is quite possible for the variance in overall consumption within each resulting bed-slot category to cause such a large overlap in the ranges of overall consumption of each category as to make the notion of discriminating between the categories on the basis of overall consumption clinically or economically unjustifiable.

METHOD

As with the Variance Hypothesis, the measures of overall resource consumption to be used are the three overall TISS measures: Total TISS Interventions, Total TISS Score and TISS Ratio. However, when testing the strength of the relationship The measures of difference in overall resource consumption between different categories will be based on the F statistic and the area under the receiver-operator characteristic (ROC) curve, although for the latter only Total TISS Score will be used on the assumption that similar ROC curves would equally be generated by either of the other TISS variables.

The justification for using the F statistic rather than the derived measure of significance is the level of sensitivity in values of the F statistic which is often lost using the measure of significance. All of the differences in means for each category are after all, as can be seen in Table 10.03, Table 10.04 and Table 10.05 above, would be interpreted as being statistically significant. The justification for using the additional measure of the area under the ROC curve is that this is able to give a broader picture of the relationship between overall resource consumption and proposed bed-slot category. The area under the ROC curve gives an indication of the extent to which the proposed classification scheme results in mutually discriminable categories where discrimination is based on the post hoc overall level of resource consumption. Given that the extent to which a classification generates mutually discriminable categories is dependent on both the variances and differences in the measure of overall resource consumption, the area under the ROC curve may thus be interpreted as providing supportive evidence or otherwise for both the Variance Hypothesis and the Difference Hypothesis.

The determination of whether or not the difference between the mean overall level of resource consumption between different categories is large enough to disconfirm the Difference Hypothesis will based on an intuitively plausible criterion. The classifications which will be considered are Chronicity, Operative Status and the TISS Cluster classification which was used in the discussion of the Variance Hypothesis above.

RESULTS

Looking first at the proposed classification based on Chronicity, there can be no doubt that there is a statistically significant difference in overall level of consumption between chronic and non-chronic bed-slots. As patients stay longer in the AICU, there physiological state normally improves in those cases where the patients are discharged alive. On the reasonable assumption, therefore, that there is a strong positive correlation between physiological state and resourcing requirements, chronic bed-slots should be characterised by lower levels of overall consumption than non-chronic bed-slots.

The evidence for such a relationship is shown in Table 10.09 below. Table 10.09 shows the relationship between overall consumption and the number of consumed and remaining bed-slots. As hypothesised, there is a significant negative correlation between the overall level of consumption and the number of consumed bed-slots, and a significant positive correlation between overall level of consumption and remaining bed-slots.

TISS Variable	Statistic	Consumed AICU bed-slots	Remaining AICU Bed-slots
Total TISS Interventions	Pearson Correlation	-0.347	0.154
	Sig. (1-tailed)	<0.001	<0.001
Total TISS Score	Pearson Correlation	-0.346	0.173
	Sig. (1-tailed)	<0.001	<0.001
TISS Ratio	Pearson Correlation	-0.329	0.184
	Sig. (1-tailed)	<0.001	<0.001

Table 10.09. Relationship between overall consumption and the number of consumed and remaining AICU bed slots

While Table 10.09 above demonstrates a clear relationship between overall level of resource consumption and Consumed AICU bed-slots and Remaining AICU bed-slots, and Table 10.03, Table 10.04 and Table 10.05 above shows a statistically significant difference between mean overall level of resource consumption and Chronicity, the question which needs to be asked with reference to the Difference Hypothesis is whether or not this relationship is strong enough to disconfirm the hypothesis.

Similarly, with the proposed classification of bed-slot according to Operative Status. Table 10.03, Table 10.04 and Table 10.05 above shows that the difference in mean overall level of resource consumption between Operative Status categories is statistically significant. But again, this does not imply that the difference is large enough to be considered as disconfirming the Difference Hypothesis which requires not statistical significance but clinical or economic significance.

Table 10.10 below shows the difference in mean overall levels of consumption for each category of each proposed bed-slot classification for each TISS variable, expressed as a percentage of the mean for whole population. It also shows the associated F statistic for each category of each proposed bed-slot classification for each TISS variable.

	Difference as percentage of mean			**F statistic**		
TISS Variable	Chronicity	Operative Status	TISS Cluster	Chronicity	Operative Status	TISS Cluster
Total TISS Interventions	7.6	7.1	27.4	167	131	3894
Total TISS Score	9.0	8.8	39.0	125	107	4502
TISS Ratio	2.2	2.5	13.0	51	57	2813

Table 10.10. Difference as percentage of mean for Chronicity, Operative Status as TISS Cluster and each aggregate TISS Variable

As may be seen from Table 10.10 above, the difference in mean overall level of consumption expressed as a percentage of the mean for the whole population is less than 10% for both of the pre-consumption classifications - Chronicity and Operative Status. It does not seem plausible to claim that this is a large enough difference to be able to justify the further categorisation of AICU bed-slots according to either of these classifications.

In the case of the TISS Cluster classification, however, the difference in mean overall level of consumption expressed as a percentage of the mean for the whole population is much greater. In particular, for overall level of consumption measured as Total TISS Score, the figure is 39%. It would seem implausible to claim that this does not represent a large enough difference between categories to justify the further classification of AICU bed-slots.

The interpretation of these results are further supported by looking at the area under the ROC curve for each proposed classification and overall level of resource consumption. Table 10.11 below shows the area under the ROC curve for each proposed classification and Total TISS Score. The ROC curves themselves are shown in Figure 10.01, Figure 10.02 and Figure 10.03 below.

Classification	**Area**
Chronicity	0.583
Operative Status	0.577
TISS Cluster	0.91

Table 10.11. Area under ROC curve for classification based on Chronicity, Operative Status and TISS Cluster

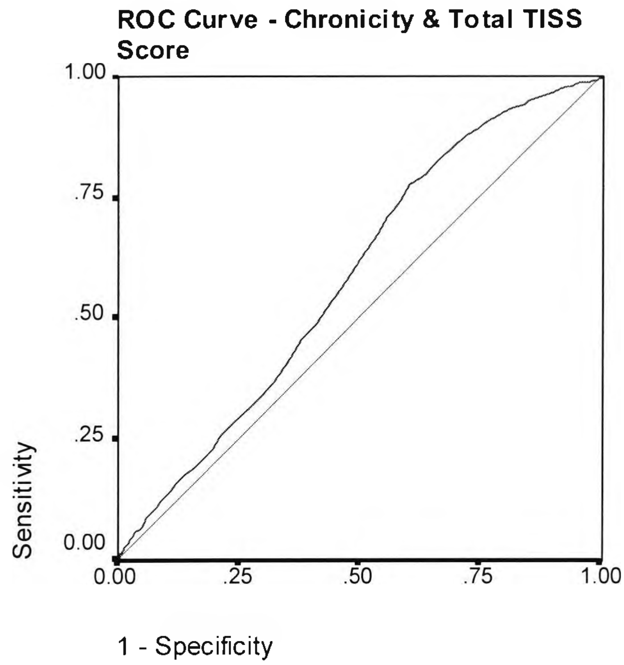

Figure 10.01. ROC Curve for Chronicity and Total TISS Score

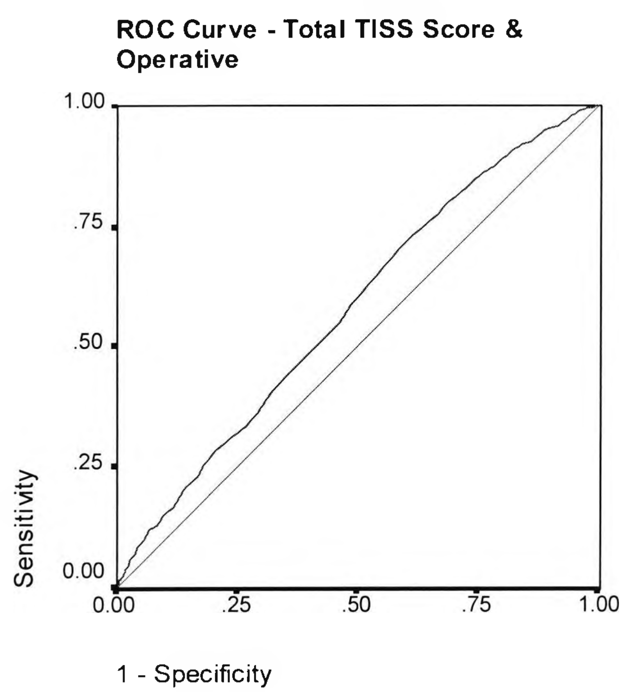

Figure 10.02. ROC Curve for Operative Status and Total TISS Score

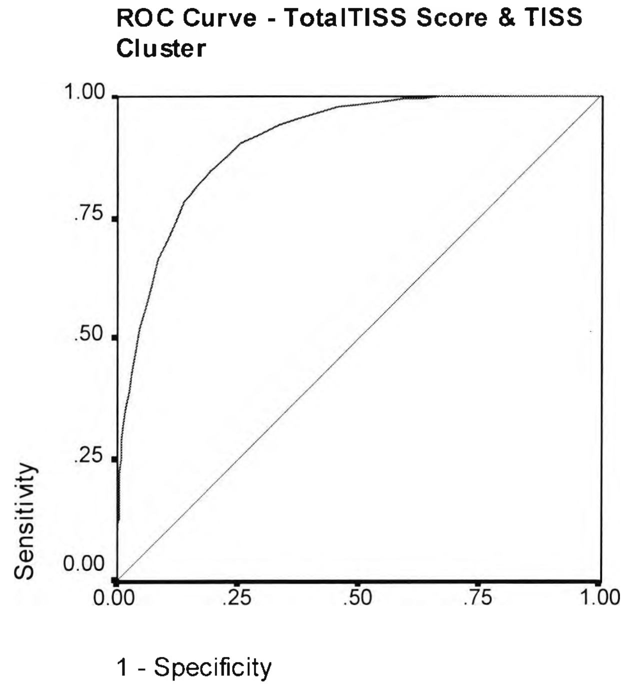

Figure 10.03. ROC Curve for TISS Cluster and Total TISS Score

As may be seen from Table 10.11, Figure 10.01 and Figure 10.02 above, the classifications based on Chronicity and Operative Status do not result in a very strong relationship between the proposed bed-slot categories and post-hoc overall resource consumption, with both areas under the ROC curve not being much greater than 0.5 which signifies no relationship between the two variables. This is in contrast to the classification based on TISS Cluster which, as should be expected given how the classification was derived, results in a much larger area under the ROC curve.

