Investigating Efficiency in the Emergency Department at Groote Schuur Hospital

Author:
Allister MOWBRAY

Supervisor:
Professor TJ STEWART

Thesis presented for the degree of Master of Science
in the Department of Statistics

University of Cape Town
May 5, 2010

Contents

Chapter 1

Introduction

1.1 The context of the research

Since democracy in 1994, the South African government has promised improved healthcare access to the general public at a minimum of cost. However, this has proven to be rather difficult, due to a number of factors including scant resources (human and otherwise), oversubscription of services, and an ineffecient, poorly understood referral system.

In its current context, the health system consists of three levels: Level 1 or primary care hospitals and clinics, Level 2 hospitals, and Level 3 hospitals. Level 1 hospitals and clinics are day hospitals which are unequipped to handle any sort of surgery or serious illnesses. Level 2 hospitals are better equipped and are able to handle some surgery, and can accommodate patients overnight. Level 3 hospitals are fully equipped to handle any sort of illness or surgery. In the current system, patients are supposed to first go to a Level 1 hospital, which is then instructed to send them to a Level 2 hospital if they cannot be accommodated. The Level 2 hospitals are then instructed to send the patients to a Level 3 hospital only if they cannot deal with them. In theory, the patient allocation should be as in Figure 1.1 (data provided by Chief Logistics Officer, Groote Schuur Hospital).

This research takes place in Groote Schuur Hospital; a state-funded Level 3 teaching hospital located in Cape Town, South Africa. The reality for a hospital such as Groote Schuur is that they are currently seeing 20-25% of all patients seeking public medical care in Cape Town instead of the less than 5% that they should be seeing. This is due to a number of reasons which will not be mentioned here. As can be imagined, this places a considerable strain on the existing resources of the hospital. As such, these hospitals are constantly seeking ways in which they can increase their efficiency whilst maintaining their current resource use.

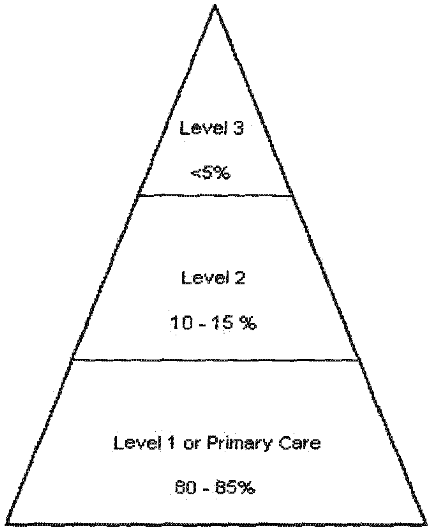

Figure 1.1: Theoretical patient allocation in the South African health system

1.2 The research projects

One of the big problems at an institution such as Groote Schuur is the lack of resources. Due to this fact, most of the existing resources are directed at 'keeping the ship afloat'. In other words, there is not a lot of effort given to updating existing methods which have worked to some degree in the past. However, precisely because of this fact, there are a number of areas which could be drastically improved if there were someone who had the time and tools to analyse them. With this in mind, the Chief Logistics Officer(CLO) at Groote Schuur approached the Operational Research group in the Depratment of Statistical Sciences at the University of Cape Town to see if there were some projects on which they could work together. After some consideration, three possible projects emerged:

- Outpatient department mapping
- Staff scheduling

- Quantifying kangaroo mother care

1.2.1 Outpatient department mapping

Over a year, some 500 000 patients attend more than 100 clinics in the outpatient unit. The problem which exists here is that some weeks are very busy, while others are very quiet, in a manner which is difficult to predict. Obviously this variation has profound effects on staffing and downstream services such as the pharmacy. It is felt that a probable cause of this problem is the current appointment system. In order to analyse this problem, we would look at building a simulation of the outpatient department, and then use this to evaluate different appointment scheduling systems, and possibly different areas of the clinics which could be improved.

1.2.2 Staff scheduling

Groote Schuur has a staffing complement of close on 2000 nurses and about 800 doctors. Due to the 24 hour service which the hospital provides, shift allocation for staff becomes a major task for all senior staff to coordinate. These shift rosters are drawn up manually, incorporating various constraints on the shifts along with requests from the staff. This obviously takes up a lot of time each month. The approach to this problem would be to formulate a linear programming model which incorporates all of the constraints and requests, and solves the problem electronically.

1.2.3 Kangaroo mother care

Kangaroo mother care is a specific method used within hospitals to care for newborn babies involving contact and support between the mother and her baby. It is understood that there is a complex relationship between the mother's willingness/ability to deliver good kangaroo mother care to her newborn infant and a host of other factors such as socio-economic standing, gravidity, age of mother, and so on. The difficulty is that the relationship between these factors and the outcome (good kangaroo mother care) is ill-defined. The problem here would be to develop a model which combines quantitative and qualitative factors, and enables the hospital to predict if the mother will be able to deliver good kangaroo mother care to her newborn infant.

1.3 Deciding on a project

After due consideration of the various projects in terms of available data, and requisite skills required, it was decided to tackle the problem of outpatient department (OPD) mapping. In order to build a simulation model of the OPD at Groote Schuur, it was necessary to spend some time there observing the various processes, and speaking to some of the staff members. In the course of this, it also became apparent that it would be beneficial to begin by using Soft Systems Methodologies (SSM) in order to structure the problem. This was mainly because the problem was not very clear from the beginning.

It soon became clear that this project was a major undertaking, and would require a lot more work than was initially thought. As mentioned before, there are over 100 clinics in the OPD, and analysing each of them would be impossible. Therefore, it was decided to downscale the project to only 1 or 2 clinics which would be analysed more closely. We then began to build a simulation model of the OPD floor itself, bearing in mind that models of the relevant clinics would be added to this initial model at a later stage. An SSM analysis was also underway at this time.

At this point, the CLO examined our progress, and thought that this simulation method would be relevant to another problem which was happening in the hospital. The Emergency Department(ED) was experiencing problems with excessive patient waiting times with no readily identifiable cause. The CLO thought that this was a more pressing problem, and one which could be more easily solved than the problem in the OPD. It must also be mentioned that we were experiencing problems with accessing certain data for the OPD problem at this stage, and that relevant data were more easily accessible for the ED problem.

Therefore, it was decided to concentrate on the ED problem, and wait for the data to become available in the OPD. As it turned out, that data never became available, and the problem in the ED proved to be of sufficient scale to warrant a full problem solving approach. So, the research in the OPD was dropped in favour of a full problem solving approach to investigating the efficiency in the ED.

1.3.1 The emergency department problem

It was decided to adopt a similar approach to tackling the problem in the ED as we had begun to use for the OPD. Therefore, we will begin by structuring the problem using SSM. It will not be necessary to conduct a full

SSM intervention; we will rather use parts of SSM in order to make our approach to the problem clearer. This will involve identifying and meeting with stakeholders within the ED itself in order to find out their views on the issues.

Once the problem has been structured, it will be necessary to begin gathering data for the purposes of building a model of the ED. There should not be the same problems with accessing data as was experienced in the OPD. However, it is likely that at least some of this data will need to be based on expert estimation. Once the data have been collected, we will be able to build a discrete event simulation model of the ED in Simul8. This model will then need to be validated through various sensitivity tests. After these have been performed, we can begin to conduct experiments on the model. These experiments should help us to determine the problem areas which are impeding the efficient working of the system.

Chapter 2

Literature Review

2.1 Introduction

Hospital outpatient and emergency departments have long been a logical area for the application of operational research(OR). Indeed, the operation of outpatient departments (OPDs) was one of the first areas in which OR was applied. As early as 1952, Bailey [13] found that doctors' consultation times could be adequately described by a Pearson distribution. Therefore, by studying doctors' idle time and patients' waiting times, an appointment system minimising waiting time could be derived.

Within a hospital, OPDs and emergency departments (EDs) are the areas which are most commonly used by members of the general public. As such, they are areas which are often under much scrutiny, and the efficiency of these departments reflect on the hospital as a whole. Perhaps this is the reason that the improvement of the operation of these departments has been a most commonly tackled subject by researchers in OR healthcare in the past.

Much of the literature in the past has concentrated on 3 areas of improvement across both emergency and outpatient departments [6]. These 3 areas are:

- Patient scheduling and admissions
- Patient routing and flow schemes
- Scheduling and availability of resources

2.2 Patient scheduling and admissions

2.2.1 Introduction

Scheduling and admissions play a vital role in controlling patient waiting times, and are relatively simple to alter. Thus, wherever possible, this is an area where many researchers have concentrated their operational improvements. Patient scheduling and admissions focus on procedures that deal with how patient appointments are scheduled in terms of when they are set, and their length of time. This generally involves rules that determine when appointments can be made, and the length of time between appointments. Obviously this can have a significant impact on how resources are utilised so as to maximise patient flow. These rules range from single-block appointments on one extreme to individual appointments on the other. Most appointment systems concentrate on modifying and combining these rules in some way. The majority of studies which have sought to alter patient scheduling rules have concentrated on OPDs [6]. This is because OPDs are the only areas in hospitals which see appointment patients on a regular basis. Altering patient scheduling rules cannot help when modeling EDs as all of the patients are necessarily walk-ins (patients with no prior appointments who literally 'walk in' off the street).

Whilst reviewing past studies in this area, it becomes clear that a typical intervention in patient scheduling involves the researchers working in close conjunction with the staff in the clinic. Through this, the analysts are able to build up an accurate simulation model for each clinic, and then verify it through showing it to the staff, and comparing it with actual data. Then, different appointment rules such as those described above can be tested within the system, and the results for each can be compared. Many interventions are staged as an ongoing process, and the analysts often try to create user-friendly front and back ends to the simulation which they can leave with the clients to use themselves. Staff are then able to use this facility to compare the different schedules, and see for themselves how the system will react.

2.2.2 Case studies

A project falling under this category is described in Harper *et al* [14]. The project in question takes place in a major Ear, Nose and Throat clinic in the United Kingdom (UK). Under the Patient's Charter, every hospital in the UK is required to see patients within 30 minutes of their appointment time. An OPD is a complex nonlinear system which is very difficult to manage, and so these departments often fail to adhere to this limit on waiting times. This particular department operates 10 clinic sessions per week and deals with 5 main types of patients. The schedules for each of these clinics have evolved over many years and were felt to be inefficient by many of the staff, leading to high patient waiting times and bottlenecks within the clinic.

It was decided to conduct a study within this department using an OR methodology. The researchers began by analysing the current outpatient system and developing accurate models of the department through close consultation with the staff. These took the form of flowcharts detailing the patient flows, resources and queuing points within the department. Since there were very little data concerning the lengths of consultations with the clinic staff, it was decided to spend a week collecting data at the clinics. The week was spent collecting data on waiting times, service times, patient flows and arrival patterns. All of this information was then built into a discrete event simulation (DES) using the Simul8 package. Then, 3 performance measures were chosen to evaluate the impact of any changes. These were:

- The average time a patient had to spend waiting from their appointment time until their first service.

- The percentage of patients who had to spend more than 30 minutes waiting for their first service (measured from their appointment time).

- The average time that a patient spends in the clinic (measured from the time they enter the clinic to the time that they leave again).

In this study, different simulation models were built for each clinic, resulting in a total of 14 simulations for the project. So, for the purposes of evaluating the effects of any changes, one clinic and hence one simulation was chosen as a control. Nine different scenarios were developed with each of them incorporating changes in: the number of patients scheduled for each clinic; the clinic starting times; and the appointment scheduling system. It was felt that these were areas where changes would have the greatest effect on patient waiting times. The 9 scenarios were:

1. Patients arrived every 5 minutes. Patients of the same type arrived sequentially.

2. Patients were scattered randomly amongst the appointment slots.

3. Patients were booked randomly, but there was a buffer period of 15 minutes where no patients were booked.

4. Patients were booked in blocks with the length of each block varying according to each patient type.
5. There were no diary patients allowed(these are last minute patients manually entered into an appointment book).
6. The schedule was not altered in this test, instead all clinics started on time.
7. Patients were booked based on an algorithm which spread the appointments out over the whole clinic session.
8. Patients were booked in large blocks scheduled at the beginning of the session.
9. A combination of scenario 7 and scenario 6.

The results for 40 runs of the control model indicated that the factor having the most influence on all 3 performance measures was the time at which the clinic started (scenarios 6 and 9). The scenarios which changed the appointment structures and not the number of patients (scenarios 7 and 9) had the most promising results. With scenario 7 it was possible to reduce patient waiting times by 10 minutes, and with scenario 9, the time to first service for patients could be reduced by 50%. Using statistical t-tests, it was also possible to see that the results for scenarios 4 to 9 were all significantly different from the unaltered control model.

Lehaney *et al* [3] described an earlier intervention which also took place within an OPD at a hospital in the UK. This hospital operated under the same Patient's Charter as mentioned earlier, meaning that patients should not wait in queues for more than 30 minutes. However, this was not happening at this hospital, and the aim of this study was to address the gaps between the expectations of patients and service providers. Lehaney *et al* adopted a different approach to that seen in the paper by Harper *et al.* They argued that most of the research previously done within this area has involved some sort of ad hoc problem structuring, but nothing on a formal basis. Therefore, they proposed using some methods from Soft Systems Methodology (SSM), a recognised 'soft' OR technique, in order to formally structure the problem, and to guide them through the whole problem-solving process.

The case was undertaken in four linked stages, with the streams of enquiry feeding into one another as the project progressed. The objectives and progress of the project were monitored throughout, and control actions were undertaken where necessary. The first stage involved cultural analysis of the problem situations. Here all of the roles of the participants, and the social and political systems were explored. This was also where the general nature of the project and its timescale were agreed upon, and some soft objectives were also set. Additionally, it was here where it was decided to use Simul8 as the simulation modeling package. From here they moved onto the second stage which involved task modeling. This was done using simple flowcharts, as these were easy for all team members to understand. A number of flowcharts were developed which helped all of the participants to gain a thorough understanding of the system. However, in the end, it was decided to drop all of the various flowcharts in favour of one flowchart which tracked the patients from entering the clinic to leaving.

The third stage involved primary task modeling, and at this stage the static flowcharts were discarded for a dynamic computer simulation. Patient waiting time and doctor utilisation were reaffirmed as the major considerations, and the simulation models were built to depict the immediate situation in the clinic, from patient arrival to patient departure. A prototyping approach was adopted to build the models, and the 10th prototype of the simulation of the department was chosen as the most accurate. This model was checked against actual data, and was found to give exactly the same results.

The project team then took all of the appointments for each doctor for the duration of the clinic and sorted them in ascending order of consultation duration. These sorted times were then used to schedule the patients for a new simulation prototype, known as Model 11, which was a slight variation on Model 10. Having shorter duration consultations occurring earlier in the clinic session allows early patients to arrive nearly simultaneously, but still not have a long waiting time. The results of the new simulation model showed a drop in waiting time of up to 40 minutes for some patients, and doctors spent 64% of the clinic duration in contact with patients. One of the clinic's doctors working as part of the project team mentioned that, in practice, it could become difficult and tiresome to predict how long individual patient consultations could take. Therefore, it was decided that a more realistic scenario may be to provide estimates of consultancy duration averages for particular patient types. The scheduling may then be based against these figures. This would probably result in a smaller reduction in queuing times, but would show improvement over the old system, with little effort on anyone's part at the hospital.