However, as with the Variance Hypothesis, for the classification based on TISS Cluster to disconfirm the Difference Hypothesis, the classification of bed-slots into TISS Cluster categories must be made within an acceptable degree of accuracy. As the results and discussion of the binary logistic regression above shows, however, this is not able to be done at the pre-consumption stage. Further, the rolling system of bed-slot categorisation, although more accurate, needs to make the assumption that the cost of misclassification is relatively low to be considered as a plausible disconfirmation of either the Variance Hypothesis or the Difference Hypothesis.

10.3. General Discussion

The aim of this study was to provide evidence which justified the use of the Bed-Slot Assumption as a useful, simplifying abstraction in modelling resource allocation in healthcare. To achieve this, the Bed-Slot Assumption was broken down into four different claims which were represented as hypotheses in the context of this study. The first of these hypotheses, The Predictability Hypothesis was proven as a theorem, while the others - The -G&F Proportion Hypothesis, the Variance Hypothesis and the Difference Hypothesis - were evaluated using data from the AICU Combined database.

From a conceptual perspective, the Predictability Hypothesis is the most important as it provides the conceptual framework in which the results of the other hypotheses are interpreted. The Predictability Hypothesis says, in effect, that whether or not a resource should be classified as generic or specific is determined by the accuracy of which the actual pattern of resource consumption may be predicted, rather than the pattern itself. Thus, the 'generic' in the expression 'generic resource' is, at least from the viewpoint of resource allocation, best seen as an epistemological concept, implying quantification over knowledge about resources at different times rather than the resources themselves.

According to the Predictability Hypothesis, each of the other hypotheses were framed so that the hypothesis could be disconfirmed if a proposition - that a patient will consume a particular resource or bed-slot - was shown to be predictable in advance of consumption. However, the extent to which something was predictable was determined as a matter of degree, with no clear threshold value beyond which the hypothesis would be considered as disconfirmed.

Only in the case of the -G&F Hypothesis was the issue of predictability not addressed, as it was shown that -G&F resources did not constitute a high enough proportion of all resources to disconfirm the hypothesis. In the case of the Variance Hypothesis and the Difference Hypothesis, it was argued that the categorisation based on the TISS Cluster classification satisfied the first criteria for disconfirming the hypotheses - that the classification resulted in a significant reduction in variance or a significant difference in overall levels of consumption between categories, respectively. In these cases, it was therefore necessary to show that the categorisation of bed-slots according to the TISS Cluster classification was able to be made at the pre-consumption stage and within an acceptable degree of accuracy.

The question of whether or not this was shown in the study remains, to some extent, open. The level of pre-consumption predictability was relatively low with 62% of bed-slots correctly classified, in which case it is implausible to claim this is sufficient to justify the further classification of AICU bed-slots according to TISS Cluster. However, a further system of prediction was used which used the previous bed-slot's component TISS variables as predictive variables along with pre-consumption variables. This rolling system of classification improved predictability to 80% of bed-slots correctly classified.

The question remained, however, whether this was accurate enough. It was argued that this could only be answered if it was possible to assign values to the cost of misclassification of bed-slots. It was further argued that the cost of misclassification in this case would be relatively high, although this was the most that could be said without being able to derive specific values, and that until such values could be assigned the two hypotheses should best be considered as remaining unconfirmed. Still, it could be plausibly claimed that, on the basis of the results thus far, that the balance of evidence is in favour of the hypotheses.

In summary, then, the Predictability Hypothesis has been shown to be true a priori, and the -G&F Hypothesis has been shown to be supported by the available data. Thus, taking these two hypotheses together, it has been shown that a general AICU bed-slot type can be identified. The Variance Hypothesis and the Difference Hypothesis have not been shown to be conclusively validated by the available data, although the balance of evidence is in their favour. Thus, it can be concluded that while there exists an AICU bed-slot type, it may not be conclusively said that this bed-slot type should be used as the basis of resource allocation instead of subtypes based, for example, on the TISS Cluster classification without a more in-depth cost-benefit analysis of adopting a further bed-slot classification than may be provided here.

11. Appendix: Operating Theatre Utilisation Analysis

11.1. Introduction

The purpose of this study is to evaluate the system of resource allocation currently in place at RBH. The evaluation criteria will be defined according to the evaluation model presented in Chapter 2. That is, the evaluation criteria will be measures of the effectiveness of the control over the allocation of resources, which in turn is considered to be closely related to the cost-effectiveness of healthcare delivery. The outcome of the evaluation will be the identification of control-limiting factors in the resource allocation process. These control-limiting factors will be classified according to whether or not they are intrinsic to the patient population or whether they are control-limiting factors of the resource allocation process.

The structure of this study is as follows. In the next section, the method which will be used in the study will be detailed, along with a description of the variables used and the various RBH databases that provided the raw data. The presentation of the results is provided in the following three sections according to the method outlined below. After a brief summary of the results, the main findings of the study are discussed according to the evaluation model developed in Chapter 2. Finally, the conclusions are summarized and discussed in relation to the central hypothesis of this thesis – that the cost-effectiveness of healthcare delivery may only be improved if managers have greater control over the allocation of healthcare resources to patients, and that such control can only realistically be achieved through the re-engineering and computerisation of the resource-allocation process.

11.2. Method

In the model presented in Chapter 2, the basis of evaluating the effectiveness of control was by comparing the actual performance of the resource allocation process with what the performance would be in a control scenario. The control scenario which was suggested in Chapter 2 was the situation where there was no control over resource allocation, where this was hypothesized as being the outcome of using the characteristics of the demand distribution as those of the consumption distribution that would occur assuming an infinite supply of available bed-slots at any one time, and then modifying the consumption distribution under the assumption of a finite supply of bed-slots at any one time. The comparative evaluation of the two cases – the actual performance versus the performance of the hypothesized control scenario – would be made on the assumption that if the standard deviation of the actual consumption distribution was less than that of the consumption distribution of the control scenario, then this would represent an increase in the effectiveness of control over the resource allocation.

The method which has been adopted in this study is that suggested in Chapter 2, with the exception that the control scenario is different. In this study, the control scenario to be used is that of the consumption distribution which would result if the intended bed-slot allocation decisions are not subsequently rescinded in the event of unintended events which imply the need for changing existing allocations. Since bed-slot allocations may only be changed prior to the initiation of the consumption of those bed-slots, as discussed in Chapter 2, this means that the control distribution is that which would result from non-monotonic bed-slot allocations. Conversely, the test distribution is that which results from monotonic bed-slot allocations.

This comparison of the intended bed-slot consumption distribution with the actual bed-slot distribution is especially pertinent to the aims of this thesis as it allows for the identification of different types of control-limiting factors. In Chapter 2 a distinction was made between epistemological and non-epistemological control-limiting factors. That is, the distinction between those control-limiting factors which arise through inherent limitations in the healthcare system such as, for example, different categories of patient being able to be admitted to the operating room suite on particular days of the week because of the availability of a surgeon only on those days of the week, and those factors which arise through a lack of knowledge about patients' resourcing requirements. In the former case, such factors will be accounted for in all bed-slot allocation decisions whether monotonic or non-monotonic, in the latter case, however, such factors can only be accounted for in monotonic bed-slot allocation decisions since it is only when the patient has actually consumed the allocated bed-slots that their resourcing requirements are known for certain.

The holding constant of non-epistemological control-limiting factors between the control and test bed-slot allocations is especially pertinent to the aims of this thesis since the main hypothesis states that increases in the effectiveness of control of resource allocation may be achieved through the computerisation and re-engineering of the existing process of resource allocation. Thus, since such computerisation and re-engineering is confined to the resource allocation process and does not extend to prescribing changes in the resources themselves, and since the resource allocation process is an informational process, any increases in the effectiveness of control gained through the computerisation and re-engineering of the process will be through the reduction in the size of the effect of epistemological rather than non-epistemological control-limiting factors.

In this section a method will be outlined for the comparison of the control and test data sets and the subsequent identification and classification of control-limiting factors. The method will be in two parts. The first develops a within-category based comparison of the two data sets. That is, a comparison based on variation in performance between different categories of bed-slot allocation in each data set. In this case two categories will be used corresponding to the month of the study period in which the bed-slot was scheduled for consumption and the day of the week in which the bed-slot was scheduled for consumption. The second part of the method makes no within-category comparison, instead looking at the overall performance characteristics of the two data sets. What follows, then, is a more detailed description of the two parts of the method, starting with the first part.

The first part of the method aims at identifying control-limiting factors that result in variation (that is, variance in the case of numeric variables; significant differences in between category frequencies in the case of categorical variables) in the values of performance variables between the different months covered by the study period or between the different days of the week (Monday to Friday), where that variation cannot be accounted for by the 'environmental' variation in the total number of allocated bed-slots per month or per week. The performance variables which will be used in the study are detailed below and are all either direct or indirect measures of workload.

In those cases where the variation in the values of a performance variable cannot be accounted for by variation in the total number of allocated bed-slots, a control-limiting factor will be assumed to be present. The determination of the type of control-limiting factor whose presence is indicated will be made on the basis of a comparison of the two data sets. This comparison can give results which are more pertinent to actually improving the resource allocation process. More specifically, comparing the same hypotheses for the same variable in each dataset can indicate the presence or absence of epistemological (E) or non-epistemological ($\neg E$) control-limiting factors. This claim requires making a fundamental assumption: that the optimal operating scenario that may be achieved within the system constraints is expressed by the non-monotonic resource allocation data. This assumption is justified by the following line of reasoning:

1. The admissions manager has full knowledge of all constraints that are present in the system; and
2. The admissions manager has full knowledge of all $\neg E$ factors that are anticipated within any given period covered by non-monotonic bed-slot allocation decisions; and
3. The admissions manager schedules patients according to their predicted healthcare requirements, and
4. the predicted healthcare resources which will be available.

Given access to this information, the admissions manager should therefore be able to schedule admissions such that the performance of the system that would result should all resource-allocation decisions not be subsequently rescinded is optimal relative to the presence of both system constraints and E factors.

Naturally, this assumption is an idealisation. Nevertheless, it remains possible to perform an evaluation by making the reasonable assumption that, although sum-optimal, it remains closer to the optimum for at least some of the performance variables than the situation for the monotonic allocation of resources.

To make the comparison between the two datasets, the following hypothesis, H_1, is tested for each variable:

H_1 = There is significant variation in the mean values/sum frequencies of variable V_{js} between the different months of the study period or the days of the week, which cannot be accounted for by random variation, where 'V_j' denotes one of the above performance variables, and 'S' denotes one of the above scheduling status categories.

Defining the null hypothesis, H_0, as the negation of H_1, then there are four possible outcomes for each variable V_j, as shown below, where 'H_i=T(V_{jS})' denotes that hypothesis H_i is confirmed for variable V_{jS}; 'V_{jn}' denotes variable V_j within the data set of non-monotonically scheduled bed-slots; 'V_{jm}' denotes variable V_j within the data set of monotonically scheduled bed-slots.