The fourth stage involved the naming and modeling of relevant issue-based systems, and in this case the system of intervention itself was named and modeled. The key stakeholders reviewed the process in which they had been engaged, and at that stage root definitions and conceptual models were used to encapsulate the process. The framework utilised the information gained during the previous stages, to produce a framework of the intervention itself. The root definition for the system being modeled was:

A system owned by a group of key stakeholders, run by analysts, key stakeholders, and other stakeholders, who use simulation modeling as an aid to develop and implement operational policy which meets both internal and external criteria.

The system which is being simulated is the patient flow through one clinic of the OPD, but the root definition is for an intervention system that utilises simulation modeling, but which is not itself being simulated.

This was a unique intervention and, despite its success, Lehaney *et al* observed that opportunities to transfer this approach to other interventions may be limited. This is largely due to the the length of time, and subsequently the amount of money, needed for an approach of this sort.

2.3 Patient routing and flow schemes

2.3.1 Introduction

One of simulation's strengths is its ability to model complex patient flows through a system. Users can then alter the patient flow rules and policies, and see how the system reacts. Although they have no control over the arrival patterns of the patients, medical staff are able to control the route through which patients travel. Therefore, by altering this route, and the flow of patients, it is theoretically possible to minimise patient waiting time and increase staff utilisation rates.

Investigating patient routing and flow schemes is a common starting point for any project trying to improve efficiency in an ED. This is because, unlike in an OPD, there is no way to apply scheduling rules. EDs also lend themselves to the introduction of a 'fast-track'. Non-urgent patients generally have the longest waiting time in emergency rooms for obvious reasons. A fast-track is simply a dedicated route for these patients. There will generally be a limited number of resources available for the fast-track (eg 1 consulting area and 1 nurse), as these patients do not require a high degree of care. By introducing a fast-track, hospitals are able to clear the area for the more urgent patients, and at the same time decrease the waiting time for both urgent and non-urgent patients. By using simulation, hospitals are able to see what would happen to patient flow if a fast-track were to be introduced in their ED without having to affect patient care in any way.

2.3.2 Case studies

Garcia *et al* [10] tested the notion of introducing a fast-track in the ED at Mercy Hospital in Miami, Florida. Recent trends in the finances of the ED had indicated the need for an operations improvement analysis, and it was decided to study the flow of patients through the ED with and without a fast-track in place.

Patients entering the ED at Mercy Hospital were sorted into 4 levels of illness acuity with level 0 being the most serious patients, and level 3 being the least serious. The length of wait was determined by the level of acuity, with level 3 patients waiting the longest. Data were collected for some of the service times, and for the total turnaround time in the ED. However, there were no data available for the doctors' service times, so these were estimated after talking with one of the resident doctors. All of the service times were then approximated by Uniform distributions. The arrival patterns were estimated from data collected from the nursing charts, and these were assigned Exponential distributions. Using these data, the researchers were able to model the ED as it stood, without the fast-track.

$$y = \sum_{i=2}^{k} [E[x_i * n * Pr(l = i)] \tag{2.1}$$

where
y = Average excess time in ED due to no fast-track
x_i = Total time in ER for a patient of acuity level i with no fast-track
n = Average no. of patients in queue
l = Patient acuity level
$k = 3$

Using Equation 2.1, they were able to get 2 new distributions for total time in the ED for the patients of acuity level 2 and 3 (who were the patients who would benefit from the introduction of a fast-track). After statistical tests using this information, it was concluded that a fast-track within the ED would be an appealing alternative. So, it was decided to then simulate the ED with and without a fast-track.

The performance measures used were:

1. the cost of implementation

2. the effect on other patients in the ED
3. the patient's length of stay

These performance measures were carefully chosen in order that the hospital could monitor that they were able to provide the same level of care for the other (more acutely ill) patients as they had previously.

When testing the effects of the fast-track, Mercy Hospital provided the project team with six different scenarios for a fast-track that they found feasible. These scenarios differed from each other concerning the time of day that they were utilised and the number of beds that they employed. (There would always be 1 nurse and no doctors available on the fast-track.) The hospital also stipulated that patients must proceed to the regular ED if the fast-track was full. The scenarios tested were as follows

1. Fast-track operational from 10:00 - 20:00, 1 bed available
2. Fast-track operational from 10:00 - 20:00, 2 beds available
3. Fast-track operational from 11:00 - 21:00, 1 bed available
4. Fast-track operational from 11:00 - 21:00, 2 beds available
5. Fast-track operational from 12:00 - 22:00, 1 bed available
6. Fast-track operational from 12:00 - 22:00, 2 beds available

Hypothesis tests were then conducted comparing the flow time for patients with and without the fast-track, and it was concluded that patients of acuity level 3 and 4 greatly benefitted from the introduction of a fast-track, without negatively impacting on the other patients. This was true for all of the scenarios tested, but it was decided that the 1st scenario was the most preferable. This was because it produced the largest reduction in waiting times relative to the amount of inconvenience caused due to having to relocate resources.

A fast-track is not the only way in which to improve patient routing and flow schemes within an ED. Kirtland *et al* presented a paper at the 1995 Winter Simulation Conference after staging an intervention at Peninsula Regional Medical Center's ED in the US where several different routing and flow scenarios were tested. The department was experiencing a decline in the departmental productivity as indicated in the comparative indicators, and the patients were getting increasingly dissatisfied with their length of stays. For these reasons, a team comprising doctors, nurses, and technicians from the ED, along with external management consultants, was formed in order to improve the department's processes. Simulation was decided upon as the method of choice.

The team collected detailed information about the operation of the ED from patient charts, ED logs, computer information systems, interviews, observations, and data collection where information was not available. The arrival pattern to the ED was taken from 3 months worth of data on an hourly basis. The researchers then organised the patient flow into flowcharts spanning 9 categories of patients, and 5 levels of increasing patient acuity. Once this information was organised, they were able to build up a simulation of the ED using MedModel simulation software, which is a specialist medical simulation package. The model was validated by comparing the simulation results to the results from 400 random patient folders taken from the previous year. These folders had details about the patients' conditions, treatments and length of stays. This historical information was also used as a baseline against which to test any changes in the system.

Once the accurate simulation representation of the ED was in place, 11 different scenarios to improve patient flow were tested. These were:

1. Setting up a fast-track system in the minor care area.
2. Using point of care testing in the ED where possible.
3. Reducing the number of technicians in the ED.
4. Reducing the number of registered nurses in the ED.
5. Taking the patient back to an open treatment room and not letting them wait in the waiting room until the staff are less busy.
6. Initiating admission room search for an inpatient as soon as a doctor determines the need to admit a patient.
7. Using an internal waiting room for patients waiting on the results of laboratory and other tests (when the ED is busy).
8. Setting up triage protocols that direct the triage nurse to order certain tests.
9. Changing around some intermediate care rooms.
10. Reducing the number of registered nurses and having 5% more patient volume.

11. Reducing patient volume by 5%.

The 3 optimal alternatives in order of patient waiting time saved were:

- setting up the fast track, which saved 15.5 minutes
- placing patients in an open treatment room, which saved 14.1 minutes
- using point of care testing, which saved 8.4 minutes

These 3 changes, which can be used in conjunction, save a total of 38 minutes of waiting time for patients, which amounts to 24% of total waiting time. Interestingly, although the fast-track did save the most patient waiting time, placing patients in an open treatment room was a very close second. This was obviously a scenario which was specific to this study, indicating that we should always look at all possible 'local' options, and the best way to do this is by working closely with the staff. Additionally, the study also helped to identify some 'best practice' alternatives which do not show significant reductions in waiting time, but will help to ease the strain when the department is busy.

Samaha *et al* [20] also tested the effects of a fast-track at Cooper Health System in South Jersey, USA in a similar manner to the previous two papers. They also found that this process would expedite non-critical patients through the system and shorten their length of stay in the ED.

Ruohonen *et al* [21] investigated the effect of introducing a new triaging system in the ED at a hospital in Finland. The ED saw around 34 000 patients annually, but this was expected to increase as a number of units within the hospital were being combined. This study developed a simulation model which demonstrated a new triaging method which sought to reduce patient waiting and throughput times.

The model was also developed using MedModel. There were data available from a previous study, and these were supplemented with data collected for 24 hours a day over a 2 week period. These were collected through the use of a special form, created for the purposes of the project, and completed by staff and patients. In this way, very accurate distributions were able to be fitted for every stage of the process. The model was then verified and validated through using its visual and numerical information, and comparing it to actual data.

The area being tested in this research was that of triaging. In the original system, as is common in EDs around the world, triaging was performed by a registered nurse. The nurse could also perform initial tests (such as blood pressure tests) but could not order tests such as x-rays, which could only be done after the patient had seen a doctor. This is thought to increase the throughput time of the patient, and cause a high degree of utilization of the specialists. The proposed change was to introduce a triage-team, consisting of three staff members, who would receive all patients. There would be a receptionist to input patient data, plus a nurse and a doctor who could order all necessary tests, and possibly more accurately define a patient's acuity level. This change was examined on two levels: firstly, on how efficiently the team could operate; and secondly, on how many acutely ill patients the team could process without slowing down the whole operation of the ED. (This was because acutely ill patients would require a longer treatment from the triage-team.) In both cases, the process times would be defined by Uniform distributions.

The triage-team method was tested using several alternative process time scenarios. The staff estimated that the process time would be somewhere between 0.5 - 1 minutes but, as this was just an estimation, several different scenarios were taken under examination. The scenarios were as follows:

1. Process time 0.5 - 1.5 minutes
2. Process time 1 - 2 minutes
3. Process time 2 - 4 minutes
4. Process time 4 - 6 minutes
5. Process time 6 - 8 minutes
6. Process time 8 - 10 minutes
7. Process time 10 - 12 minutes
8. Process time 12 - 14 minutes

All developed scenarios were tested and the results were compared to the existing operation, concentrating on the average throughput time of all patients. The results showed that if the operation is as effective as staff had estimated, there would be a 26% reduction of the average throughput time. The results also indicated that the operation would become more effective if the process time is under 12 - 14 minutes.

Medeiros *et al* [5] also looked at a similar type of idea in their research in the US. However, they proposed introducing a physician only at the triage stage in an ED. This would allow testing to be done at an earlier stage and allow more accurate definitions of a patient's acuity levels. They found that it reduced a patient's length of stay by more than 23%. Holm *et al* [7] also proposed introducing a physician at the triage stage in a study done in Norway. Their results showed that the overall length of stay was not significantly reduced, but they argued that patients seeing a physician at an earlier stage in the process found it more reassuring.

Introducing physician triage is an interesting concept, but it is one that few hospitals in developing countries can afford to do with their lack of doctors.

2.4 Scheduling and availability of resources

2.4.1 Introduction

While much of the research into reducing patient waiting time has concentrated on scheduling the patients themselves, it is also possible to approach this problem from the other side, by looking at scheduling doctors and nurses in a more efficient manner. This approach has been used when looking at EDs because, as mentioned previously, the patient arrival patterns cannot be altered and so different approaches to the problem need to be sought. If the patient arrivals cannot be altered, then the problem of scheduling staff becomes a combinatorial problem. However, this is not a straight-forward linear programming problem as the patient arrivals can never be modeled as deterministic, and so a stochastic constraint needs to be introduced.

2.4.2 Case studies

When using this approach in an OPD, researchers often combine it with applying some patient scheduling rules as well. In a paper presented at the 2006 Winter Simulation Conference Wijewickrama *et al* [2] described an intervention made in an OPD of a Japanese hospital. Japan has an increasingly aging society, and long waits in OPDs followed by short consultations are becoming a major problem.

The OPD in question saw 3 different types of patients: appointment patients, same-day appointment patients, and new patients. As is common in Japan, the majority (86%) were appointment patients. Additionally, there were 10 types of patient categories depending on the consultation required. The primary performance measures considered were the waiting times for each of the 10 categories of patient. In addition, an index-weighted average patient waiting time (W), was calculated using the following formula:

$$W = \frac{(W_A * n_A) + (W_S * n_S) + (W_N * n_N)}{(n_A + n_S + n_N)} \tag{2.2}$$

where
W_A = Average patient waiting time for appointment patients
W_S = Average patient waiting time for same day appointment patients
W_N = Average patient waiting time for new patients
n_A = No. of appointment patients
n_S = No. of same day appointment patients
n_N = No. of new patients

The OPD had an existing Access database which had information about patient arrival times and service times based on the patient type and category. This provided the researchers with the core of their data requirements, and any other information was acquired through interviews with staff members. All of this data was used to create a special purpose data generator in Visual Basic which fed data into a simulation which was created using Arena.

Using the OptQuest optimizer, which is built into the Arena simulation package, Wijewickrama *et al* then searched for the best Doctor Scheduling Mix (DSM) possible. OptQuest uses the *tabu* search and scatter heuristic to find the best value for one or multiple objective functions. The optimization problem structure was formulated as follows:

Minimise

$$Z = W \tag{2.3}$$

subject to

$$x_1 + x_2 + x_3 + x_4 + x_5 + x_6 + x_7 + x_8 + x_9 + x_{10} \leq 31$$

$$x_1, x_2, x_3, x_4, x_5, x_6, x_7, x_8, x_9, x_{10} \leq 1$$

where
x_i = No. of doctors allocated for ith consultation service

Then, 10 different appointment scheduling (AS) rules were examined. The AS rules were as follows:

1. Patients are appointed with an interval of the average consultation time.
2. 2 patients are appointed at a time with an interval of twice the average consultation time.
3. Patients are appointed with an interval of the average consultation time plus a fraction of the standard deviation.
4. Patients are appointed with an interval of the average consultation time less a fraction of the standard deviation.
5. This is the same as the first rule, except the first 5 patients are appointed early, and every patient after this is appointed late.
6. This is the same as the fifth rule, except patients are appointed earlier and later respectively.
7. 2 patients are appointed in a block at the beginning, and then individually in average consultation times.
8. 3 patients are appointed in a block at the beginning, and then individually in average consultation times.
9. 4 patients are appointed in a block at the beginning, and then individually in average consultation times.
10. The first patient is appointed. Then after that, patients are appointed with an interval of the average consultation time plus an alternative fraction of the standard deviation.

These were all found to reduce waiting time. However, rule 10 was the best performer, reducing waiting time by 59.65% against the base case. All of these rules were then combined with some of the best DSMs identified by the optimizer. The final result reduced waiting time by 59.95% or 31 minutes per patient.

Centeno *et al* [9] adopted a slightly different approach when looking to reduce staffing costs in an emergency department at a US hospital. They focused on first building a simulation model which established an acceptable length of stay (LOS) within the unit and the number of staff required to operate it. They then looked to apply a linear programming approach which minimised the cost of the staff subject to a number of constraints.

Building the simulation model was a relatively straight-forward affair as the staff had logged in all the patient data over a period of time for the researchers. This enabled them to fit accurate distributions for the arrival patterns for all types of patients and all of the treatment service times. Once the simulation model was verified and validated, the conditions for the experiments were established. They needed to calculate the sample size required to achieve a reliability level of +/- 3.61 when building a 95% confidence level. They were able to establish that they would require 38 replications.

Then, an optimization integer linear programming model was used to find the optimal number of staff (registered nurses) needed to work each shift. The objective function seeks to minimise the labour cost for the registered nurses. The model was:

Minimise

$$Z = \sum_{i=1}^{k} c_i X_i \qquad (2.4)$$

subject to

$$\begin{array}{ll} \sum_{i=1}^{k} X_i \geq a_j & \forall j = 1, 2, ..., n \\ X_i \geq 1 & \forall i = 1, 2, ..., k \end{array}$$

where
X_i = No. of nurses working shift i
c_i = Salary cost for a nurse during shift i
a_j = No. of registered nurses required per period
i = Index for shifts
j = Index for periods
k = Maximum no. of shifts
n = Maximum no. of periods

In Equation 2.4, a_j is calculated according to the simulation model. In other words, it is the number of nurses required in the department in order for the patients to experience a reasonable LOS as calculated by the simulation model. This also enables management to determine the exact number of staff necessary to achieve a specific LOS goal, along with the cost of that staff. Another added bonus is that no-one needs to manually generate the nursing schedule which should save more valuable person hours.

In a similar fashion to the previous two cases, Ahmed *et al* [8] presented a model which integrated simulation with optimization to design a decision support tool for the operation of an ED unit at a governmental hospital in Kuwait. Their primary objective was to evaluate the impact of various staffing levels on service efficiency.

Gunal *et al* [11] adopted an approach to staffing which did not seek to alter the scheduling rules, but rather tested the effects of multi-tasking by doctors. This approach is not purely a scheduling problem as the previous three papers were, but also draws in some elements from patient routing and flow schemes. The research took place in an ED in the UK and focused on the multi-tasking behaviour and experience level of medical staff.

The unit in which they worked was mid-sized, seeing approximately 45 000 patients annually. The idea was to develop a DES of the department's activities with the intention that this served as a generic model of an ED that could be parameterised to fit a range of such departments in different hospitals. The data which they used to parameterise and validate their model were obtained from the electronic patient admission data which the hospital kept. These were supplemented with data from paper-based patient cards which were completed by medical staff and recorded every detail of patient treatment. Using all of this information, they were able to fit distributions for the arrival patterns and service times, and use these to develop a DES of the ED by using Micro Saint Sharp.

Doctors and nurses are scant resources in EDs and, most of the time, they treat multiple patients concurrently. In this study, the idea was to fragment a doctor or nurse into mini doctors, and thus examine the effect of multi-tasking on patient waiting time. Initially, senior doctors were split into 6, junior doctors were split into 4, and nurses were split into 2. These numbers were known as 'Multi-tasking Factors (MTF)'.

The model was validated using actual data from the previous year, and then experimentation could commence. The first experiments merely altered the MTF values. The base model was 2-4-6 (as described previously), and this was altered to 1-1-1, 1-4-6, and 2-7-7. Interestingly, the 2-7-7 result almost converged with the 2-4-6 result, suggesting that more multi-tasking by doctors may not increase the performance. However, decreasing the nurse MTF by 1 led to a significantly worse performance. The worst performance of all was when all of the MTFs were set to 1, suggesting that the multi-tasking of staff, in general, is a real determinant of performance.

The next area of experimentation focused on the effect of other factors on the performance of the system as a whole. The parameters which were altered in the model were: the treatment times, x-ray service times, and physical cubicle capacities. The scenarios tested were as follows:

1. It was assumed that experienced doctors could treat patients quicker. In this scenario, all doctors worked at the same pace, but base case multi-tasking was retained.

2. The process times for x-rays and other tests were increased by 10 minutes.

3. The proportion of patients requiring x-rays was decreased.

4. More cubicles were introduced into the ED area.

The results showed that scenarios 1, 3, and 4 demonstrated increased performance, whereas scenario 2 showed a decreased performance. This is not really surprising considering the changes in inputs. However, the interesting part of this research was the effect of building multi-tasking into the model. This is a unique approach which makes complete sense in an ED environment.

2.5 Other Cases

2.5.1 Introduction

The last two cases which will be included do not fall under any of the previous three headings. The first case can be said to examine multi-service health facilities and health planning, while the second compares two different hospital policies within an ED, and could technically be included in Section 2.3, but seems more appropriate here . These studies were both performed as part of masters dissertations at UK universities, and are therefore both a little less sophisticated than those with which we have previously dealt. However, parts of both of them, particularly with regard to data capture, should still be relevant to the research being done here.

2.5.2 Case studies

Ashton *et al* [15] conducted an OR study of the operation of an NHS walk-in centre in Liverpool in 2001. Walk-in centres are nurse-led and provide community-based ambulatory care and information about how to access other services. This particular centre shared a building with an NHS Direct (emergency advice centre), a GP practice, and various other primary and community health care services. All of these services were being relocated, and the aim of the study was to assist in planning the scheduled move by examining how the different services located in the centre could be accommodated and operated to the best overall effect. Particular attention would be given to modeling expected patient flows through the new centre, and to assessing how different levels of demand would affect patient waiting times and the number of patients in the waiting room.

A basic problem structuring approach involving a stakeholder analysis and system description was undertaken initially. This initial problem structuring helped the researchers to view the health centre as consisting of three distinct elements: the patients, the walk-in centre, and the clinics. They then used flow diagrams in order to help them to understand the activities associated with these three elements. Using these flow diagrams, they were able to develop an initial simulation in Micro Saint. This simulation consisted of three main parts: one for generating patient arrivals, one for the clinics, and one for the walk-in centre.

The data used for the simulation were gathered from a variety of sources. The arrival pattern for walk-in centre patients was based on NHS system reports. The arrangements for triage and details about the number of seriously ill patients arriving were obtained through discussion with a senior nurse. Times for treatment, triage, and reception were taken from service reports compiled for the Department of Health. However, these were slightly modified after discussion with the nurses involved. The clinic arrivals were based on information gathered through discussion with clinic staff. These staff also provided supplementary information regarding the time spent by patients at reception.

The model was verified through comparison with historical data. Additionally, patients were 'tracked' through the system to ensure that they were following the correct routes. Extreme condition tests were also undertaken and simplifying conditions input to aid the verification process. Validation checks were carried out by checking the logic of flow diagrams with the clients, and confirming number of patients and length of queues with staff members. Various sensitivity tests were also carried out on the model.

This model was not intended to be used to implement any major changes. Rather it was intended to be used as a tool to see how the centre was coping, and to examine the effects of any minor changes. These changes included the introduction of a shared reception, examining different triage procedures, changing the timing of certain clinics, and altering some of the appointment scheduling rules. The performance measure used to evaluate the effect of any changes was the number of patients waiting in various parts of the centre. Overall, the simulation performed here proved to be very valuable for the client, as they were able to use it to evaluate the changes detailed above, and see how the new centre would cope.

Davies [16] describes the use of DES in testing two different hospital policies in an ED in England. The first policy is called 'See and Treat', and this is where patients with minor ailments are seen by one doctor or ENP (Emergency Nurse Practitioner, a highly qualified nurse) who sees, treats and discharges the patient in one go. The other policy is 'See' and 'Treat', and this is where patients with minor ailments are seen by a doctor or an ENP, but are then treated by a less qualified (and less well paid) nurse. At that stage in this particular ED, the ENPs were attempting to operate 'See and Treat', but the doctors were not. The proposed new system would change the arrangements so that doctors and ENPs would both operate in the same way.

Simulation models of the current and proposed systems were built in Simul8. The data were provided by the ED and they purported to show the time at which the patient was seen at various points in the pathway. However, the data proved to only be consistently reliable for the arrival time and the leaving time. Times for the different stages in treatment had to be estimated from the records that appeared to be complete. These times were then fitted to distributions and the parameters were entered into the simulation program. There appears to be a distinct lack of confidence in the data from the authors. However, as they point out, both simulations use the same times for the different activities, so the comparative results should be credible.

The performance measures used to test the model were:

- The average time spent in the system
- % of patients with treatment completed in four hours
- The average time spent in the waiting room
- The average time spent in the doctor queue

The results indicated that the proposed system would have a substantial effect in reducing patient waiting time in the three performance measures which compared the time spent in the system. Additionally, a higher percentage of patients completed their treatment in four hours in the proposed system.

2.6 Summary

While none of the papers covered in this chapter are completely relevant to our particular problem, it is a useful exercise to try and extract the approaches which seemed to produce the best results, and would seem relevant to the work which is being undertaken in this study.

The first useful process is that of structuring the problem effectively. While effective problem solving might not solve the problem for you, it helps immensely in simplifying the problem and pointing the problem solvers in the right direction. Ashton *et al* [15] used a very simple problem solving approach; performing a stakeholder analysis and system description. This simple problem structuring was nevertheless very effective in giving them a direction in which to work. Lehaney *et al* [3] used a more sophisticated approach in SSM which structures the problem more effectively, but also takes a lot more time. It would be helpful to try and find a compromise between the two in any problem structuring approach. Linked into the problem structuring part of the problem is that of creating flow diagrams. Many of the researchers chose to create flow charts representing the system before recreating this as a computer simulation. This is useful as it helps the researchers as well as their clients to visualise the system which they are modeling.

Another approach which was effectively used in six of the papers (Wijewickrama *et al* [2], Harper *et al* [14], Garcia *et al* [10], Gunal *et al* [11], Ruohonen *et al* [21], and Kirtland *et al* [1]) was that of designing and testing different scenarios within the simulation models. The scenarios ranged from different appointment scheduling rules to different staffing arrangements to different operating hours. The type of scenario tested did not really matter. The important lesson to take from this is that this is one of simulation strengths; the ability to test a range of different scenarios and compare the results without having to physically alter anything.

Linked in with the comparison of various results are the performance measures which are chosen to compare them. Different performance measures used included: the average wait of a patient in the waiting room, the percentage of patients seen within a time limit, the cost of implementation of a new system, and the effect any new system would have on other patients in the unit. Clearly, the performance measures used are dependent on what is being tested. However, every study was concerned with the average time spent by a patient in the unit, so this would seem to be the most important performance measure.

Another important issue to emerge from this chapter is that of data collection. In three of the studies (Wijewickrama *et al*[2], Harper *et al*[14], and Kirtland *et al* [1]) the data collection was not a problem. This was due to either having access to meticulously kept records, or to having large teams working on data collection over a reasonable period of time. All of the other studies had access to limited data, and had to estimate some of it themselves. These estimations were generally based on discussion with a senior staff member familiar with the system, or by simple observation and then extrapolation. In these cases, service times are often estimated by Uniform distributions. Obviously this means that the model which is built will not be a perfectly accurate representation of the reality. However, if the data are carefully collected and checked, then the model produced should be adequate for the purposes of research. This is particularly relevant to the research being undertaken here, as this takes place in a development context where data are notoriously hard to come by.

Chapter 3

Problem Structuring

3.1 Introduction

In Chapter 2, mention was made of a study done by Lehaney *et al* [3]. In this study, the researchers noted that most of the previous OR interventions done in OPDs and EDs had involved some sort of 'ad-hoc' problem structuring, but that there was very little formal problem structuring in evidence. In their intervention, they sought to first structure the problem clearly using Soft Systems Methodology (SSM), and then attempted to tackle the resulting problem using simulation. This approach of combining methods is used widely in OR interventions. Mingers and Ormerod have both always been strong proponents of using more than one method when tackling a problem, arguing that some methods are better suited to different parts of the problem-solving process [19].

SSM is designed to address 'messy' problem situations which are unstructured, and in which problems may be unidentified or viewed in many different ways by the problem participants. SSM "articulates a process of enquiry which leads to the action, but that is not an end point unless you choose to make it one" [4]. SSM is a learning process, and can be visualised as a sequence of stages, some of which can stand alone. Clearly, it would seem that SSM is a technique which would be strong at the problem structuring phase of the problem-solving process. However, when it was used by Lehaney *et al* [3] in their intervention, they found that while it was very helpful, it was also very time-consuming, and landed up costing their clients a lot of money.

It is not always necessary to adopt the whole SSM approach in order to structure a problem in an adequate fashion. Mingers [19] reasons that it is fairly straightforward to detach pieces of methodology at the level of techniques. He says that this "is particularly useful when enhancing a whole methodology...with techniques from another". Indeed, one of the examples which he uses is decomposing SSM, in particular the *root definition* technique along with its 'lower level tool' *CATWOE*. In other words, it could be advantageous to use root definitions to structure parts of the problem without having to do a complete SSM intervention.

Therefore, it seems as though it would be beneficial to the problem-solving process if we were to attempt to formally structure the problem. However, for the reasons cited by Lehaney *et al*, a full SSM intervention will not be attempted. Nevertheless it seems that it would still be beneficial to adopt some of the techniques of SSM in order to help us to understand the problem better.

3.2 Brief explanation of the methods used

A full SSM intervention consists of three stages or analyses: Analysis One (the intervention itself); Analysis Two (social); and Analysis Three (political). For the problem which is faced here, we will be concentrating on parts of Analysis One and some of its associated tools and techniques. A brief explanation of the approach to be adopted follows.

Analysis One aims to bring together three elements - the methodology, the use of the methodology by a practitioner, and the situation - in a particular relationship. The practitioner will adapt the principles and techniques of the methodology to organise the task of addressing and intervening in the situation, aiming at taking action to improve it [4]. This is done by taking a careful look at the situation faced, and then looking to define important roles which will help in the problem solving process. These are as follows:

- Client - the person who caused the intervention to happen

- Practitioner - the person conducting the investigation

- Issue Owners - people concerned about or affected by the situation and the outcome of the effort to improve it

After these roles have been identified, the practitioner will look to construct *root definitions* for each issue owner. Root definitions of the system are summaries of the problem as seen from a specific worldview. The most insight to a problem is seen to be generated by entertaining many different possibilities, and root definitions enable us to compare these possibilities.

There are various ways in which to construct root definitions, and a useful tool to use is to consciously consider the elements of the mnemonic CATWOE which is explained in Figure 3.1 (adapted from [4]).

C	Customer	The victims/beneficiaries of the activity
A	Actor	Who would do the activities?
T	Transformation process	What is the purposeful activity?
W	Weltanschauung	View of the world which makes this definition meaningful
O	Owner	Who could stop this activity?
E	Environmental Constraints	What constraints in the environment are taken as given?

Figure 3.1: Formulation of Root Definitions

After root definitions have been constructed, the SSM learning cycle encourages the analyst to ask what measures of performance could best be used to judge the operation of the notional system. By consciously considering these criteria, the analyst can then judge the system which they have created [4]. The relevant criteria are as follows:

- criteria for the efficacy of the transformation
- criteria for the efficiency of the transformation
- criteria for the effectiveness of the transformation

Therefore, for our current problem, we will aim to define the roles of each person involved in the problem, and then look to construct a root definition for each issue owner. Then, we will look at some criteria by which we can judge our system. This will give us a clearer idea of the problem which we are faced with, and should help us to begin to formulate a path to solve it.