Eq.11.01 $H_1 = T(V_n^j)$ and $H_0 = T(V_m^j)$;

Eq.11.02 $H_0 = T(V_n^j)$ and $H_1 = T(V_m^j)$;

Eq.11.03 $H_0 = T(V_n^j)$ and $H_0 = T(V_m^j)$;

Eq.11.04 $H_1 = T(V_n^j)$ and $H_1 = T(V_m^j)$.

Relating these combinations to the identification of epistemological and non-epistemological control-limiting factors, it is reasonable to infer that whenever the hypothesis H_0 is confirmed for non-monotonically scheduled bed-slots, then this indicates the absence of a non-epistemological control-limiting factor. The reasoning behind this inference is that the presence of a non-epistemological control-limiting factor would be indicated by there being variation in measures of implied workload for non-monotonically scheduled bed-slots between different months or days of the week which could not be accounted for by 'noise' in the system. Negating this assumption therefore results in the above inference. The presence of an epistemological control-limiting factor can only be definitively confirmed if a) there is no non-epistemological control-limiting factor, and b) the hypothesis H_1 is confirmed for monotonically scheduled bed-slots. The reasoning behind this inference is that the presence of a control-limiting factor of either type would be indicated by there being variation in measures of workload for monotonically scheduled bed-slots between different months or days of the week which could not be accounted for by 'noise' in the system. Therefore, if the presence of a non-epistemological control-limiting factor is ruled out on the basis of the hypothesis H_0 being confirmed for non-monotonically scheduled bed-slots, then this means that any control-limiting factor has to be epistemological. In general, the identification of control-limiting factor may be made according to the following inferences:

Eq.11.05 $H_1 = \mathrm{T}(V_n^j) \rightarrow \exists f(\neg Ef)$

Eq.11.06 $[H_0 = \mathrm{T}(V_n^j) \wedge H_1 = \mathrm{T}(V_m^j)] \rightarrow \exists f(Ef)$

Eq.11.07 $H_0 = \mathrm{T}(V_n^j) \rightarrow \neg\exists f(\neg Ef)$

Eq.11.08 $[H_1 = \mathrm{T}(V_n^j) \wedge H_1 = \mathrm{T}(V_m^j)] \rightarrow [\exists f(Ef \vee \neg Ef)]$

Eq.11.09 $H_0 = \mathrm{T}(V_m^j) \rightarrow \neg\exists f(Ef)$

It will be noted that these inference rules are not able to determinately identify either E or $\neg E$ control-limiting factors in many cases. In fact, there are three possibilities when the double null hypotheses are ignored, as shown in Table 11.01 below, where each combination is numbered 1 though to 3:

Non-Monotonic	Monotonic	Combination
H_0=T(V_n)	H_1=T(V_m)	C1
H_1=T(V_n)	H_0=T(V_m)	C2
H_1=T(V_n)	H_1=T(V_m)	C3

Table 11.01. Classification of control-limiting factors based on the presence or absence of control-limiting factors

According to the expressions above, the combination C1 positively identifies E control-limiting factors and C2 and C3 positively identifies $\neg E$ control-limiting factors.

Combination C2 should, in theory, exist only as a mathematical construct since it would indicate that actual performance evidences more effective control over resource allocation than was intended in the non-monotonic allocation of bed-slots. If this situation were not to come about by chance (which remains a mathematical possibility given the nature of epidemiological proof), the only explanation that could be given to such a scenario would be the existence of some other control process complementary to, but independent of, the resource allocation process described in this chapter (see Section 3.5 below).

Combination C3 is most likely to indicate the presence of non-epistemological factors, although these may also exist undetected in other combinations. This possibility can be largely excluded if one makes the assumption that the presence of a non-epistemological factor would be typified by both an control-limiting factor being present in each data set, and that the strength of the effect is similar in each case.

Throughout this study, the confirmation or disconfirmation of H_0 and H_1 has been determined by the level of significance measured by the p-value. However, confirmation of H_0 or H_0 based on a simple binary correlation with the statistical confirmation is too blunt. Therefore, in this study a three-valued logic is used according to three value-ranges of the p-value, in which case, presence of each of the above three possible combinations can be positively excluded, positively verified or non-positively verified. In Table 11.02 these three valuations are denoted by the numbers 0, 1 and 2 respectively:

Scheduling Status		Combination		
Non-monotonic	**Monotonic**	**C1**	**C2**	**C 3**
p<0.05	p<0.05	0	0	2
p<0.05	0.1>p>0.05	0	1	1
p<0.05	p>0.1	0	2	0
0.1>p>0.05	p<0.05	1	0	1
0.1>p>0.05	0.1>p>0.05	0	0	0
0.1>p>0.05	p>0.1	1	0	1
p>0.1	p<0.05	2	0	0
p>0.1	0.1>p>0.05	1	0	1
p>0.1	p>0.1	0	0	2

Table 11.02. Classification of control-limiting factors based on levels of statistical significance

The statistical tests that will be used in the identification of control-limiting factors will be dependent on the nature of the performance variable. For numeric variables, a normal distribution will be assumed and the range of statistical tests used will correspondingly be restricted to parametric type tests. For numeric variables, the hypothesis which will be tested is that there is a significant difference in the mean workload per procedure between the different months covered by the study period or between different days of the week. For categorical variables, the hypothesis which will be tested is that there is a significant difference in the proportions of bed-slot allocations of the different categories between different months or days of the week than would be expected by the total numbers of bed-slot allocations of each category and the total number of bed-slot allocations for each month or each day of the week. In addition to these hypotheses, the hypothesis that there are differences in the mean number of bed-slot allocations per day or workload per day between different months or day of the week will also be tested for each numeric and binary categorical performance variable.

The method used to identifying the different types of control-limiting factor will be to take each data set individually, starting with the data set of non-monotonic bed slot allocations. For each data set, after a preliminary presentation of frequencies, the performance variables will be considered in turn first testing for artefacts between the different months covered by the study period, second testing for artefacts between the days of the week. The identification and classification of the different types of control-limiting factors will be undertaken after the summary of results has been presented.

The second part of the method involves a direct comparison of the two data sets. Here the aim is to identify control-limiting factors which exist between the two categories of bed-slot allocation (monotonic and non-monotonic) rather than within them. The main hypothesis that will be tested for each numeric and binary categorical performance variable in this part of the method is that there are differences in the variance of the distribution for bed-slot allocations per day or workload per day between monotonic and non-monotonic bed-slot allocations. In addition to this the hypothesis that there are differences in the mean number of bed-slot allocations per day or workload per day between monotonic and non-monotonic bed-slot allocations will also be tested. Further, for numeric variables, the hypothesis will be tested that there is a significant difference in the mean workload per procedure between the monotonic and non-monotonic bed-slot allocations. For categorical variables, the hypothesis will be tested that there is a significant difference in the proportions of bed-slot allocations of the different categories between monotonic and non-monotonic bed-slot allocations than would be expected by the total numbers of bed-slot allocations of each category and the total number of bed-slot allocations for monotonic and non-monotonic bed-slot allocations.

The statistical tests which will be used in the second part of the method will be dependent on the nature of the performance variable and the statistical characteristic of the variable which is being tested. The test for difference in variance between different distributions that will be used is Levene's test for homogeneity of variance. Although a normal assumption will be assumed in each case, it is worth noting that this test does not assume normality. In addition to testing for homogeneity of variance, Independent Samples T-Test will also be used for comparison of mean workloads between the two data sets.

11.3. Data

The data used in this study is for the period 1st August 1998 to 31st March 1999 inclusive. The data was initially divided into two data sets – one containing all of the monotonic bed-slot allocations, the other containing all of the non-monotonic bed-slot allocations. These two data sets were then combined into one data set where the additional categorical variable Scheduling Status was evaluated for each record according to whether the bed-slot allocation was monotonic or non-monotonic. Each record in each data set corresponds to the bed-slot allocations made for a single patient. Only those patients allocated a bed-slot in the operating room suite are included in the study. The variables included in the study are summarised in Table 11.03 below.

Variable	**Type**	**Range**
AICU Bed-slot Allocation Status	Categoric	(Allocated, Not Allocated)
Fast-track Status	Categoric	(Fast-track, Non Fast-track)
Number AICU Bed-slots Allocated per Day	Numeric	N
Number Fast-tracks per Day	Numeric	N
Number Paediatrics per Day	Numeric	N
Number PICU Bed-slots Allocated per Day	Numeric	N
Number Urgent Bed-slots Allocated per Day*	Numeric	N
Operating Room	Categoric	(1, 2, 3, 4, 5)
Operative Category	Categoric	(Cardiac, Thoracic, Other)
OR Workload per Day	Numeric	R
OR Workload per Procedure	Numeric	R
Patient Type	Categoric	(Adult, Paediatric)
PICU Bed-slot Allocation Status	Categoric	(Allocated, Not Allocated)
Post-operative Location	Categoric	(Adult Ward, Paediatric Ward, RR, AICU, PICU)
Pre-operative Location	Categoric	(Adult Ward, Paediatric Ward, RR, AICU, PICU)
RR Workload per Day	Numeric	R

RR Workload per Procedure	Numeric	R
Urgency*	Categoric	(Urgent, Non-Urgent)

*Variables only included in monotonic bed-slot allocations data set.

Table 11.03. Performance variables used in the study

It will be assumed that most of the variable names are sufficiently descriptive to require no further explanation. Each variable in the above table can be considered as being a candidate in the performance analysis of the resource allocation system in that they either measure overall levels of resource utilisation in a unit, or they provide information as to how this level of utilisation comes about. This is obviously true in the case of the variables RR Workload per Procedure, OR Workload per Procedure, RR Workload per Day and OR Workload per Procedure. However, the variables AICU Bed-slot Allocated, Number AICU Bed-slots Allocated per Day, Number PICU Bed-slots Allocated per Day, PICU Bed-slot Allocated, Post-operative Location and Pre-operative Location can also be considered as direct measures of resource utilisation in that they signify the allocation of bed-slots in other units. All of the other variables, while not being general measures of performance, nevertheless are important considerations in any evaluation of the resource allocation system in that they provide a breakdown of the consumption of bed-slots. For example, Operating Room breaks down the OR Workload per Procedure or OR Workload per Day in terms of the proportion of the total bed-slots that were from each operating room.