3.3 Applying the methods

The ED at Groote Schuur has always struggled to move patients through in a timeous fashion. Patients who do not have serious complaints often have a waiting time of more than 6 hours, and even those with serious problems sometimes have to wait for up to an hour, causing delays and congestion all around. Like any long-lived problem, everyone involved has their own opinion as to what the issues are. It is not possible to simply add more space or staff to the current ED, as those resources are not currently available. So basically, everyone can agree that a problem with excessive waiting time and congestion exists, and that something needs to be done about it using the existing resources of the department.

An initial meeting was held with the Chief Operating Officer (COO) and the Chief Logistics Officer (CLO) at Groote Schuur in order to discuss the problem at hand. At this meeting, a basic plan of action was formulated, and we were granted access to the ED so that we could begin to see what the issues were. After the meeting, we spent a lot of time in the ED meeting some of the staff members and patients, and observing the inner workings of the department. This observation and interaction enabled us to identify the stakeholders and some of the more pertinent issues. From here, we were able to start building up an initial structure for the problem.

The issue owners are as follows:

- Dr Brey *CLO, Groote Schuur Hospital*
- Dr Linda *COO, Groote Schuur Hospital*
- Emergency Department Staff
- Emergency Department Patients

The clients are:

- Dr Brey
- Dr Linda

The practitioner is:

- Allister Mowbray, aided by Professor TJ Stewart.

3.3.1 Root definitions

For this research, a slightly different approach was adopted in order to obtain the root definitions for the issue owners. The CATWOE approach was still used, but seeing as many of the elements for all of the issue owners were very similar, it was decided to rather concentrate on these differences.

Let us first consider an abridged CATWOE diagram which will illustrate the shared elements between the CLO, the ED staff, and the ED patients. This is shown in Figure 3.2.

C	the patients in the ED
A	the staff in the ED
T	inefficient system $\longrightarrow$ efficient system
O	the hospital management
E	the existing physical structure within the ED

Figure 3.2: Common CATWOE elements for CLO, ED staff, and ED patients

The only element which has been excluded from this diagram are the *weltenschauung* or worldview of each issue owner. This is probably the most important part of the diagram, because each person's worldview makes their definition of the problem unique to themselves. However, this diagram shows us that the other elements of the problem are shared by the majority of the issue owners for this problem. All of these issue owners want to improve the ED from an inefficient system to an efficient system, and they are all faced with the same customers, actors, owners, and environmental constraints.

Even though these issue owners agree on the structure of the problem, their worldviews are all still quite different, and this will result in them all having different root definitions. This will be illustrated in the following section.

The CLO

The CLO was our initial contact at Groote Schuur who alerted us to the problem which existed. He has been involved in various efforts to improve efficiency within the hospital, but lacked the necessary time and skills to tackle this problem. Although he does not work directly in the ED on a day-to-day basis, he has been involved in efforts to improve the existing patient care by trying to restrict the number of patients entering the ED, and by moving around the current resources. He is concerned with reducing waiting times for patients, and improving conditions for staff and patients by reducing the congestions within the ED.

All of these concerns lie within a broader view of improving efficiency in the hospital in general, and his weltenschauung should be: "a decreased waiting time and decreased congestion in the ED and the rest of the hospital is desirable". Therefore a root definition for the CLO could be defined as follows:

A hospital-owned and staffed system to deliver efficient service with the minimum of patient waiting time and reduced congestion, in order to provide quality medical care, an improved working environment, and to provide a benchmark which other systems within the hospital can be compared to.

The ED staff

The ED staff refers to the doctors and nurses who work in the department on a daily basis. It has been decided to group all of the staff together because, after speaking with a number of them, it became clear that their issues were all much the same. The staff are concerned with the quality of service which they are able to offer to the patients. Additionally, a longer waiting time for the patients means a more congested ED which makes their work harder to do. This affects the quality of their service, and it also creates a stressful working environment. A few of the staff had suggestions on ways in which they thought that the ED could be improved which demonstrates that they want to be proactive in finding a solution to the issues at hand. The problem is that it is difficult for them to find the source of these issues.

So, the weltenschauung for the staff should be: "any work that allows staff to provide quality care to the patients while clearing congestion and providing a less stressful working environment is desirable". Then, the root definition for the ED staff could be defined as follows:

A system to deliver efficient service with the minimum of patient waiting time and reduced congestion, in order to provide quality medical care and a less stressful working environment.

The patients

The patients are the people who come into the ED off the street on a daily basis. Their primary concern is to receive quality medical care with the minimum of waiting time. They are only concerned with the efficiency of the system insofar as it relates to their experience within the ED. Therefore, if they receive prompt and efficient service then they are happy, but they probably are not really concerned about other patients' experiences.

So, the weltenschauung for the patients should be: "any decrease in waiting times for patients is desirable". Then, the root definition for the patients could be defined as follows:

A system to deliver efficient service with the minimum of waiting time for patients in order to provide quality medical care for them.

The COO

Defining a root definition for the COO has been left until last. This is because both her environmental constraints, as well as her weltenschauung will be different from the other issue owners. As the COO for Groote Schuur, her outlook on the problem is how it fits into the bigger picture.

She is obviously concerned with the working conditions of her staff, and would like the ED to become less congested. However, her primary concern is with moving the patients through the system at a faster rate than is currently happening. This is because reduced patient waiting times will have more of an effect on improving the service which the hospital offers than reduced congestion would.

As she is more intimately involved in budgetary and staffing matters than any of the other issue owners, she would also be more concerned than any one else that any solution must not exceed current resource use. Also, seeing as she does not work in the ED on a daily basis, her concern for the day-to-day running of any solution would not be as high as amongst the other issue owners.

Finally, the biggest difference between the COO and the other issue owners is that she operates in a wider system. While most of the issue owners only operate within the walls of the ED or, at most, the walls of the hospital, she must be concerned with how the hospital as a whole operates within the provincial or even national health system. Therefore, any solution needs to be taken in this context.

The CATWOE elements for the COO could be as in Figure 3.3.

C	the patients in the ED
A	the staff in the ED
T	inefficient system $\longrightarrow$ efficient system
W	a decreased patient waiting time in the ED using current resources and complying with current provincial/national health guidelines is desirable
O	the hospital management and provincial health authorities
E	the existing physical structure and resources within the ED and the existing provincial/national health guidelines

Figure 3.3: CATWOE elements for the COO

Then, a root definition for Dr Linda could be defined as follows:

A hospital-owned and staffed system to deliver efficient service with the minimum of patient waiting time, using current resources, within provincial/national health guidelines in order to provide quality medical care and to reflect well on the hospital as a whole.

3.3.2 Summary of analysis

The CATWOE and root definition analyses revealed that, although the stakeholders have different worldviews, their basic outlook is the same. This is that it is beneficial and desirable to everyone to transform the ED from an inefficient system into an efficient system.

As revealed previously in the chapter, the differences in problem perception between the issue owners arise primarily from their different worldviews and the environmental constraints which each of them are subject to. This difference in worldviews is an obvious point and has already been dealt with. However, taking a closer look at the different environmental constraints which each stakeholder is subject to could be quite revealing.

The staff, including the CLO, have the current physical structure of the ED as an environmental constraint. The COO has this same constraint coupled with an additional constraint of not exceeding current resource use and acting within governmental guidelines. The patients have the environmental constraint of having to accept whatever changes anyone else chooses to make in the ED.

This means that, it should be taken as a given that the actual physical structure of the ED cannot be altered insofar as building a new department or making major renovations within the existing department. However, what could possibly be altered is the layout within the existing ED. Obviously, for the staff members, resources could also be added in the form of more staff members or more equipment or both. This is because they do not have to deal with the repercussions, monetary or otherwise, of such a decision. However, for the COO, this becomes a thornier issue, as these are problems which she would have to deal with.

Now, we can begin to consider some criteria for measuring the efficacy, efficiency, and effectiveness of the system. Criteria for efficacy tell us whether the transformation is producing its intended outcome [4]. The intended outcome is to improve efficiency by reducing waiting times and congestion. Therefore, if we were to measure waiting times within the system, and these were substantially reduced, this would be an indication that the system was efficacious. The reduction of congestion would be more difficult to measure, and could only really be achieved by talking to the staff who work in the ED.

Criteria for efficiency tell us whether the transformation is being achieved with the minimum use of resources [4]. For the problem at hand, this would mean that any changes which are being proposed should not exceed current resource use.

Criteria for effectiveness tell us whether the transformation is helping to achieve some higher-level aim [4]. The higher-level aim for this problem would be to improve the service being offered by the hospital. If the criteria for efficacy and efficiency are met, whilst still providing quality care to the patients, this would mean that the overall service of the ED has been improved, which would necessarily mean that the service being offered by the hospital as a whole has been improved.

Therefore it would seem that, by reducing waiting times within the ED using existing resources, we would be satisfying all three criteria measuring the performance of the model. At this stage it becomes clear that, while the root definition and CATWOE analyses helped to define the boundaries of the problem, the usefulness of these techniques ends here. An approach is needed which will help us to evaluate different alternatives within these boundaries, while not actually physically altering anything. After analysing the problem and reading extensive literature on similar interventions, it becomes clear that discrete event simulation is the most appropriate technique to use in order to tackle this problem further.

Chapter 4

Data Collection

4.1 Physical layout of the emergency department

The ED in Groote Schuur is located in the New Main Building, and the physical layout is depicted in Figure 4.1.

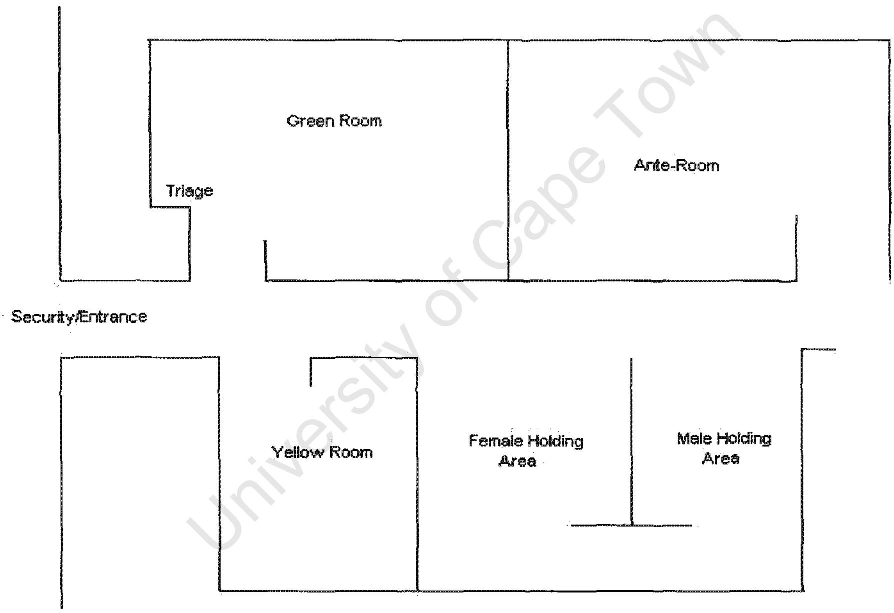

Figure 4.1: The physical layout of the Emergency Department

All patients enter past the security guards who are instructed only to allow emergency patients and medical staff into the ED. Ambulance patients are taken straight to the ante-room, but everyone else has to go through triage, which is a little section of the green room. Once patients are triaged, they then either remain in the green room, or get sent on to the yellow room or the ante-room depending on their condition. The green room and the yellow room are for less serious patients, and contain only chairs for the patients to sit and wait on. The ante-room is for more serious patients and contains only trolleys. The holding area is a temporary area for patients who are either waiting for test results or waiting to be admitted into a hospital ward. It is broken up into separate male and female sections. The holding area is equipped to contain both chairs and trolleys, with the less serious patients sitting in the chairs, while the ante-room contains only trolleys. The holding area is equipped to hold 6 trolleys, and the ante-room 10 trolleys. However, in reality there are often 14 or more trolleys in each of these areas.

4.1.1 Patients in the emergency department

Through observation of each area of the ED, and also through discussion with staff, we were able to obtain average counts for the number of patients in each area. Obviously, this number varies throughout the day, and certain areas are busier at certain times of the day. However, the following table merely gives an average number of patients throughout the day which is sufficient for our modeling purposes.

Area	Ave. no. patients
Green	6
Yellow	6
Ante	16
Holding	14

4.2 Patient flows in the emergency department

As mentioned previously, all patients entering the ED, except for those arriving by ambulance, must pass through the triage area. It is here that a nurse will attend to the patient; their vitals are checked, then the nurse makes a quick assessment of the acuity of the patient's condition, and assigns them a colour code depending on this level of acuity. There are four colour codes representing different levels of acuity:

- Red - the most severely ill patients. These patients are generally experiencing seizures, have burns to the face or have inhaled smoke, or have severely low glucose levels.

- Orange - the severely ill patients. These patients generally have compound fractures, bad burns, have been vomiting blood, have experienced a reduced level of consciousness, or are experiencing a psychosis.

- Yellow - the moderately ill patients. These patients generally have controlled haemorrhages, dislocations, closed fractures, or have been experiencing vomiting.

- Green - the mildly ill patients. These are patients experiencing mild illness symptoms who do not fall under any of the above categories.

4.2.1 Red and orange coded patients

As mentioned above, these are the most severely ill patients who enter the ED. Red and orange coded patients have been grouped together because, although the patients are experiencing different illness acuity levels, the patient flow and treatment is essentially the same for both groups of patients. These patients either arrive by themselves or in an ambulance. However, after consultation with the CLO at the hospital, it was decided to not include ambulance arrivals as a separate entity, as this would not add anything to the model. After the initial triaging, these patients are sent to the ante-room. The ante-room is equipped to hold 10 trolleys (although as previously mentioned this number is often exceeded), and there are always at least 2 doctors on duty here. Patients in the ante-room are attended to according to how severe their illness is, with the unstable patients being treated first.

All patients entering the ante-room are entered into the ante-room treatment book, and given an initial examination. After this, there are a variety of ways in which a patient can be treated. However, after speaking with ED doctors, it was decided to include only the most common treatment paths in the interests of simplicity. Therefore, after their initial consult, ante-room patients will either require bloodtests, x-rays, x-rays and bloodtests, or they will require no further testing. Stable patients can then be discharged straight from the ante-room, but they are generally sent into the holding area first, either to wait for their test results or to wait for a specialist. Patients who go to the holding area from the ante-room will generally be admitted to a ward after they have seen a specialist or have received their test results, as these tend to be the patients requiring more serious treatment. However, some of them will go home from the holding area if the results of their tests are favourable.

Figure 4.2 is a flowchart which shows the flow of ante-room patients through the ED. The *Treatment Finished* block is not another treatment

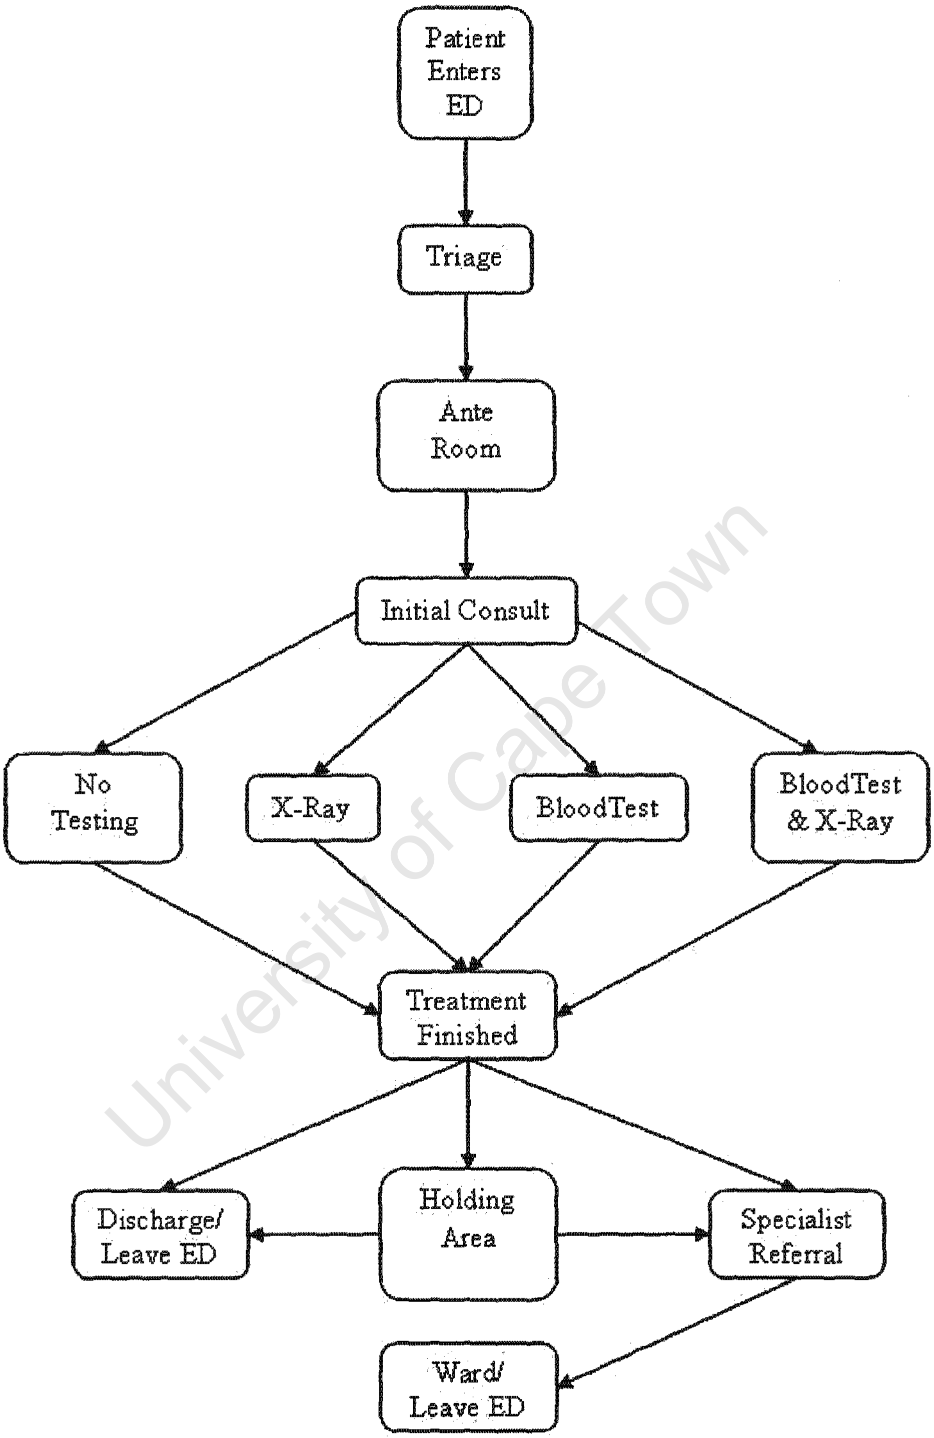

Figure 4.2: Flowchart depicting red/orange coded patients' movements through the ED

which the patients undergo, but is rather included for the purpose of making the flowchart more aesthetically pleasing. This flowchart is intended to only model flows within the ED, so there is only one entry point into the system, and the only exit points are for a patient to be admitted into a ward, or to be discharged. The event of a patient dying has been included in the *Discharge* exit.

In this flowchart, red and orange coded patients are treated together, because they have the same basic flow within the system. However, in the model, there will be a separate initial consult for red and orange coded patients. This is because red coded patients are much more critical than orange coded patients and need to be seen as soon as possible. Orange coded patients, while also critical, can be left waiting for longer than red coded patients. Also, the treatment of red coded patients is normally longer than that of orange coded patients. After the initial treatments, their flows are exactly the same.

4.2.2 Yellow coded patients

The yellow coded patients are those who are moderately sick. All of these patients arrive in the ED through their own means of transport (no ambulance arrivals). After the initial triaging, these patients are directed to the yellow waiting room. This room is operational between 08:00 and 23:00 every day. If a yellow coded patient arrives outside these hours, they are sent to the ante-room which is operational 24 hours a day. All yellow coded patients are then entered into the yellow room treatment book. They then wait on chairs to be attended by the doctor. Patients are seen on a first-in first-out basis, and there is one doctor attending to patients in the yellow room.

After an initial waiting period, yellow coded patients are seen by the doctor in the yellow room, who then determines the requirements of the patient. Again, there are a number of treatment paths which could theoretically be taken by the patient at this stage. However, after speaking to the ED doctors, it was decided to only include the most common treatments in the interests of simplicity. Therefore, at this stage in our system, yellow coded patients can go for x-rays, bloodtests, or a specialist referral. Yellow coded patients who are referred to a specialist will either come back at a later stage to see the specialist and hence leave the ED system (less severe patients), or they will wait in the holding area for the specialist (more severe patients). The patients who are sent to the holding area to wait for a specialist are generally admitted to a ward at a later stage. Patients waiting for x-ray or bloodtest results will either wait in the yellow room or in the holding area depending on the severity of their condition.

Figure 4.3 is a flowchart which shows the flow of yellow coded patients through the ED. Once again, the *Treatment Finished* block is included for the purpose of making the flowchart more aesthetically pleasing, and is not intended to be interpreted as another treatment.

4.2.3 Green coded patients

The green coded patients are the least ill out of all the patients in the ED. As with the yellow coded patients, these patients will all arrive in the ED through their own means of transport. After the initial triaging, these patients are directed into the green waiting room. This room is operational between 10:00 and 18:00 every day. Green coded patients who arrive outside of these hours are either directed to the yellow room (during its operating hours), or to the ante-room (when the yellow room is also closed). All green coded patients are then entered into the green room treatment book, and they wait on chairs to see the doctor. Patients are seen on a first-in first-out basis, and there is one doctor attending to patients in the green room.

After an initial waiting period, green coded patients are seen by the doctor in the green room, who then determines the requirements of the patient in a similar fashion to that of the yellow patients. Once again, there are a number of possible treatment paths from here, but for simplicity's sake only three have been included in the model (x-rays, bloodtests, or a specialist referral). However, because the green coded patients are not suffering from a severe illness, all green coded patients who are referred to a specialist will return to see the specialist at a later stage. Therefore, they are for all intents and purposes out of the system after being referred to a specialist, and so this option is not included in the flowchart. For the same reason, any green coded patient who undergoes a bloodtest will come in at a later stage to receive their results, rather than wait around and clog up the ED.

Figure 4.4 is a flowchart which shows the flow of green coded patients through the ED. As explained before, there are essentially only two treatment paths for green coded patients in the system: *Initial Consult* by itself and *Initial Consult* followed by *X-rays*. The initial consult here would also include any bloodtests. This is because the actual taking of blood for a test does not take long and can be included as part of the consult. As green coded patients do not wait in the system for their results, an explicit path

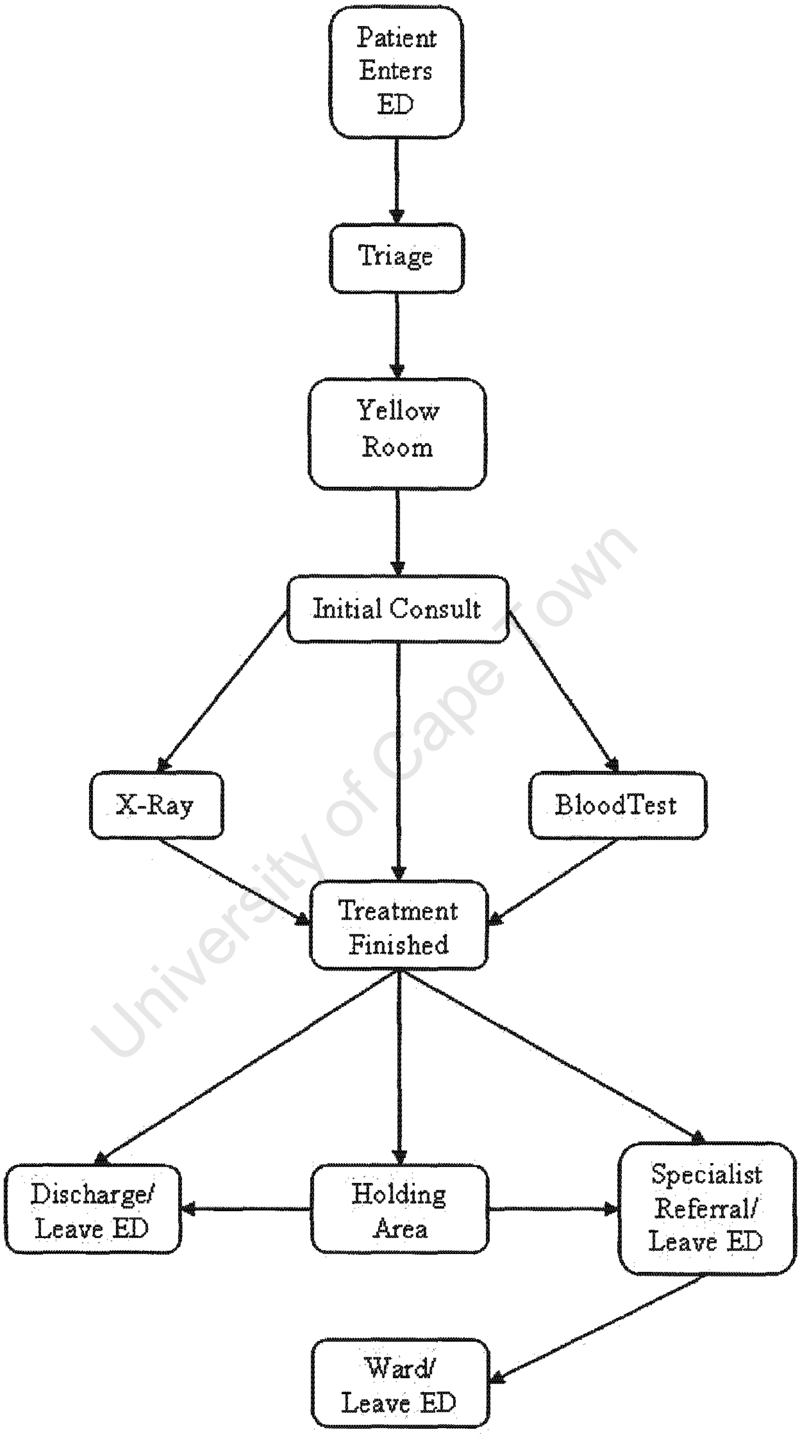

Figure 4.3: Flowchart depicting yellow coded patients' movements through the ED

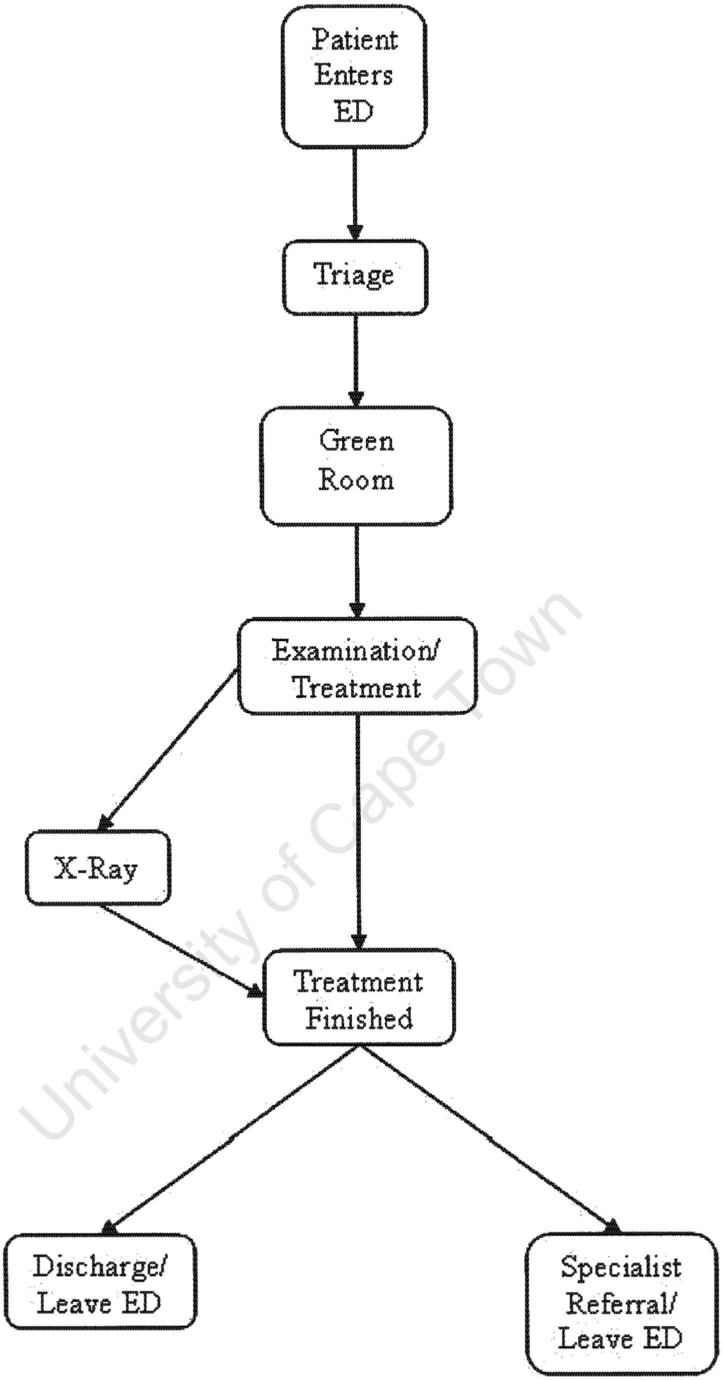

Figure 4.4: Flowchart depicting green coded patients' movements through the ED

for patients who are waiting for their test results does not need to be included here as it needs to be for the other acuity levels.