The variables RR Workload per Procedure and RR Workload per Day estimate the number of RR bed-slots that are allocated per procedure/day. The number of RR bed-slots is assumed to be proportional to the length of time required to recover patients from the procedure performed. All major cardiac (i.e. fast-track) procedures were assumed to require 4 times as much work than all other procedures, in which case one RR bed-slot was assumed to be equivalent to the amount of resources needed to recover a non-major cardiac patient. Similarly, the variables OR Workload per Procedure and OR Workload per Day estimates the number of OR bed-slots that are allocated per procedure/day. All minor cardiac and all thoracic procedures were assumed to require half the work as all major cardiac procedures, in which case one bed-slot was assumed to be equivalent to the amount of resources needed to perform a non-major cardiac procedure. This calibration for RR and OR bed-slots may seem to be very coarse and discreet for a variable, which in reality is a continuous measurement with a wide range of values. Nevertheless, the values which have been used serve as an initial approximation which is sensitive to different types of procedures, and in that sense is a more accurate measure than simply looking at procedure frequency. Moreover, with OR resources being allocated as blocks of time which are based on the same calibration, the evaluation is actually quite accurate.

In measuring the variables for the number of bed-slots allocated for AICU and PICU, a bed-slot is only considered as being allocated if it results in a de novo admission to the unit. Measuring all admissions could not be considered as an accurate measure of resource utilisation since those patients which originate in AICU/PICU and are then re-admitted a few hours later upon discharge from OR would not have been out of AICU/PICU long enough to have effectively freed bed-slots within those units for another admission.

Each of the variables listed above represent one dimension along which the performance of a resource allocation system implemented at RBH may be evaluated. As argued in Chapter 2, the overall measure of the cost-effectiveness of healthcare delivery is a composite measure involving both measures of operating efficiency and effectiveness. As such, the evaluation of a resource allocation system – particularly one allocating resources in a highly interdependent progressive-care system – will involve comparison with an optimal operating scenario, where such a scenario cannot be defined in terms of one variable alone. Moreover, what is considered optimal varies according to the objectives of the healthcare organisation. The evaluation which is presented here, therefore is deficient in an important respect. Namely, that it is a fundamental assumption that the optimal operating scenario is expressed by the non-monotonic resource allocation system. As will be shown, however, this assumption is undermined to some degree with the possible existence of Type 2 control-limiting factors which imply that the non-monotonic resource allocation is sub-optimal. Nevertheless, it remains possible to perform an evaluation by making the reasonable assumption that, although sum-optimal, it remains closer to the optimum for at least some of the performance variables than the situation for the monotonic allocation of resources.

In the next two sections, the two different data sets are analysed according to the first part of the method, starting with the data set of non-monotonic bed-slot allocations. Next, the two data sets are compared directly according to the second part of the method. Because of the large number of results generated by the different analyses presented here, there will be a summary of results after all of the analyses are completed, followed by a discussion of the results and a summary of the conclusions. The identification and classification of control-limiting factors will be undertaken in Section 6, Summary of Results.

11.4. Scheduling Status: Non-monotonic bed-slot allocation

The data used in this section of the study is that relating to non-monotonically allocated bed-slots in the RBH high-dependency environment. That is, bed-slots whose allocation is made at a stage in the resource allocation process where the allocation may subsequently be revoked at a later time.

The approach adopted here is to first examine the relationship between the dependent variables listed in the Method section above and the independent variable Month, which measures the month covered by the study period and in which the patient was scheduled for admission to the operating room suite. Second, the same dependent variables will be examined in relation to the independent variable Day of the Week, which measures the day of the working week (i.e., Monday to Friday) in which the patient was scheduled for admission to the operating room suite. Before either of these tasks are undertaken, however, the summary processing statistics and frequencies are presented, beginning with the processing statistics which are shown in Table 11.03 below.

	Month	**Day of the Week**	**Operating Room**	**Patient Type**	**Operative Category**
Valid	1624	1624	1619	1621	1622
Missing	0	0	5	3	2
	Pre-Operative Location	**Post-operative Location**	**Fast-track Status**	**AICU Bed-slot Status**	**PICU Bed-slot Status**
Valid	1616	1601	1624	1600	1601
Missing	8	23	0	24	23

Table 11.03. Processing statistics for non-monotonic bed-slot data set

The total number of procedures per operating room is shown in Table 11.04 below.

Operating Room	Frequency	Valid Percent
2	527	32.55
1	311	19.21
4	300	18.53
3	287	17.73
5	194	11.98
Total	**1619**	**100.00**

Table 11.04. Total number of procedures per operating room

The total number of procedures per patient type is shown in Table 11.05 below.

Patient Type	Frequency	Valid Percent
Adult	1302	80.32
Paediatric	319	19.68
Total	**1621**	**100.00**

Table 11.05. Total number of procedures per patient type

The total number of procedures per procedure type is shown in Table 11.06 below.

Operative Category	Frequency	Valid Percent

Cardiac	983	60.60
Thoracic	613	37.79
Other	26	1.60
Total	**1622**	**100.00**

Table 11.06. Total number of procedures per operative category

The total number of procedures per AICU bed-slot allocation status is shown in Table 11.07 below.

AICU Bed-slot Status	Frequency	Valid Percent
No AICU bed-slot allocation	1044	65.25
AICU bed-slot allocation	556	34.75
Total	**1600**	**100.00**

Table 11.07. Total number of procedures per AICU bed-slot allocation status

The total number of procedures per PICU bed-slot allocation status is shown in Table 11.08 below.

PICU Bed-slot Status	Frequency	Valid Percent
No PICU bed-slot allocated	1412	88.19
PICU bed-slot allocated	189	11.81
Total	**1601**	**100.00**

Table 11.08. Total number of procedures per PICU bed-slot allocation status

The total number of procedures per fast track status is shown in Table 11.09 below.

Fast-track Status	Frequency	Valid Percent
Non fast-tracked	1419	87.38
Fast-tracked	205	12.62
Total	**1624**	**100.00**

Table 11.09. Total number of procedures per fast-track status

The total number of procedures per pre-operative location is shown in Table 11.10 below.

Pre-Operative Location	Frequency	Valid Percent
Adult Ward	1307	80.88
Paediatric Ward	296	18.32
PICU	13	0.80
Total	**1616**	**100.00**

Table 11.10. Total number of procedures per pre-operative location

The total number of procedures per post-operative location is shown in Table 11.11 below.

Post-operative Location	Frequency	Valid Percent
AICU	557	34.79
Adult Ward	488	30.48
RR	249	15.55
PICU	202	12.62
Paediatric Ward	105	6.56
Total	**1601**	**100.00**

Table 11.11. Total number of procedures per post-operative location

The total number of procedures per month is shown in Table 11.12 below.

Month	Frequency	Valid Percent
Sep 98	230	14.16
Oct 98	226	13.92
Mar 99	221	13.61
Jan 99	208	12.81
Nov 98	207	12.75
Aug 98	197	12.13
Feb 99	189	11.64
Dec 98	146	8.99
Total	**1624**	**100.00**

Table 11.12. Total number of procedures per month

11.5. Scheduling Status: Monotonic bed-slot allocation

The data used in this section of the study is that relating to monotonically allocated bed-slots in the RBH high-dependency environment. That is, bed-slots whose allocation is made at a stage in the resource allocation process where the allocation may not subsequently be revoked at a later time.

The approach adopted here is, as in the preceding section, to first examine the relationship between the dependent variables listed in the Method section above and the independent variables Month and Day of the Week. Before either of these task are undertaken, however, the summary processing statistics and frequencies are presented, beginning with the processing statistics which are shown in Table 11.14 below.

	Month	Day of the Week	Operating Room	Urgency Status	Patient Type	Operative Category
Valid	1657	1657	1635	1656	1653	1654
Missing	0	0	22	1	4	3

	Pre-operative Bed-slot Allocation	Post-operative Bed-slot Allocation	Fast-track Status	AICU Bed-slot Allocation	PICU Bed-slot Allocation
Valid	1655	1650	1657	1657	1657
Missing	2	7	0	0	0

Table 11.14. Processing statistics for monotonic bed-slot data set

The total number of procedures per operating room is shown in Table 11.15 below.

Operating Room	Frequency	Valid Percent
2	531	32.48
1	323	19.76
3	288	17.61
4	271	16.57
5	222	13.58
Total	1635	100.00

Table 11.15. Total number of procedures per operating room

The total number of procedures per patient type is shown in Table 11.16 below.

Patient Type	Frequency	Valid Percent

Adult	1317	79.67
Paediatric	336	20.33
Total	1653	100.00

Table 11.16. Total number of procedures per patient type

The total number of procedures per procedure type is shown in Table 11.17 below.

Operative Category	Frequency	Valid Percent
Cardiac	980	59.25
Thoracic	641	38.75
Other	33	2.00
Total	1654	100.00

Table 11.17. Total number of procedures per operative category

The total number of procedures per AICU bed-slot allocation status is shown in Table 11.18 below.

AICU Bed-slot Allocation	Frequency	Valid Percent
AICU Bed-slot Allocation	1177	71.03
No AICU Bed-slot Allocation	480	28.97
Total	1657	100.00

Table 11.18. Total number of procedures per AICU bed-slot allocation status

The total number of procedures per PICU bed-slot allocation status is shown in Table 11.19 below.

PICU Bed-slot Allocation	Frequency	Valid Percent
No PICU bed-slot allocation	1487	89.74
PICU bed-slot allocation	170	10.26
Total	1657	100.00

Table 11.19. Total number of procedures per PICU bed-slot allocation status

The total number of procedures per fast track status is shown in Table 11.20 below.

Fast-track Status	Frequency	Valid Percent
Non fast-tracked	1489	89.86
Fast-tracked	168	10.14
Total	1657	100.00

Table 11.20. Total number of procedures per fast-track status

The total number of procedures per pre-operative location is shown in Table 11.21 below.

Pre-operative Bed-slot Allocation	Frequency	Valid Percent
Adult Ward	1250	75.53
Paediatric Ward	311	18.79
AICU	63	3.81
PICU	26	1.57
RR	5	0.30
Total	1655	100.00

Table 11.21. Total number of procedures per pre-operative location

The total number of procedures per post-operative location is shown in Table 11.22 below.

Post-operative Bed-slot Allocation	Frequency	Valid Percent

Adult Ward	574	34.79
AICU	541	32.79
RR	202	12.24
PICU	195	11.82
Paediatric Ward	138	8.36
Total	1650	100.00

Table 11.22. Total number of procedures per post-operative location

The total number of procedures per level of urgency is shown in Table 11.23 below.

Urgency Status	Frequency	Valid Percent
Non-urgent	1548	93.48
Urgent	108	6.52
Total	1656	100.00

Table 11.23. Total number of procedures per urgency status

The total number of procedures per month is shown in Table 11.24 below.