4.3 Patient arrivals

In any ED, the patient arrivals are the most critical unknown variable in the system. If there are a large number of patients entering the ED, and they can be said to be acting more or less independently, then the number of patients entering the ED during a unit of time could be modeled by a non-homogeneous Poisson process with λ dependent on the time of day [18].

Data were made available from Groote Schuur Hospital which detailed the patient arrival numbers by the hour for a period of 1 month. These data were collected over the month of July, 2009. Using this information, we were able to estimate a value for λ for every hour of the day by using the average arrivals for that hour over the month.

The data analysis began by plotting the estimated values of λ for each hour as a continuous process as shown in Figure 4.5. As can be seen in Figure 4.5, there is a definite trend in the arrival patterns. Arrivals slow down slightly after 00:00 until about 07:00 where they start to pick up, exhibiting a steady growth pattern until they peak at 11:00 where λ is exactly 8. After 11:00, the arrivals exhibit a steady decline at a substantially reduced rate than the initial upward trend. There is no evidence of any major peaks during this downward trend with only a brief increase at about 20:00.

This behaviour is fairly simple to explain. Firstly, as can be expected, there is little traffic outside of business hours. This can be explained by two main reasons: firstly, Groote Schuur is situated quite far away from its "target market", and so people are only going to go there at night time if they are really ill; secondly, the type of people who rely on Groote Schuur are generally the same people who rely on public transport which is generally non-existent at night time in Cape Town. Therefore, at the beginning of the day, there is a marked increase in the number of arrivals, peaking at about 11:00. After this peak, there is a steady drop off in arrivals with only a minor peak at about 20:00. This is about the time that public transport starts shutting down and would possibly represent people coming in after work but before public transport closes.

Looking at Figure 4.5, the arrival pattern transitions fairly smoothly from one value to the next, apart from the anomaly at 20:00. This indicates that the time unit of 1 hour is adequate for the purposes of this study.

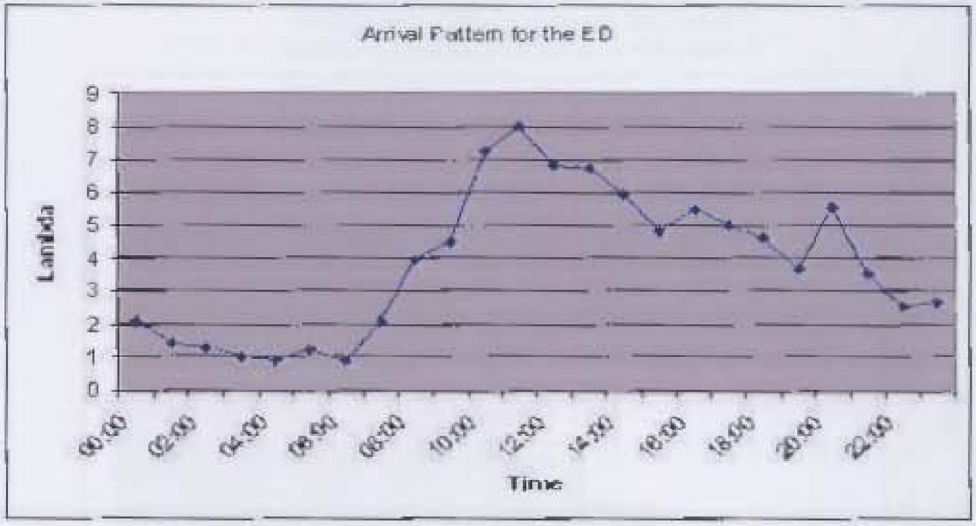

Figure 4.5: The estimated values of λ for each hour

If there were lots of areas where the transitions were not smooth, it would indicate that the time unit was too large, and a smaller unit would be needed.

The next step was to construct graphs comparing the actual cumulative distribution against the theoretical cumulative distribution in order to get some feel for how the data fit the proposed distributions. The details of all of this are included in Appendix A, but an example of the Statistica output for the 12:00 increment is given in Figure 4.6.

Finally, a more statistically formal approach was adopted. Using Statistica, we then tested the goodness of fit using a Pearson χ^2 goodness-of-fit test.

The Pearson χ^2 statistic is as follows:

$$\chi^2 = \sum_{i=0}^{n} \frac{(O_i - E_i)^2}{E_i} \tag{4.1}$$

where
O_i =No. of patients arriving each hour in the system
E_i =Expected values according to the Poisson model with fixed λ
n = Days in July 2009

The null distribution of the test statistic can be shown to be approximated by the χ^2 distribution with degrees of freedom (df) being given by:

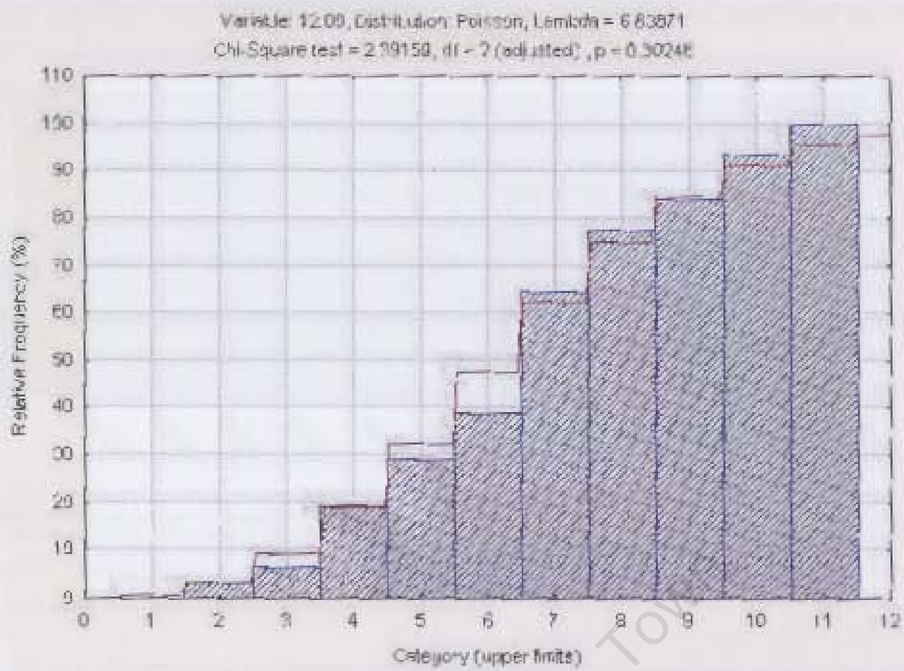

Figure 4.6: Actual cumulative distribution vs theoretical cumulative distribution for 12:00

df =number of cells - number of independent parameters fitted - 1

The 'number of cells' above refers to the number of occurrences of each number of patients which arrived in the given hour. This is explained in more detail in Appendix A. If the p-value for χ^2 is low, this is an indication that the null hypothesis (which in our case is that the data follow a Poisson distribution) should be rejected [18]. Table 4.1 is a summary of this data analysis, and includes the value of λ for each hour with the χ^2 values and corresponding p-values.

Looking at the p-values in Figure 4.1 for the χ^2 statistic, it would seem that the assumption that the arrivals follow a Poisson process is justified. Only one observation has a p-value of less than 0.1. As statistical logic dictates that in a test of this size at least one observation should be significant, we should not be concerned with this.

It should be mentioned at this point that these arrival data were collected as part of a CLINICOM survey. After consulting with the ED staff and the CLO, it was estimated that these arrivals underestimated the true arrivals by some 30%. Therefore, when building the actual model, these values will

Table 4.1: The estimated values of λ for each hour

Hour	λ	χ^2	p-value
00:00	2.06	1.01	0.32
01:00	1.42	0.52	0.77
02:00	1.29	0.60	0.44
03:00	1.03	0.98	0.32
04:00	0.94	1.93	0.16
05:00	1.26	0.15	0.70
06:00	0.87	0.79	0.37
07:00	2.10	1.20	0.27
08:00	3.93	2.78	0.25
09:00	4.48	1.85	0.60
10:00	7.26	1.35	0.51
11:00	8.00	2.33	0.31
12:00	6.84	2.39	0.30
13:00	6.71	0.51	0.77
14:00	5.94	4.89	0.09
15:00	4.84	2.37	0.31
16:00	5.48	0.90	0.64
17:00	5.03	0.30	0.86
18:00	4.65	1.23	0.54
19:00	3.65	1.94	0.38
20:00	5.55	3.15	0.21
21:00	3.52	2.34	0.31
22:00	2.55	0.12	0.94
23:00	2.71	3.40	0.18

be increased by 30% in order to more accurately reflect reality.

4.4 Patient splits

Referring back to Figures 4.2, 4.3 and 4.4, it can be seen that there are a number of occasions where the patients are split up according to their proposed treatment or path in the system. The first split happens after the patient is triaged, when the nurse decides how ill they are, and assigns them a colour code accordingly. Then, looking at Figure 4.2, ante-room patients get split after their initial examination, when a doctor decides on their course of treatment. After this treatment, there is another split when patients can either get discharged, have a specialist referral, or go to the holding area. Similarly for yellow and green coded patients, there is a split when the doctor decides on their treatment path, and there is a further split when they are either sent home, referred to a specialist, or sent to the holding area (for yellow coded patients only).

Earlier in this chapter, reference was made to an 'ante-room treatment book', a 'yellow room treatment book', and a 'green room treatment book'. Additionally, there is also a 'holding area treatment book'. These are books which reside in each of these particular rooms which contain patient particulars, illness details, and their destination when they leave the room. The details of these books were made available in order to help this research. However, all of the patient details are not always entered into these books. For example, it is very rare for all the details of a patient's treatment to be recorded here. Nevertheless, they do contain a record of every patient passing through every room in the ED, and they also generally contain details about where they are sent. We were able to use information from these books for the first week of July 2009 in order to estimate percentages for some of the patient splits mentioned before. When there was no information about where the patient was sent, that patient was ignored, as it is reasonable to assume that these omissions were done on a completely random basis with no inherent bias. It would also not be possible to include a path for these omitted patients in the final simulation model.

Patient Destination after Ante-Room	
Discharge	Holding Area/Specialist Referral
30%	70%

The table above lists the destinations of patients from the ante-room broken up into percentages. These data were for the first week of July, and the details were taken from 163 patients. There were 322 patients who went through the ante-room in that week in a ratio of 86:14 for orange coded patients:red coded patients. However, as mentioned before, there was no information of patients' destinations for a number of patients, and so these were not included in the final percentages. The event of a patient dying in the ante-room has been included under *Discharge*. Additionally, as the patients wait for specialists in the holding area, these 2 destinations have been grouped together.

Patient Destination after Yellow Room			
Discharge	Specialist Referral	Holding Area	Ante-room
49%	16%	16%	19%

The above table lists the destinations of patients from the yellow room broken up into percentages. These data were for the first week of July, and the details were taken from 150 patients. There were a total of 219 patients who went through the yellow room during this period, but there was no record of a destination after the yellow room for 69 of them.

Patient Destination after Green Room	
Discharge	Specialist Referral
89%	11%

The above table lists the destinations of patients from the green room broken up into percentages. These data were for the first week of July, and the details were taken from 174 patients. There was a record of every patient's destination in this book, so there was no need for any omissions.

Patient Destination after Holding Area	
Discharge	Ward Admission
50%	50%

The table above lists the destinations of patients from the holding area broken up into percentages. These data were for the first week of July, and the details were taken from 108 patients. There were a total of 218 patients passing through the holding area in this period, but there was no record of a destination after the holding area for 110 of them. It was also observed that over the time period in question, 32% of all patients in the ED passed through the holding area at some stage.

Total ED Patient Split			
Red	Orange	Yellow	Green
6%	36%	32%	26%

The above table lists the percentage of patients who were classified into each acuity level during the first week of July. These data were simply the number of patients in each room's book for the period in question. There was no question of a patient's records being omitted here, as all that was required was a record of the patient having passed through the room.

As there were no details for the treatment of patients in these books, it was necessary to estimate the patient splits for the various treatments which could be undertaken. This was done after careful consultation with some of the doctors in the ED. The following table lists these treatment splits:

Room	Treatment	X-Ray	Blood Test	X-Ray & BloodTest	Total
Green	60%	40%	-	-	100%
Yellow	50%	15%	35%	-	100%
Ante	10%	20%	25%	45%	100%

4.5 Service times

The last important piece of information needed in order to build a model of the ED are the service times for each activity within the department. These proved to be the hardest data to find. However, after consultation with the doctors, we were able to establish minimum and maximum treatment times for some of the service points which would enable us to construct uniform distributions around the times. For the activities which could be considered as 'pure waiting times', such as waiting for x-ray results, waiting for bloodtest results, and waiting in the holding area, we had to use the average waiting times for each of these.

Activity	Time
Triaging	2 - 7 minutes
Green Treatment	10 - 30 minutes
Yellow Treatment	15 - 40 minutes
Orange Treatment	30 - 150 minutes
Red Treatment	30 - 180 minutes
Holding Area Wait	780 minutes
X-Ray Wait	30 minutes
Bloodtest Wait	290 minutes

4.6 Summary of data

Looking at the data collected and described in this chapter, it is obvious to see that they can be split into two categories: accurate, verifiable data; and data based on expert opinion, but which is not directly verifiable. The arrival pattern, patient paths, and the patient destination splits fall under the first category, while the treatment splits and service times fall under the second. The average number of patients in each room within the ED was collected through personal observation, but as this only took place over one day and through discussion with staff, this should be regarded as non-verifiable data.

We have striven to make the non-verifiable data as accurate as possible, but we can never be absolutely sure that this is the case. Therefore, when looking ahead to model verification and experimentation, we should try to test the effects that any changes in these values will produce in the model.

Chapter 5

Model Building

5.1 Justifying the choice of simulation package

After the basic problem structuring stage of the problem solving process, it was decided that a logical choice of technique for the actual problem solving stage would be discrete event simulation (DES). In 1999, Jun *et al* [6] surveyed over 100 papers where DES had been used in health care clinics. They describe DES as "one tool available to health care decision-makers that can assist in examining new ways to improve efficiency and reduce costs. DES is an OR technique that allows the end user .... to assess the efficiency of existing health care delivery systems, to ask 'what if' questions, and to design new systems. DES can also be used to forecast the impact of changes in patient flow, to examine resource needs ...., or to investigate the complex relationships among the different model variables." Their paper goes on to describe more than 100 successful interventions in health care clinics using DES.

The interventions described in Jun *et al* all used different forms of DES, as simulations can be modeled in a variety of ways, including using programming languages or low level simulation packages. However, the problem with these approaches is that, while they may be more powerful in many ways, they are not transparent and easy to use. Visual Interactive Simulation (VIS) allows the user to see the movement of patients through the clinic on the computer screen. This is obviously very transparent, and has great practical appeal when building a model which needs to be validated by the clinic staff who are probably not familiar with DES at all. Both Harper *et al* [14] and Lehaney *et al* [3] describe using the Simul8 simulation package in their interventions. Simul8 is an easy to use, inexpensive VIS which has been used in successful health care clinic interventions in the past. As the author also had some past experience in the use of Simul8, it seemed a logical choice to use here.