Month	Frequency	Valid Percent
Mar 99	238	14.36
Oct 98	236	14.24
Nov 98	225	13.58
Aug 98	208	12.55
Jan 99	206	12.43
Feb 99	204	12.31
Sep 98	184	11.10
Dec 98	156	9.41
Total	1657	100.00

Table 11.24. Total number of procedures per month

11.6. Results

The following summary of results is broken down into those results for a) non-monotonic bed-slot allocations; b) monotonic bed-slot allocations, and c) a comparison of monotonic and non-monotonic bed-slot allocations.

11.6.1. Scheduling Status: Non-monotonic bed-slot allocation

The summary of results for each variable tested from the data set of non-monotonic bed-slot allocations is shown in Table 11.25 below. The first column lists the dependent variables which were tested; the second and third columns list the independent variables, i.e., month and day of the week respectively.

Dependent Variable	p for Month	p for Day of the Week
AICU Bed-slot Allocation Status	$p<0.05$	$p<0.01$
Fast-track Status	NS	$p<0.01$
Number AICU Bed-slots Allocated per Day	$p<0.01$	$p<0.01$

Number Fast-tracks per Day	NS	p<0.01
Number Paediatrics per Day	NS	p<0.01
Number PICU Bed-slots Allocated per Day	NS	p<0.01
Operative Category	NS	p<0.01
OR Workload per Day	NS	p<0.01
OR Workload per Procedure	NS	p<0.01
Patient Type	NS*	p<0.01
PICU Bed-slot Allocation Status	NS	NS
RR Workload per Day	NS*	p<0.01
RR Workload per Procedure	NS*	p<0.01

'NS' = not significant; *trend towards significance (0.1 > p >0.05).

Table 11.25. Summary of results of non-monotonic data set

11.6.2. Scheduling Status: Monotonic bed-slot allocation

The summary of results for each variable tested from the data set of monotonic bed-slot allocations is shown in Table 11.26 below. The first column lists the dependent variables which were tested; the second and third columns list the independent variables, i.e., month and day of the week respectively.

Dependent Variable	p for Month	p for Day of the Week
AICU Bed-slot Allocation Status	p<0.01	NS
Fast-track Status	p<0.01	p<0.01
Number AICU Bed-slots Allocated per Day	p<0.01	NS*
Number Fast-tracks per Day	p<0.01	p<0.01
Number Paediatrics per Day	NS	NS*
Number PICU Bed-slots Allocated per Day	NS	p<0.01
Operative Category	NS	p<0.05
OR Workload per Day	p<0.05	p<0.01
OR Workload per Procedure	NS	p<0.01
Patient Type	NS	p<0.05
PICU Bed-slot Allocation Status	NS	p<0.01
RR Workload per Day	p<0.01	p<0.01
RR Workload per Procedure	p<0.01	p<0.01
Number Urgent Bed-slots Allocated per Day	NS	NS
Urgent Bed-slot Allocation Status	NS	NS

'NS' = not significant; *trend towards significance (0.1 > p >0.05).

Table 11.26. Summary of results of monotonic data set

11.6.3. Scheduling Status: Monotonic and non-monotonic bed-slot allocation

The summary of results for each variable tested from the data sets of both monotonic and non-monotonic bed-slot allocations is shown in Table 11.27 below. The first column lists the dependent variables which were tested; the second column lists the level of significance for each dependent variable against the independent variable Scheduling Status.

Variable	Scheduling Status	Mean	SD	p (Levene's Test)
Number of Paediatric Bed-slot Allocations per Day	Non-monotonic	1.98	1.31	0.074
	Monotonic	2.02	1.45	

OR Workload per Day	Non-monotonic	15.75	3.95	0.015
	Monotonic	14.99	5.18	
RR Workload per Day	Non-monotonic	9.34	4.35	0.096
	Monotonic	8.54	4.92	
Number Fast-tracked per Day	Non-monotonic	1.31	1.01	0.413
	Monotonic	1.01	1.11	
Number AICU Bed-slots Allocated per Day	Non-monotonic	3.48	1.30	0.117
	Monotonic	2.89	1.51	
Number PICU Bed-slots Allocated per Day	Non-monotonic	1.18	0.83	0.202
	Monotonic	1.02	0.94	

Table 11.27. Comparison of mean performance variables between monotonic and non-monotonic data sets

Table 11.28 below shows a comparison of means between the two datasets of monotonic and non-monotonic scheduled bed-slots. For each dependent variable the mean value is compared between the two datasets and the level of significance indicated for the comparison.

Dependent Variable	p for Scheduling Status
AICU Bed-slot Allocation Status	p<0.01
Fast-track Status	p<0.05
Number AICU Bed-slots Allocated per Day	p<0.01*
Number Fast-tracks per Day	p<0.05*
Number Paediatrics per Day	NS*
Number PICU Bed-slots Allocated per Day	NS*
Operative Category	NS
OR Workload per Day	NS*
OR Workload per Procedure	p<0.01*
Patient Type	NS
PICU Bed-slot Allocation Status	NS
RR Workload per Day	NS*
RR Workload per Procedure	NS*
Month	NS
Day of the Week	NS

'NS' = not significant; *Equal variances not assumed.

Table 11.28. Summary of results of comparisons between non-monotonic and monotonic data sets

Table 11.29 below shows the differences in mean workload measure for each (numeric) dependent variable listed in the first column between non-monotonic and monotonic bed-slot allocations. As can be seen from the table, the hypothesis that there is a difference in the mean workload between monotonic bed-slot allocations and non-monotonic bed-slot allocations is confirmed for the dependent variables Number Fast-tracked per Day, Number AICU Bed-slots Allocated per Day and OR Workload per procedure. Each of these workload measures showed a significant reduction in workload from the non-monotonic to the monotonic bed-slot allocations.

Dependent Variable	Non-Monotonic	Monotonic	Difference	p* (2-tailed)

Number Paediatrics per Day	1.981	2.024	-0.043	NS
OR Workload per Day	15.745	14.994	0.751	NS
RR Workload per Day	9.340	8.542	0.798	NS
Number Fast-tracked per Day	1.314	1.012	0.302	p<0.05
Number AICU Bed-slots Allocated per Day	3.475	2.892	0.583	p<0.01
Number PICU Bed-slots Allocated per Day	1.181	1.024	0.157	NS
OR Workload per Procedure	1.563	1.505	0.058	p<0.01
RR Workload per Procedure	0.897	0.856	0.041	NS

*p values between the two cases of equal variances being assumed and equal variances not being assumed in no case affected the overall judgements of significance. 'NS' = not significant.

Table 11.29. Summary of results of comparisons in means between non-monotonic and monotonic data sets

Table 11.30 below shows the F values and corresponding level of significance for the hypothesis that there is no difference in the variance between the workload distributions of the monotonic and non-monotonic bed-slot allocations for each numeric measure of workload. As can be seen from Table 11.30, there is no significant difference in the variance in workload distributions for the dependent variables OR Workload per Day and OR Workload per Procedure using Levene's Test for the Homogeneity of Variance.

Dependent Variable	**F**	***St.Dev.*** [a]	***St.Dev.*** [b]	**Sig.**
Number AICU Bed-slots Allocated per Day	2.476	1.31	1.45	NS
Number Fast-tracked per Day	0.672	3.95	5.18	NS
Number of Paediatric Bed-slot Allocations per Day	3.215	4.35	4.92	NS*
Number PICU Bed-slots Allocated per Day	1.634	1.01	1.11	NS
OR Workload per Day	6.016	1.30	1.51	p<0.05
RR Workload per Day	2.785	0.83	0.94	NS*
OR Workload per Procedure	26.094	0.50	0.50	p<0.01
RR Workload per Procedure	9.493	1.27	1.16	p<0.05

'NS' = not significant; * trend towards significance; a) variance for non-monotonic bed-slot allocations workload distribution; b) variance for monotonic bed-slot allocations workload distribution.

Table 11.30. Summary of results of comparisons in variances between non-monotonic and monotonic data sets

11.6.4. Identification and Classification of Control-Limiting Factors

Throughout this study two data sets have been used corresponding to whether the bed-slot allocations which constitute the records of each data set are made non-monotonically or monotonically. The analysis of either of these data sets in isolation may give indications as to the effectiveness of the resource allocation process in terms of mean workloads per month or variations in mean workloads throughout the week, and so on. However, a comparison of the two datasets as described above can provide a means to identify and classify control-limiting factors into epistemological and non-epistemological factors, thus providing an important insight into not only the areas of weakness in the operational management and control of the healthcare system, but also into the causes of those weaknesses.

Applying the system of classification developed above to the results of the study for the testing of H_1 as it applies between the different months of the study period, the results are as follows:

	Scheduling Status		Combination		
Dependent Variable	Non-monotonic BSA	Monotonic BSA	C1	C2	C3
AICU BSA	P<0.05	p<0.05	0	0	2
Fast-track Status	p>0.1	p<0.05	2	0	0
Number AICU Bed-slots Allocated per Day	P<0.05	p<0.05	0	0	2
Number Fast-tracks per Day	p>0.1	p<0.05	2	0	0
Number Paediatrics per Day	p>0.1	p>0.1	0	0	0
Number PICU Bed-slots Allocated per Day	p>0.1	p>0.1	0	0	0
OR Workload per Day	p>0.1	p<0.05	2	0	0
OR Workload per Procedure	p>0.1	p>0.1	0	0	0
PICU BSA	p>0.1	p>0.1	0	0	0
RR Workload per Day	0.1>p>0.05	p<0.05	1	0	1
RR Workload per Procedure	0.1>p>0.05	p<0.05	1	0	1

Table 11.31. Identification of control-limiting factors between different months of the study period.

Table 11.31 above shows the positive presence of E control-limiting factors for the variables Fast-track Status, Number Fast-tracks per Day and OR Workload per Day. It shows the positive presence of $\neg E$ control-limiting factors for the variables AICU Bed-slot Allocation Status and Number AICU Bed-Slots Allocated per Day.

Applying the system of classification to the results of the study for the testing of H_1 as it applies between the days of the week, the results are as follows:

	Scheduling Status		Combination		
Dependent Variable	Non-monotonic BSA	Monotonic BSA	C1	C2	C3
AICU BSA	P<0.05	p>0.1	0	2	0
Fast-track Status	P<0.05	p<0.05	0	0	2
Number AICU Bed-slots Allocated per Day	P<0.05	0.1>p>0.05	0	1	1
Number Fast-tracks per Day	P<0.05	p<0.05	0	0	2
Number Paediatrics per Day	P<0.05	0.1>p>0.05	0	1	1
Number PICU Bed-slots Allocated per Day	P<0.05	p<0.05	0	0	2
OR Workload per Day	P<0.05	p<0.05	0	0	2
OR Workload per Procedure	P<0.05	p<0.05	0	0	2
PICU BSA	p>0.1	p<0.05	2	0	0
RR Workload per Day	P<0.05	p<0.05	0	0	2
RR Workload per Procedure	P<0.05	p<0.05	0	0	2

Table 11.32. Identification of control-limiting factors between different days of the week.