5.2 Explanation and setup

Simul8 uses different five different icons in order to build up a simulation. These icons are:

- Work Entry Points
- Storage Areas
- Work Centres
- Work Exit Points
- Resources

Work entry points are locations where the patients will enter. Storage areas are areas where patients queue and wait for service. Work centres are areas where patient treatment is undertaken. Work exit points are locations where patients will exit, and resources are staff members.

A simulation can be as complex and interwoven as the modeler chooses to make it. For this particular problem, it was felt that it would be simplest to confine the simulation to the walls of the ED. In other words, any treatments or results obtained from outside the ED would be exogenous to the simulation. Additionally, as discussed in Chapter 4, certain simplifications were made to the treatment paths of patients.

For the purposes of this simulation, resources were only explicitly included in the model in areas where doctors were needed to float between work centres. In areas with only one doctor, the work centre acted as the doctor.

5.2.1 Service times and distributions

As mentioned in Section 4.5, it was difficult to obtain service times for all of the patient treatments. However, we were able to obtain estimates of consulting times for some of the treatments, and estimations of the waiting times for the rest of the service points which are provided in Section 4.5.

Additionally, we were able to obtain counts for the average number of patients waiting in each area of the ED through observation and discussion with staff which are given in Section 4.1.1.

The estimates of the consulting times were provided in the form of minimum and maximum times for each consult. As it was assumed that there was an equal probability of the consultation taking any time between this maximum and minimum, these can be estimated with a Uniform distribution.

However, in order to fit a distribution where the only information which existed was the expected average waiting time, it was decided to experiment with different distributions and parameters in order to find the most accurate combination. This would be done by running multiple (40) runs of the simulation, and then comparing the values for key performance indicators with the values obtained in reality. These indicators were: the average queuing time, and the average queue size. Additional indicators which would also be measured were: the minimum (non-zero) queuing time, and the maximum queuing time. These were included in order to see if there were obvious inconsistencies when using the distribution. The simulation would be run for a warm-up period of 2 weeks in order to reach a steady state. After that, the simulation would be run over a period of 31 days in order to collect the data.

Over the trial of 40 simulation runs, Simul8 provides an average result for each indicator, along with a low 95% range and a high 95% range. These ranges give the range in which the long term average will be on 95% of the times that the prediction is made. Obviously, the more runs which are made, the tighter these ranges become. Therefore, in order to make sure that the values which were used provided consistent waiting and service times, a consistency indicator was defined. This is known as the Result Relative Consistency or RRC and is defined as follows:

$$\text{Result Relative Consistency} = \frac{ResultAve - Low95\%Range}{ResultAve}$$

A low RRC value would indicate a more consistent result. As completely accurate distributions were not possible to obtain for any of the waiting times, it was decided that it would be easiest if the same distribution were used for each service time, with only the parameter changing. After experimenting with different Normal and Erlang distributions, it was decided to use an Erlang 15 distribution for the service times. The Erlang distribution was felt to be more appropriate than the Normal distribution as it has a positive value for all real numbers greater than 0, whereas the Normal distribution can have negative values depending on the size of the variance. There was not a big difference in consistency between the different Erlang distributions provided in Simul8, but the Erlang 15 distribution fitted the data the best out of the Erlang distributions.

5.3 The model

5.3.1 Arrivals

There is only one physical entry point into the ED, so it was decided to only use one entry point into the simulation. Chapter 4 described the fitting of a Poisson distribution to the arrival patterns experienced by the ED. Simul8 allows one to set up their own distribution, tailoring it specifically to meet the needs of the situation. It was therefore possible to set up an arrivals distribution for the ED using a Poisson distribution with λ changing with the hour. Arrival patterns in Simul8 are calculated by inter-arrival times in minutes, which would be $\frac{60}{\lambda}$. This then enabled an entry point to be set up which exactly mimicked that of the actual ED at Groote Schuur. As described in Section 4.3, it was decided to use a 30% higher arrival rate than the one recorded there in order to more accurately depict reality. The times used were as in Table 5.1.

5.3.2 Start of patient flow

After the patients have entered the ED through the entry point, they immediately enter the queue to be triaged. Triaging is the point where the patient's vitals are checked, and the nurse decides on how acute the patient's ailment is. It is a very quick process, and there are very seldom any lines formed for it. The estimation of the minimum and maximum time for triaging is given in Section 4.5. Using these values, we were able to use a Uniform distribution to model the triaging process. At this point, the patients are also split into green, yellow, orange, and red coded patients according to the percentages given in Section 4.4.

Figure 5.1 displays a screenshot from Simul8 depicting the arrival point, triage, and subsequent split. Green and yellow coded patients enter their respective queues for treatment, while red and orange coded patients enter the ante-room which in itself is a queue or storage area in the model. In the model, the orange and red parts of the ante-room are treated separately in order to allow different treatment times for each type of patient. However, in reality, there is only one actual ante-room. The dummy nodes are not relevant here, and will be discussed in the next section.

Table 5.1: The inter-arrival times for each hour

Hour	Inter-arrival Time
00:00	22.36
01:00	32.52
02:00	35.77
03:00	44.71
04:00	49.34
05:00	36.69
06:00	52.99
07:00	22.01
08:00	11.73
09:00	10.29
10:00	6.36
11:00	5.77
12:00	6.75
13:00	6.88
14:00	7.78
15:00	9.54
16:00	8.42
17:00	9.17
18:00	9.94
19:00	12.66
20:00	8.32
21:00	13.13
22:00	18.11
23:00	17.03

5.3.3 Shift times

The three different areas which patients enter after being triaged have different operating hours which need to be built into the model. These operating hours are as follows:

- Green Room - 10:00 to 18:00
- Yellow Room - 08:00 to 23:00
- Ante-Room - 24 Hours

An additional consideration is that, in reality, the yellow and green coded patients arriving from about 07:00 are redirected from the ante-room into their respective areas, and they start to form lines there. This means that

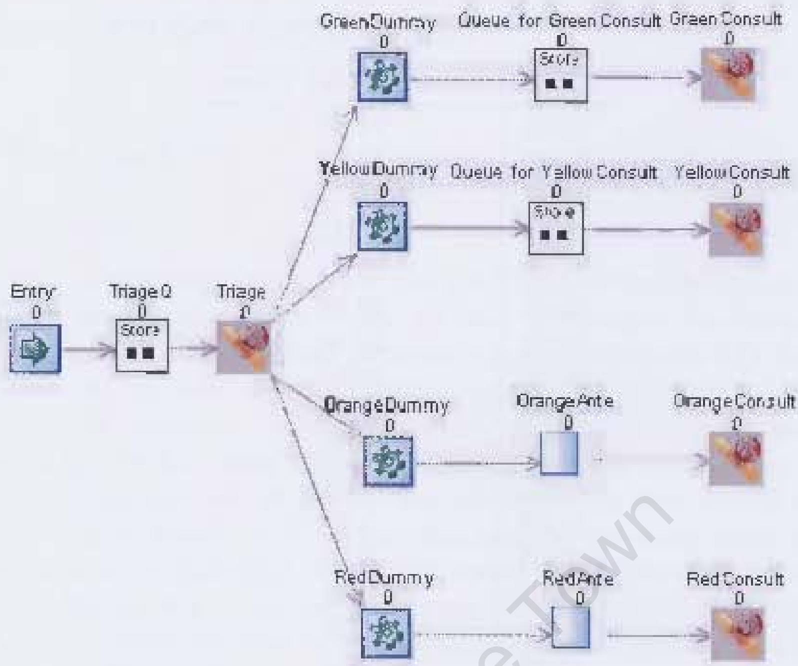

Figure 5.1: Simul8 screenshot of arrivals

there is already a line formed when the doctor arrives to start seeing patients. In order to build this into the simulation, it was necessary to build shifts into the Simul8 model. Figure 5.1 shows that, after triaging, yellow and green coded patients are sent through dummy nodes before they enter the actual queues for their relevant rooms. These dummy nodes only become operational from 07:00, allowing the initial queues in the system to be formed. Then, the *GreenConsult* and *YellowConsult* nodes are only operational during the times given above, which simulates the doctors arriving and working their shifts. When the dummy nodes are not operational, the patients are re-routed into available queues.

However, one of the issues with handling the shifts in this manner is that, when running the model, patients would build up in the yellow and green lines overnight which is not an accurate representation of reality. In order to counter this, in the model the yellow room closes at 21:00 and the green room closes at 16:00. This is actually similar to reality because this is when the triage nurses start redirecting these patients away from these rooms. In the model, the doctors are then operational for 3 hours after the official shift closing as indicated previously. This does not generally happen in reality, but this is just to make sure that any stray patients are seen to, and do not end up staying in line in the model overnight.

Another issue with incorporating the shifts in this manner is that it affected the initial patient split percentages which were discussed in Section 4.4. Because the ante-room was operating for 24 hours, the total percentage of patients which passed through it was higher than the split given in Section 4.4. In order to accomodate this, it was necessary to alter the initial patient splits until the appropriate percentages were again obtained. This was done in a similar fashion to that described in Section 5.2.1.

		Low 95% Range	Average Result	High 95% Range
GreenDummy	Number Completed Jobs	958.19	965.95	973.71
YellowDummy	Number Completed Jobs	1173.96	1182.33	1190.70
OrangeDummy	Number Completed Jobs	1321.82	1330.83	1339.83
RedDummy	Number Completed Jobs	208.92	213.63	218.33

Table 5.2: Results for the split percentages from 40 runs of the model

Table 5.2 shows the results for the number of each type of patient passing through the ED for 40 runs of the model ('completed jobs' indicates patients passing through the node). The average percentages of each patient type are summarised in the following table, with the splits having been tailored using the technique mentioned earlier in order to match those given in Section 4.4. The numbers on the left are the actual splits, and those on the right are the figures used in the model.

Total ED Patient Split			
Red	Orange	Yellow	Green
6% | 2.75%	36% | 15.75%	32% | 27%	26% | 54.5%

5.3.4 Doctors in the ante-room

As there was only one doctor in each of the yellow and green rooms, it was simple to simulate them as just one work centre which the patients all passed through. However, in the ante-room, doctors work in shifts which are summarized in the table below:

Anteroom doctor shifts		
Shift	Time	No. doctors
Day	08:00-16:00	6
Evening	16:00-23:00	4
Night	23:00-08:00	2

In order to accommodate this, a resource known as *Doctor* was introduced into the model. All of the *Initial Consult* work centres in the anteroom require a *Doctor* resource before they can operate. The *Doctor* resource works on a shift basis according to the above table.

5.3.5 Green coded patient flow

Green coded patients go straight into the green room queue. Patients waiting in this queue are waiting to be seen by a doctor in the green room for their initial consult which takes between 15 and 40 minutes (see Section 4.5). This consult is therefore modeled on a Uniform distribution with the minimum set at 15 and maximum at 40.

After this initial consult, patients either exit the green room or, if they have had an x-ray taken, they will wait for this result. Patients being discharged are either completely discharged, or they are given a specialist referral, which means they exit the system for the time being, but return at a later stage. Figure 5.2 displays a screenshot from Simul8 depicting the green coded patients' pathway through the system.

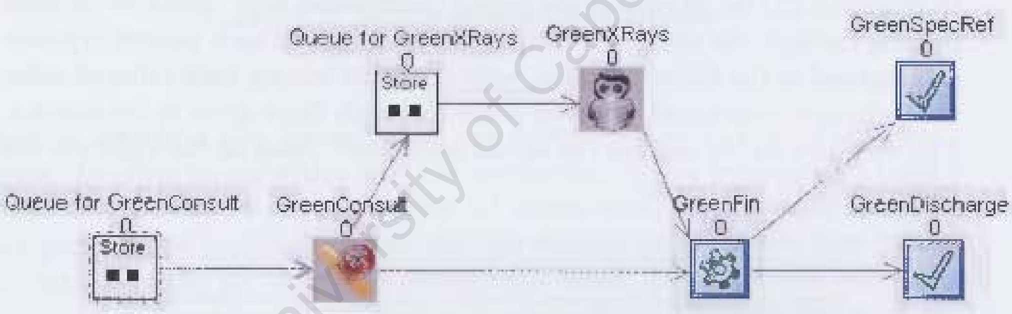

Figure 5.2: Simul8 screenshot of green pathway

The wait for x-ray results is generally about 30 minutes (refer to Section 4.5). Using the logic outlined in Section 5.2.1, the service time for green x-rays was modeled on an Erlang 15 distribution with a parameter value of 25. Figure 5.3 displays a screenshot from Simul8 with the results from the key performance indicators for 40 runs of the simulation:

Table 5.3 shows that the average wait for x-ray results is just less than 5 minutes. If the parameter value of 25 is added to this time, we obtain an expected waiting time of just less than 30 minutes. The RRC for the wait for x-ray results is 0.05 which indicates that these results are very consistent.

		Low 95% Range	Average Result	High 95% Range
Queue for GreenConsult	Average queue size	3.96	4.04	4.13
	Minimum (non-zero) Queuing Time	11.29	15.95	20.60
	Average Queuing Time	184.05	186.75	189.46
	Maximum Queuing Time	419.21	437.99	456.77
Queue for GreenXRays	Average queue size	0.04	0.04	0.04
	Minimum (non-zero) Queuing Time	0.13	0.18	0.23
	Average Queuing Time	4.36	4.56	4.76
	Maximum Queuing Time	49.95	53.82	57.69

Table 5.3: Simul8 results for green patients

The average wait for a consult is just more than 3 hours, and the RRC for this result is 0.01 which again indicates that this result is very consistent. The average queue length for green coded patients is just over 4 people. The average number observed in the green queue in reality was about 6 people (refer to Section 5.7). This discrepancy can be explained by the shift system used in the model. As explained in Section 5.3.3, the model actually runs the green room for longer than it operates in reality in order to attend to any stray patients. This means that there will be periods where there are very few patients in the line, which would account for the slight underestimation.

5.3.6 Yellow coded patient flow

Yellow coded patients go straight into the yellow room queue after triaging. Patients waiting in this queue are waiting to be seen by a doctor in the yellow room for their initial consult which takes between 15 and 40 minutes (see Section 4.5). This consult is therefore modeled on a Uniform distribution with the minimum set at 15 and maximum at 40.

After this initial consult, patients either exit the yellow area, wait for x-ray or blood-test results, or wait for a specialist referral. Less severe patients waiting to see a specialist will leave the system and come back at a later stage, while the more severe patients will wait in the holding area or ante-room. Similarly for patients waiting for test results; less severe patients will wait in the yellow room, and more severe patients will wait in the holding area or ante-room. Figure 5.3 depicts the yellow coded patients' pathway through the system.