Table 11.32 above shows the positive presence of E control-limiting factors for the variable PICU Bed-slot Allocation Status. It shows the positive presence of $\neg E$ control-limiting factors for the variables Fast-track Status, Number of Fast-tracks per Day, Number of PICU Bed-Slots Allocated per

Day, OR Workload per Day, OR Workload per Procedure, RR Workload per Day and RR Workload per Procedure.

11.7. Discussion

The effective control over admissions in any healthcare unit is of critical importance, both in order to contain costs, as well as to be able to satisfy patients' healthcare requirements in a timely manner. Admissions control becomes even more important in the context of progressive-care systems, where the effective control of admissions to one unit depends on the effective control of admissions in other units within the system.

In measuring the effectiveness of control over admissions to a progressive-care system, the main measures of interest relate to the extent to which the workload in each of the component healthcare units of the system can be maintained within a range of acceptable values. Wide fluctuations in workload, either between months or between different days of the working week, represents the under-utilisation of resources at one time and the over-utilisation of resources at other times, as well as a possible decrease in the quality of healthcare which is delivered.

To inform the process of improving admissions control through the development of better scheduling information systems, it is necessary to identify the nature of any artefacts that limit the extent to which the control of admissions may be optimised. It is this line of reasoning that leads to the classification of system artefacts above into artefacts and artefacts and the statistical method proposed to identify those artefacts.

The results of the analyses on workload by month suggest that, while the control of workload in the paediatric intensive-care unit is relatively good, the other units have both E and $\neg E$ control-limiting factors present, which results in fluctuations of workload between different months. In particular, it suggests that there are fluctuations in the workload of operating theatres whose primary cause is one or more artefacts. This could be due to unanticipated numbers of certain types of operative procedures being performed in particular months or an abnormal amount of urgent operative procedures in certain months, for example. In either case, the cause is epistemological in nature – there is a lack of knowledge regarding future events. This is in contrast to the cause being a problem in staff scheduling, where a disproportionate number of surgeons might choose to take their holiday in a particular month, for example, or a problem in resource planning, where an operating theatre may be closed for refurbishment in a particular month. In both of these examples, such perturbations in workload would be caused by events that are planned in advance.

The results of the analyses on workload by day of the week suggest that the control of workload in all units is affected by $\neg E$ control-limiting factors. This is to be expected, since the scheduling of many staff resources operates on a weekly basis. This is particularly the case with the scheduling of surgeons, who tend to specialise in particular types of operative procedure and tend also to perform those procedures on particular days of the week. Thus, for example, the presence of a $\neg E$ artefact in the workload by day of the week for the paediatric intensive-care unit is almost certainly due to the availability of surgeons to treat paediatric cases being unevenly distributed throughout the week. This, incidentally, may also explain the presence of $\neg E$ artefacts in the workload by day of the week for the other healthcare units, since paediatric patients are less likely than adult patients to be admitted to the post-operative recovery room, which therefore results in a reduced workload for that unit on those days when a large number of paediatric patients are admitted. Also, since paediatric patients tend to require more complex operative procedures, the workload by day of the week for the operating theatres will be similarly affected.

Thus, while the method proposed here does not answer the question as to what the cause is of a particular series of workload data, it does narrow the list of possible candidates by identifying the cause as being either epistemological or non-epistemological in nature. It is also able to form an important component in the evaluation of any scheduling information system, with the causes of any improvements in the control of admissions post-implementation of a new information system, as measured by workload data, being able to be identified by formal means.

Before the method proposed here may be used to provide input into any evaluation study, a more precise and quantitative set of decision criteria need to be developed. In particular, the measurement of statistical significance used in the pilot study described above fails to give an absolute strength of effect measurement, the alpha value of any statistical test being dependent on the sample size and the number of categories.

12. Operational Model Data Attributes

12.1. Current Operational Model Data Attributes

COP BED SLOT ID

COP Bed Slot ID is a component attribute of the Bed Slot object class and measures a unique identifier to identify individual bed-slots.

COP ACTUAL BED SLOT STATUS TIME = [T]

COP Actual Bed Slot Status Time = [T] is a component attribute of the Bed Slot object class and measures the actual scheduling status of the bed slot at time T. Projected status is a categorical measure evaluated as Unallocated-Unoccupied, Allocated-Occupied, Allocated-Unoccupied.

COP ACTUAL LABOUR COMPONENT [N] STATUS TIME = [T]

COP Actual Labour Component [N] Status Time = [T] is a component attribute of the Bed Slot object class and measures whether or not the bed slot comprised a labour component of type N during a standardised period of time beginning at time T, where 'N' could refer to nursing labour, clinician labour, and so on.

COP ACTUAL TISS COMPONENT [N] STATUS TIME = [T]

COP Actual TISS Component [N] Status Time = [T] is a component attribute of the Bed Slot object class and measures whether or not the bed slot comprised the TISS Component N during a standardised period of time beginning at time T, where N could be any TISS component which denotes the consumption of a particular healthcare resource or group of resources.

COP BED SLOT UNIT NAME

COP Bed Slot Unit Name is a component attribute of the Bed Slot object class and measures the unique identifier of the unit of which the bed slot is a component.

COP PATIENT SCHEDULED MONOTONIC BED SLOT TIME = [T]

COP Patient Scheduled Monotonic Bed Slot Time = [T] is a component attribute of the Bed Slot object class and measures the patient which has been monotonically allocated the bed slot at time T.

COP PATIENT SCHEDULED NON-MONOTONIC BED SLOT TIME = [T]

COP Patient Scheduled Non-Monotonic Bed Slot Time = [T] is a component attribute of the Bed Slot object class and measures the patient which has been non-monotonically allocated the bedslot at time T.

COP PATIENT TYPE [P] ADMISSIBLE

COP Patient Type [P] Admissible is a component attribute of the Bed Slot object class and measures whether or not the patient type P may consume the healthcare resources represented by the bed-slot.

COP PROJECTED BED SLOT STATUS TIME = [T]

COP Projected Bed Slot Status Time = [T] is a component attribute of the Bed Slot object class and measures The projected scheduling status of the bed slot at time T. Projected status is a categorical measure evaluated as Unallocated-Unoccupied, Allocated-Occupied, Allocated-Unoccupied.

COP PROJECTED LABOUR COMPONENT [N] STATUS TIME = [T]

COP Projected Labour Component [N] Status Time = [T] is a component attribute of the Bed Slot object class and measures whether or not the bed slot is projected to comprise a labour comopnent of type N during a standardised period of time beginning at time T, where 'N' could refer to nursing labour, clinician labour, and so on.

COP PATIENT HOSPITAL NUMBER

COP Patient Hospital Number is a component attribute of the Patient object class and measures A unique identifier for patients admitted to RBH.

COP ACTUAL ADMISSION TIME UNIT [U]

COP Actual Admission Time Unit [U] is a component attribute of the Patient object class and measures the patient's actual time of admission to unit U.

COP ACTUAL LENGTH OF STAY UNIT [U]

COP Actual Length of Stay Unit [U] is a component attribute of the Patient object class and measures the patient's actual length of stay in unit U.

COP ADMITTING CONSULTANT

COP Admitting Consultant is a component attribute of the Patient object class and measures the consultant clinician responsible for the patient and their treatment. In those cases where a consultant surgeon is responsible for the patient, this attribute may be set to 0 as applicable.

COP ADMITTING SURGEON

COP Admitting Surgeon is a component attribute of the Patient object class and measures the surgeon who is responsible for the patient and on whose waiting list the patient is entered for surgery. This attribute only applies to surgical patients

COP PATIENT ADMISSION DIAGNOSIS

COP Patient Admission Diagnosis is a component attribute of the Patient object class and measures the diagnosis which was made when the patient was first admitted to the hospital.

COP PATIENT CLINICAL ATTRIBUTE [N]

COP Patient Clinical Attribute [N] is a component attribute of the Patient object class and measures the set of clinical attributes which represent and evaluate the patient's physiology and pathological condition when the measurement was made. These attributes are assumed to be static and require no updating according to the progression or treatment

COP PATIENT CLINICAL ATTRIBUTE [N] TIME = T

COP Patient Clinical Attribute [N] Time = T is a component attribute of the Patient object class and measures the set of clinical attributed which represent and evaluate the patient's physiology and pathological condition at time T. These attributes may be simple attributes such as Blood Pressure at Time T, or they may be more complex attributes such as Parsonnet.

COP PATIENT CURRENT DIAGNOSIS

COP Patient Current Diagnosis is a component attribute of the Patient object class and measures the current diagnosis of the patient. It is assumed that this diagnosis may be different from the diagnosis which was made at admission. It is not derived from any other attributes' values and is edited by the processor ClinicalStaff.

COP PATIENT CURRENT LOCATION

COP Patient Current Location is a component attribute of the Patient object class and measures the current unit which is treating the patient

COP PATIENT DATE OF BIRTH

COP Patient Date of Birth is a component attribute of the Patient object class and measures the patient's date of birth.

COP PATIENT DEMOGRAPHIC ATTRIBUTE [N]

COP Patient Demographic Attribute [N] is a component attribute of the Patient object class and measures the values for a set of demographic attributes of the patient. These attributes could measure, for example, Next of Kin, Gender, Age, etc. All of these attributes are not represented individually for reasons of simpllicity.

COP PATIENT HOME ADDRESS

COP Patient Home Address is a component attribute of the Patient object class and measures the home address of the patient.

COP PATIENT HOSPITAL ADMISSION DATE

COP Patient Hospital Admission Date is a component attribute of the Patient object class and measures the date at which the patient was admitted to the hospital

COP PATIENT NAME

COP Patient Name is a component attribute of the Patient object class and measures the name of the patient

COP PATIENT PROJECTED DISCHARGE TIME

COP Patient Projected Discharge Time is a component attribute of the Patient object class and measures the most current projection of the time at which the patient is expected to be able to be discharged from the unit in which they are currently being treated.

COP PATIENT SCHEDULING STATUS

COP Patient Scheduling Status is a component attribute of the Patient object class and measures the patient's current scheduling status

COP UNIT NAME

COP Unit Name is a component attribute of the Unit object class and measures a a unique identifier for each unit.

COP ACCEPT PATIENT TYPE [P] FROM UNIT [U]

COP Accept Patient Type [P] From Unit [U] is a component attribute of the Unit object class and measures whether or not the unit may admit patients of type P from unit U.

COP ACTUAL LABOUR COMPONENT [N] PER BED SLOT TIME = [T]

COP Actual Labour Component [N] per Bed Slot Time = [T] is a component attribute of the Unit object class and measures the actual amount of the labour component N per bed slot within the unit at time T, where 'N' could refer to nursing labour, clinician labour, and so on.

COP ACTUAL NUMBER ALLOCATED-OCCUPIED BED SLOTS TIME = [T]

COP Actual Number Allocated-Occupied Bed Slots Time = [T] is a component attribute of the Unit object class and measures the total actual number of bed slots within the unit which are occupied by a patient at time T

COP ACTUAL NUMBER ALLOCATED-UNOCCUPIED BED SLOTS TIME = [T]

COP Actual Number Allocated-Unoccupied Bed Slots Time = [T] is a component attribute of the Unit object class and measures the total actual number of bed slots within the unit which are allocated to a patient, but not occupied by any patient at time T

COP ACTUAL NUMBER UNALLOCATED-UNOCCUPIED BED SLOTS TIME = [T]

COP Actual Number Unallocated-Unoccupied Bed Slots Time = [T] is a component attribute of the Unit object class and measures the total actual number of bed slots within the unit which are neither allocated nor occupied by any patient at time T

COP ACTUAL OCCUPANCY RATE TIME = [T]

COP Actual Occupancy Rate Time = [T] is a component attribute of the Unit object class and measures the proportion of bed slots within the unit at time T which are occupied.

COP ACTUAL TISS COMPONENT [N] PER BED SLOT TIME = [T]

COP Actual TISS Component [N] per Bed Slot Time = [T] is a component attribute of the Unit object class and measures the actual amount of TISS component N per bed slot within the unit which is consumed per unit time at a time T, where N could be any TISS component which denotes the consumption of a particular healthcare resource or group of resources.

COP CLINICAL DIRECTOR NAME

COP Clinical Director Name is a component attribute of the Unit object class and measures the name of the clinical director of the unit with overall responsibility for all aspects of the delivery of health care to patients within the unit.

COP DISCHARGE PATIENT TYPE [P] TO UNIT [U]

COP Discharge Patient Type [P] To Unit [U] is a component attribute of the Unit object class and measures whether or not patient type P may be discharged from the unit to unit U.

COP Maximum Actual Number Bed Slots

COP Maximum Actual Number Bed Slots is a component attribute of the Unit object class and measures the maximum number of bed slots available within the unit at any given time that may be allocated to any of the admissible patient types

COP Maximum Actual Number Bed Slots Patient Type [P]

COP Maximum Actual Number Bed Slots Patient Type [P] is a component attribute of the Unit object class and measures the maximum number of available bed slots at any given time which may be allocated to patients of type P within the unit.

COP Mean Actual Occupancy Rate Period = [P]

COP Mean Actual Occupancy Rate Period = [P] is a component attribute of the Unit object class and measures the average (mean) actual occupancy rate for the bed slots within the unit during a period of time T.

COP Operational Manager Name

COP Operational Manager Name is a component attribute of the Unit object class and measures the name of the person responsible for the overall operational management of the unit.

COP Projected Labour Component [N] per Bed Slot Time = [T]

COP Projected Labour Component [N] per Bed Slot Time = [T] is a component attribute of the Unit object class and measures the projected amount of labour component N per bed slot within the unit at time T, where 'N' could refer to nursing labour, clinician labour, and so on.

COP Target Labour Component [N] per Bed Slot

COP Target Labour Component [N] per Bed Slot is a component attribute of the Unit object class and measures the target amount of labour component N per bed slot within the unit, where 'N' could refer to nursing labour, clinician labour, and so on.

COP Target TISS Component [N] per Bed Slot

COP Target TISS Component [N] per Bed Slot is a component attribute of the Unit object class and measures the target amount of TISS component N per bed slot within the unit, where N could be any TISS component which denotes the consumption of a particular healthcare resource or group of resources.

COP TISS COMPONENT [N] PER BED SLOT DISCREPANCY TIME = [T]

COP TISS Component [N] per Bed Slot Discrepancy Time = [T] is a component attribute of the Unit object class and measures the difference between the projected and actual average consumption of the TISS component N per bed slot within the unit per unit of time starting at time T.

12.2. Proposed Operational Model Data Attributes

POP BED SLOT ID

POP Bed Slot ID is a component attribute of the Bed Slot object class and measures a unique identifier to identify individual bed-slots

POP ACTUAL BED SLOT STATUS TIME = [T]

POP Actual Bed Slot Status Time = [T] is a component attribute of the Bed Slot object class and measures the actual scheduling status of the bed slot at time T. Scheduling status is a categorical measure evaluated as Unallocated-Unoccupied, Allocated-Occupied, Allocated-Unoccupied.

POP ACTUAL LABOUR COMPONENT [N] STATUS TIME = [T]

POP Actual Labour Component [N] Status Time = [T] is a component attribute of the Bed Slot object class and measures whether or not the bed slot comprised a labour component of type N during a standardised period of time beginning at time T, where 'N' could refer to nursing labour, clinician labour, and so on.

POP ACTUAL TISS COMPONENT [N] STATUS TIME = [T]

POP Actual TISS Component [N] Status Time = [T] is a component attribute of the Bed Slot object class and measures whether or not the bed slot comprised the TISS Component N during a standardised period of time beginning at time T, where N could be any TISS component which denotes the consumption of a particular healthcare resource or group of resources.

POP BED SLOT UNIT NAME

POP Bed Slot Unit Name is a component attribute of the Bed Slot object class and measures the unique identifier of the unit of which the bed slot is a component.

POP PATIENT SCHEDULED MONOTONIC BED SLOT TIME = [T]

POP Patient Scheduled Monotonic Bed Slot Time = [T] is a component attribute of the Bed Slot object class and measures the patient which has been monotonically allocated the bed slot at time T.

POP PATIENT SCHEDULED NON-MONOTONIC BED SLOT TIME = [T]

POP Patient Scheduled Non-Monotonic Bed Slot Time = [T] is a component attribute of the Bed Slot object class and measures the patient which has been non-monotonically allocated the bedslot at time T.

POP PATIENT TYPE [P] ADMISSIBLE

POP Patient Type [P] Admissible is a component attribute of the Bed Slot object class and measures whether or not the patient type P may consume the healthcare resources represented by the bed-slot

POP PROJECTED BED SLOT STATUS TIME = [T]

POP Projected Bed Slot Status Time = [T] is a component attribute of the Bed Slot object class and measures The projected scheduling status of the bed slot at time T. Projected status is a categorical measure evaluated as Unallocated-Unoccupied, Allocated-Occupied, Allocated-Unoccupied.

POP PROJECTED LABOUR COMPONENT [N] STATUS TIME = [T]

POP Projected Labour Component [N] Status Time = [T] is a component attribute of the Bed Slot object class and measures whether or not the bed slot is projected to comprise a labour comopnent of type N during a standardised period of time beginning at time T, where 'N' could refer to nursing labour, clinician labour, and so on.

POP PROJECTED TISS COMPONENT [N] STATUS TIME = [T]

POP Projected TISS Component [N] Status Time = [T] is a component attribute of the Bed Slot object class and measures whether or not the bed slot is projected to comprise the TISS Component N during a standardised period of time beginning at time T, where N could be any TISS component which denotes the consumption of a particular healthcare resource or group of resource

POP PATIENT HOSPITAL NUMBER

POP Patient Hospital Number is a component attribute of the Patient object class and measures A unique identifier for patients admitted to RBH

POP ACTUAL ADMISSION TIME UNIT [U]

POP Actual Admission Time Unit [U] is a component attribute of the Patient object class and measures the patient's actual time of admission to unit U.

POP ACTUAL LENGTH OF STAY UNIT [U]

POP Actual Length of Stay Unit [U] is a component attribute of the Patient object class and measures the patient's actual length of stay in unit U

POP ADMITTING CONSULTANT

POP Admitting Consultant is a component attribute of the Patient object class and measures the consultant clinician responsible for the patient and their treatment. In those cases where a consultant surgeon is responsible for the patient, this attribute may be set to 0 as applicable.

POP ADMITTING SURGEON

POP Admitting Surgeon is a component attribute of the Patient object class and measures the surgeon who is responsible for the patient and on whose waiting list the patient is entered for surgery. This attribute only applies to surgical patients.

POP DISCREPANCY ADMISSION TIME UNIT [U] TIME = [T]

POP Discrepancy Admission Time Unit [U] Time = [T] is a component attribute of the Patient object class and measures the difference between the patient's actual time of admittion to unit U and the projected time of admission measured at time T.

POP DISCREPANCY LENGTH OF STAY UNIT [U] TIME = [T]

POP Discrepancy Length of Stay Unit [U] Time = [T] is a component attribute of the Patient object class and measures the difference between the patient's actual length of stay in unit U and the projected length of stay measured at time T.

POP PATIENT ADMISSION DIAGNOSIS

POP Patient Admission Diagnosis is a component attribute of the Patient object class and measures the diagnosis which was made when the patient was first admitted to the hospital.

POP PATIENT CLINICAL ATTRIBUTE [N]

POP Patient Clinical Attribute [N] is a component attribute of the Patient object class and measures the set of clinical attributes which represent and evaluate the patient's physiology and pathological condition when the measurement was made. These attributes are assumed to be static and require no updating according to the progression or monitoring of

POP PATIENT CLINICAL ATTRIBUTE [N] TIME = T

POP Patient Clinical Attribute [N] Time = T is a component attribute of the Patient object class and measures the set of clinical attributed which represent and evaluate the patient's physiology and pathological condition at time T. These attributes may be simple attributes such as Blood Pressure at Time T, or they may be more complex attributes such as Parsonnet

POP PATIENT CURRENT DIAGNOSIS

POP Patient Current Diagnosis is a component attribute of the Patient object class and measures the current diagnosis of the patient. It is assumed that this diagnosis may be different from the diagnosis which was made at admission.

POP PATIENT CURRENT LOCATION

POP Patient Current Location is a component attribute of the Patient object class and measures the current unit which is treating the patient

POP PATIENT DATE OF BIRTH

POP Patient Date of Birth is a component attribute of the Patient object class and measures the patient's date of birth.

POP PATIENT DEMOGRAPHIC ATTRIBUTE [N]

POP Patient Demographic Attribute [N] is a component attribute of the Patient object class and measures the values for a set of demographic attributes of the patient. These attributes could measure, for example, Next of Kin, Gender, Age, etc. All of these attributes are not represented individually for reasons of simpllicity.

POP PATIENT HOME ADDRESS

POP Patient Home Address is a component attribute of the Patient object class and measures the home address of the patient.

POP PATIENT HOSPITAL ADMISSION DATE

POP Patient Hospital Admission Date is a component attribute of the Patient object class and measures the date at which the patient was admitted to the hospital

POP PATIENT NAME

POP Patient Name is a component attribute of the Patient object class and measures the name of the patient

POP PATIENT PROJECTED DISCHARGE TIME

POP Patient Projected Discharge Time is a component attribute of the Patient object class and measures the most current projection of the time at which the patient is expected to be able to be discharged from the unit in which they are currently being treated.

POP PATIENT SCHEDULING STATUS

POP Patient Scheduling Status is a component attribute of the Patient object class and measures the patient's current scheduling status

POP PROJECTED ADMISSION TIME UNIT [U] TIME = [T]

POP Projected Admission Time Unit [U] Time = [T] is a component attribute of the Patient object class and measures the projected time when the patient is expected to be admitted to unit U at time T. In those cases where the patient is not expected to be admitted to unit U, the attribute may be assumed to be evaluated as 0. In recording the time when this projection is

POP PROJECTED LENGTH OF STAY UNIT [U] TIME = [T]

POP Projected Length of Stay Unit [U] Time = [T] is a component attribute of the Patient object class and measures the patient's projected length of stay in unit U at time T.

POP UNIT NAME

POP Unit Name is a component attribute of the Unit object class and measures a a unique identifier for each unit.

POP ACCEPT PATIENT TYPE [P] FROM UNIT [U]

POP Accept Patient Type [P] From Unit [U] is a component attribute of the Unit object class and measures whether or not the unit may admit patients of type P from unit U.

POP ACTUAL LABOUR COMPONENT [N] PER BED SLOT TIME = [T]

POP Actual Labour Component [N] per Bed Slot Time = [T] is a component attribute of the Unit object class and measures the actual amount of the labour component N per bed slot within the unit at time T, where 'N' could refer to nursing labour, clinician labour, and so on.

POP ACTUAL NUMBER ALLOCATED-OCCUPIED BED SLOTS TIME = [T]

POP Actual Number Allocated-Occupied Bed Slots Time = [T] is a component attribute of the Unit object class and measures the total actual number of bed slots within the unit which are occupied by a patient at time T

POP ACTUAL NUMBER ALLOCATED-UNOCCUPIED BED SLOTS TIME = [T]

POP Actual Number Allocated-Unoccupied Bed Slots Time = [T] is a component attribute of the Unit object class and measures the total actual number of bed slots within the unit which are allocated to a patient, but not occupied by any patient at time T

POP ACTUAL NUMBER UNALLOCATED-UNOCCUPIED BED SLOTS TIME = [T]

POP Actual Number Unallocated-Unoccupied Bed Slots Time = [T] is a component attribute of the Unit object class and measures the total actual number of bed slots within the unit which are neither allocated nor occupied by any patient at time T

POP ACTUAL OCCUPANCY RATE TIME = [T]

POP Actual Occupancy Rate Time = [T] is a component attribute of the Unit object class and measures the proportion of bed slots within the unit at time T which are occupied.

POP ACTUAL TISS COMPONENT [N] PER BED SLOT TIME = [T]

POP Actual TISS Component [N] per Bed Slot Time = [T] is a component attribute of the Unit object class and measures the actual amount of TISS component N per bed slot within the unit which is consumed per unit time at a time T, where N could be any TISS component which denotes the consumption of a particular healthcare resource or group of resources.

POP ALLOCATED-OCCUPIED BED SLOTS DISCREPANCY TIME = [T]

POP Allocated-Occupied Bed Slots Discrepancy Time = [T] is a component attribute of the Unit object class and measures the difference between the projected total number of allocated and occupied bed slots at time T and the actual number of allocated and occupied bed slots at time T in the unit.

POP ALLOCATED-UNOCCUPIED BED SLOTS DISCREPANCY TIME = [T]

POP Allocated-Unoccupied Bed Slots Discrepancy Time = [T] is a component attribute of the Unit object class and measures the difference between the projected number of allocated and unoccupied bed slots within the unit at time T and the actual number of allocated and unoccupied bed slits within the unit at time T.

POP CLINICAL DIRECTOR NAME

POP Clinical Director Name is a component attribute of the Unit object class and measures the name of the clinical director of the unit with overall responsibility for all aspects of the delivery of health care to patients within the unit.

POP DISCHARGE PATIENT TYPE [P] TO UNIT [U]

POP Discharge Patient Type [P] To Unit [U] is a component attribute of the Unit object class and measures whether or not patient type P may be discharged from the unit to unit U.

POP LABOUR COMPONENT [N] PER BED SLOT DISCREPANCY TIME = [T]

POP Labour Component [N] per Bed Slot Discrepancy Time = [T] is a component attribute of the Unit object class and measures the difference between the projected and actual average amount of the labour component N that is consumed per bed slot per unit of time at time T within the unit.

POP MAXIMUM ACTUAL NUMBER BED SLOTS

POP Maximum Actual Number Bed Slots is a component attribute of the Unit object class and measures the maximum number of bed slots available within the unit at any given time that may be allocated to any of the admissible patient types

POP MAXIMUM ACTUAL NUMBER BED SLOTS PATIENT TYPE [P]

POP Maximum Actual Number Bed Slots Patient Type [P] is a component attribute of the Unit object class and measures the maximum number of available bed slots at any given time which may be allocated to patients of type P within the unit.

POP MEAN ACTUAL OCCUPANCY RATE PERIOD = [P]

POP Mean Actual Occupancy Rate Period = [P] is a component attribute of the Unit object class and measures the average (mean) actual occupancy rate for the bed slots within the unit during a period of time T.

POP MEAN OCCUPANCY RATE DISCREPANCY PERIOD = [P]

POP Mean Occupancy Rate Discrepancy Period = [P] is a component attribute of the Unit object class and measures the difference between the average (mean) projected occupancy rate and the average (mean) actual occupancy rate for bed slots within the unit during a period of time P

POP MEAN PROJECTED OCCUPANCY RATE PERIOD = [P]

POP Mean Projected Occupancy Rate Period = [P] is a component attribute of the Unit object class and measures the average (mean) projected occupancy rate for the bed slots within the unit during a period of time T.

POP OCCUPANCY RATE DISCREPANCY TIME = [T]

POP Occupancy Rate Discrepancy Time = [T] is a component attribute of the Unit object class and measures the difference between the projected and actual proportion of bed slots in the unit which are occupied at time T.

POP Operational Manager Name

POP Operational Manager Name is a component attribute of the Unit object class and measures the name of the person responsible for the overall operational management of the unit.

POP Projected Labour Component [N] per Bed Slot Time = [T]

POP Projected Labour Component [N] per Bed Slot Time = [T] is a component attribute of the Unit object class and measures the projected amount of labour component N per bed slot within the unit at time T, where 'N' could refer to nursing labour, clinician labour, and so on.

POP Projected Number Allocated-Occupied Bed Slots Time = [T]

POP Projected Number Allocated-Occupied Bed Slots Time = [T] is a component attribute of the Unit object class and measures the total projected number of bed slots within the unit which are occupied by a patient at time T.

POP Projected Number Allocated-Unoccupied Bed Slots Time = [T]

POP Projected Number Allocated-Unoccupied Bed Slots Time = [T] is a component attribute of the Unit object class and measures the total projected number of bed slots within the unit which are allocated to a patient, but not occupied by any patient at time T

POP Projected Number Unallocated-Unoccupied Bed Slots Time = [T]

POP Projected Number Unallocated-Unoccupied Bed Slots Time = [T] is a component attribute of the Unit object class and measures the projected total number of unallocated and-unoccupied bed slots within the unit at time T.

POP Projected Occupancy Rate Time = [T]

POP Projected Occupancy Rate Time = [T] is a component attribute of the Unit object class and measures the projected proportion of allocated and occupied bed slots within the unit at time T.

POP Projected TISS Component [N] per Bed Slot Time = [T]

POP Projected TISS Component [N] per Bed Slot Time = [T] is a component attribute of the Unit object class and measures the total projected amount of the TISS component N consumed per bed slot within the unit per unit of time at time T, where N could be any TISS component which denotes the consumption of a particular healthcare resource or group of resources.

POP TARGET LABOUR COMPONENT [N] PER BED SLOT

POP Target Labour Component [N] per Bed Slot is a component attribute of the Unit object class and measures the target amount of labour component N per bed slot within the unit, where 'N' could refer to nursing labour, clinician labour, and so on.

POP TARGET TISS COMPONENT [N] PER BED SLOT

POP Target TISS Component [N] per Bed Slot is a component attribute of the Unit object class and measures the target amount of TISS component N per bed slot within the unit, where N could be any TISS component which denotes the consumption of a particular healthcare resource or group of resources.

POP TISS COMPONENT [N] PER BED SLOT DISCREPANCY TIME = [T]

POP TISS Component [N] per Bed Slot Discrepancy Time = [T] is a component attribute of the Unit object class and measures the difference between the projected and actual average consumption of the TISS component N per bed slot within the unit per unit of time starting at time T.

POP UNALLOCATED-UNOCCUPIED BED SLOTS DISCREPANCY TIME = [T]

POP Unallocated-Unoccupied Bed Slots Discrepancy Time = [T] is a component attribute of the Unit object class and measures the difference between the actual and projected number of unallocated and unoccupied bed slots within the unit at time T.

13. Appendix: Related Publications

1. Hughes M, Carson E R, Morgan C, Page A, Summers R. Requirements for a computerised operating room scheduling assistant, in Proceedings of the 8th International IMEKO Conference on Measurement in Clinical Medicine, edited by R Magjarevic. Zagreb: Korema 1998; 12/3-12/6.

2. Hughes M, Carson, E R, Morgan C, Makhlouf M, Summers R. A Petri net based model of patient flows in a progressive patient care system, in Proceedings of the 20th Annual International Conference of the IEEE Engineering in Medicine and Biology Society. Piscataway, NJ: IEEE, 1998.

3. Hughes M, Carson E R, Morgan C, Makhlouf M. Modelling a progressive-care system using a coloured timed Petri net. Trans Inst Meas Contr 2000.

4. Hughes M, Carson E R, Makhlouf M, Morgan C. Extending the computersation of healthcare delivery, in Proceedings of the 4th IFAC Symposium on Modelling and Control in Biomedical Systems 2000.

5. Hughes M, Carson E R, Morgan C, Makhlouf M. Development and evaluation of computerised control systems in healthcare delivery. Control Engineering Practice 2001.

6. Hughes M, Carson E, Morgan C, Silvester P. Evaluating Admissions Control in a Surgical Progressive-Care System. IFAC 2002.