The wait for x-ray results is generally about 30 minutes, and the wait for bloodtest results is generally about 290 minutes (refer to Section 4.5). Using the logic outlined in Section 5.2.1, the service time for yellow x-rays was modeled on an Erlang 15 distribution with a parameter value of 30,

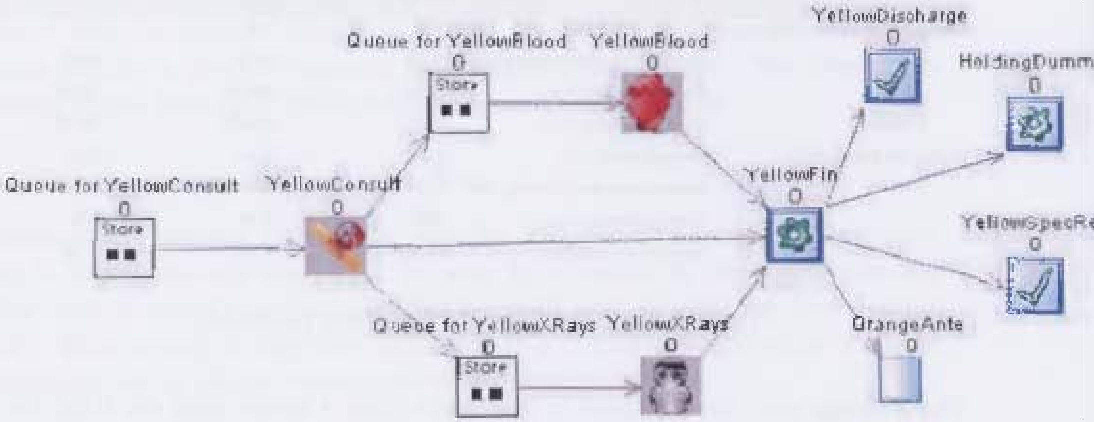

Figure 5.3: Simul8 screenshot of yellow pathway

while the service time for yellow bloodtests was modelled on an Erlang 15 distribution with a parameter value of 86. Table 5.4 displays a screenshot from Simul8 with the results from the key performance indicators for 40 runs of the simulation:

		Low 95% Range	Average Result	High 95% Range
Queue for YellowConsult	Average queue size	5.39	5.67	5.94
	Minimum (non-zero) Queuing Time	1.97	2.76	3.55
	Average Queuing Time	204.35	213.39	222.43
	Maximum Queuing Time	819.42	856.66	893.90
Queue for YellowXRays	Average queue size	0.00	0.00	0.00
	Minimum (non-zero) Queuing Time	0.65	0.31	1.17
	Average Queuing Time	0.89	0.93	1.08
	Maximum Queuing Time	26.50	29.35	32.21
Queue for YellowBlood	Average queue size	1.67	1.87	2.06
	Minimum (non-zero) Queuing Time	1.27	1.70	[illegible]
	Average Queuing Time	180.43	199.47	218.51
	Maximum Queuing Time	802.85	878.54	954.24

Table 5.4: Simul8 results for yellow patients

The average queuing time for bloodtest results is about 199 minutes which, when added to the parameter value of 86, results in a total bloodtest turnaround time of about 285 minutes. The RRC for this result is 0.08 which indicates that the result is consistent. Similarly, if the parameter value for x-rays is added to the average waiting time, it gives a total wait of just over 30 minutes with an RRC of 0.09. The average wait for the initial consult is about 213 minutes, with an RRC of 0.04. In total, there are just more than 7 people queuing on average in the yellow room in the model which is similar to the reality.

5.3.7 Ante-room patient flow

After triage, ante-room patients are sent straight into the ante-room, where they wait on a trolley if there is one available. Ante-room patients are a combination of red and orange coded patients, and the more seriously ill patients are seen first. The table in Section 4.5 shows that the initial consult is between 30 and 180 minutes for red coded patients and between 30 and 150 minutes for orange coded patients. These are both modeled as Uniform distributions in the model.

After the initial consult, patients either have: no more testing, x-rays, bloodtests, or x-rays and bloodtests. After this, patients will then exit the ante-room; either into the holding area, or to be discharged. Figure 5.4 depicts the ante-room patients' flow through the system. This looks more complicated than for the other patients because we have included multiple service points for both orange and red coded patients. This is not how the ante-room operates in reality, but this allows us to simulate the fact that there are multiple doctors in the ante-room seeing patients. The *Doctor* resource can be seen on the left of the diagram between the *OrangeAnte* and *RedAnte* work centres. This resource will float between the *OrangeConsult* and *RedConsult* work centres whenever a patient arrives at one of them.

As indicated previously in this chapter, the wait for x-ray results is about 30 minutes, and the wait for bloodtest results is about 290 minutes. Logically, the wait for both x-ray and bloodtest results should also be about 290 minutes. All of these service times are modeled as Erlang 15 distributions with parameter values of 27, 83.75 and 51 for x-rays, bloodtests, and x-rays and bloodtests respectively. We would expect the parameter values for bloodtests, and x-rays and bloodtests to be very similar, as we are essentially trying to simulate the same length of time passing. However, 25% of anteroom patients only get their blood tested, whereas 45% of anteroom patients have both x-rays and blood taken. Therefore, the parameter value for bloodtests only is necessarily higher than that for both x-rays and bloodtests because there are less patients passing through the "bloodtest only" work centre. Table 5.5 displays a screenshot from Simul8 with the results from the key performance indicators for 40 runs of the simulation.

When the parameter value for x-rays is added to the waiting time from Table 5.5, we obtain a total x-ray turnaround time of just over 30 minutes with an RRC of 0.05. Using similar calculations, we can obtain a bloodtest

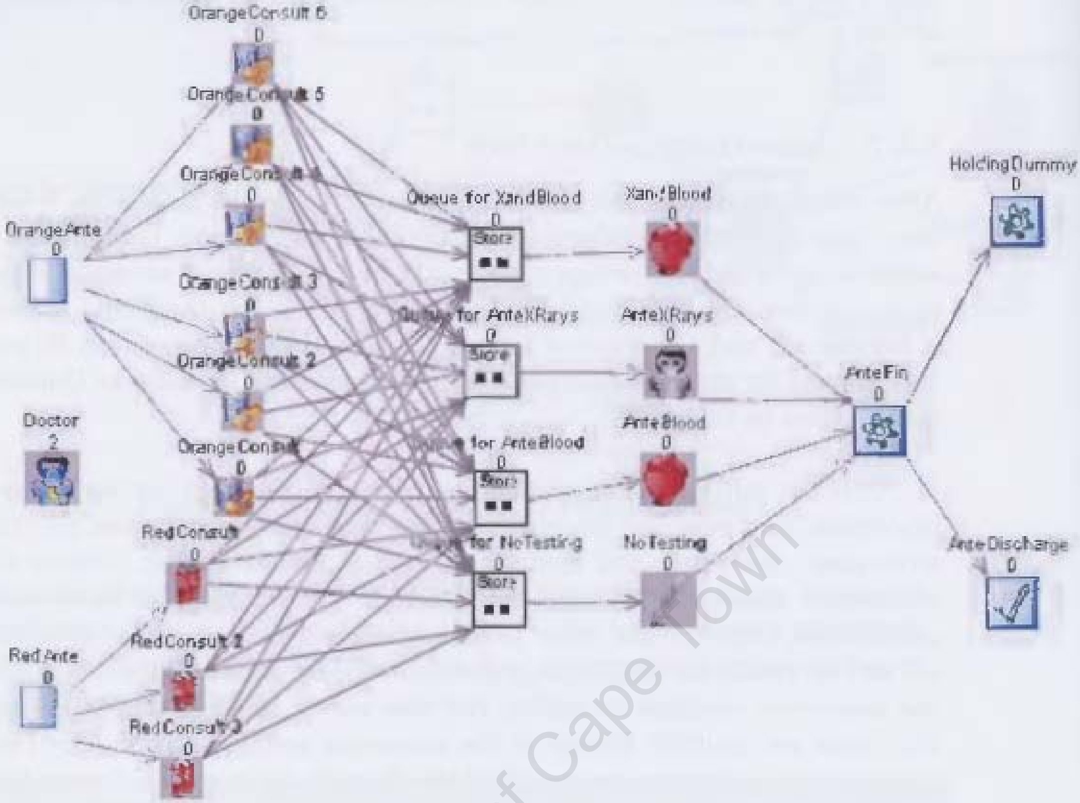

Figure 5.4: Simul8 screenshot of ante-room patients' flow through system

turnaround time of just less than 290 minutes with an RRC of 0.09, and an x-ray and bloodtest turnaround time of just over 290 minutes with an RRC of 0.12. The RRC values for all of these results indicate that they are all consistent, and should represent reality adequately in the model.

Looking at the queuing time for red coded patients, we can see that the average is 15.37 with an RRC of 0.03. This is very low in comparison to other waits in the system, and is a fair indication of reality. However, red coded patients can only be prioritized in the simulation up to a point. Any red coded patient arriving is given priority to see the next available doctor. However, in the simulation, doctors will not leave other patients in order to attend to a highly critical patient. In reality, this is not always the case: doctors in the ante-room will sometimes leave less critical patients in order to stabilise a highly critical arrival. This was impossible to build into the model, and so the average waiting time here is probably slightly higher in the model than in reality. This should not really affect the queuing time

		Low 95% Range	Average Result	High 95% Range
OrangeAnte	Average queue size	4.02	4.13	4.23
	Minimum (non zero) Queuing Time	0.10	0.14	0.15
	Average Queuing Time	115.68	118.23	120.78
	Maximum Queuing Time	524.10	536.41	548.73
RedAnte	Average queue size	0.07	0.07	0.08
	Minimum (non-zero) Queuing Time	0.12	0.16	0.20
	Average Queuing Time	14.84	15.37	15.90
	Maximum Queuing Time	96.19	102.39	108.60
Queue for AnteXRays	Average queue size	0.03	0.03	0.03
	Minimum (non zero) Queuing Time	0.29	0.43	0.56
	Average Queuing Time	3.38	3.54	3.71
	Maximum Queuing Time	54.56	57.79	61.01
Queue for AnteBlood	Average queue size	1.86	2.06	2.25
	Minimum (non zero) Queuing Time	0.76	1.11	1.45
	Average Queuing Time	187.67	205.59	223.51
	Maximum Queuing Time	842.33	918.03	993.73
Queue for XandBlood	Average queue size	3.71	4.29	4.86
	Minimum (non-zero) Queuing Time	0.42	0.59	0.77
	Average Queuing Time	208.85	239.12	269.38
	Maximum Queuing Time	778.05	847.92	917.79

Table 5.5: Simul8 results for ante-room patients

for orange coded patients in the model which is about 118 minutes with an RRC of 0.02.

If we add up all of the average queue sizes in the ante-room, the total comes to about 10 patients. In reality, there are normally about 16 patients in the ante-room. However, this can be explained by the fact that some patients who actually should be in the holding area are instead waiting in the ante-room due to the fact that there is not enough space in the holding area. Therefore 10 patients is about the average number of actual ante-room patients in the ante-room.

5.3.8 Holding area

Patients enter the holding area from the yellow room or from the ante-room. There are 2 types of patients in the holding area: those waiting for a bed in the main hospital, and those undergoing *trial of therapy*. Those undergoing *trial of therapy* are waiting for results or are simply under observation, and they will generally be discharged from the holding area if they are deemed healthy. Those waiting for a bed in the main hospital generally need to see a specialist first before they can be admitted. Therefore, from the holding area, patients are either discharged, or they wait for a specialist, after which they will get sent to a ward in the main hospital, and hence leave our system.

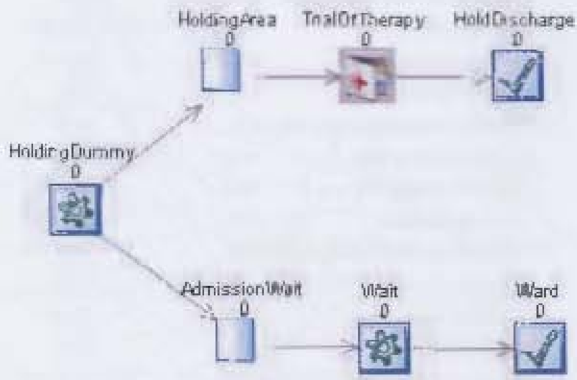

Figure 5.5: Simul8 screenshot of the holding area of the system

The average wait for a patient who is eventually discharged from the holding area is about 13 hours. The average wait for patients who are admitted into the main hospital from the holding area (the specialist referrals) is normally about 2 hours less than this. This is because the patients who are eventually discharged will receive all of their treatment in the holding area, while the other patients receive their treatment in the main hospital.

Table 5.6 displays a screenshot from Simul8 with the results from the key performance indicators for 40 runs of the simulation. The service times for both *TrialOfTherapy* and *Wait* (refer to Figure 5.5) are modeled on an Erlang 15 with the parameters 61.25 and 60.25 respectively.

		Low 95% Range	Average Result	High 95% Range
AdmissionWait	Average queue size	7.31	9.73	12.15
	Minimum (non-zero) Queuing Time	7.40	40.95	74.50
	Average Queuing Time	452.51	594.10	735.70
	Maximum Queuing Time	1233.22	1456.49	1679.76
HoldingArea	Average queue size	9.47	11.57	13.72
	Minimum (non-zero) Queuing Time	-3.88	20.66	45.20
	Average Queuing Time	592.34	719.52	846.70
	Maximum Queuing Time	1530.26	1740.17	1950.07

Table 5.6: Simul8 results for holding area patients

Adding the parameter value to the time given in Table 5.6 we obtain a total waiting time for holding area patients undergoing *trial of therapy* of just over 13 hours with an RRC of 0.17. Similarly, the wait for the specialist referrals is just less than 11 hours, with an RRC of 0.23. These are clearly the most inconsistent results in the model. However, they are adequate for the purposes of this model.

The average number of people in the holding area according to the model is just over 21 people, while the table in Section 4.1.1 indicates that, on average, there are only about 14 people there. This can be explained by noting that the overflow from the holding area usually lands up in the ante-room, as described in the previous section.

Appendix B provides a summary of all of the parameter values discussed in the previous four sections.

5.3.9 The complete model

After putting together the various components of the simulation model which were described previously, the final model is depicted in Figure 5.6.

In order to test the validity of this simulation model, a few basic tests were done using the parameters previously described in the chapter. Using the treatment books described in Chapter 4, a count of the number of patients in each section of the ED over a month was constructed. These data are shown in Table 5.7.

Table 5.7: Number of real patients in each treatment book over a month in the ED

Total	Green	Yellow	Ante	Holding
3909	951	1198	1760	1192

If Figure 5.6 is examined closely, it can be seen that some yellow coded patients are admitted into the ante-room after being through the yellow room first. These are generally cases where the doctor decides that the patient is actually more seriously ill than was initially thought in triage. Therefore, these patients would be counted in both the yellow treatment book, and the ante-room treatment book, resulting in a double count. Referring to Section 4.4 it can be seen that 16% of all yellow coded patients are admitted into the ante-room. Over a month, this would be about 152 patients. Therefore Table 5.8 presents a revised count of patients.

Then, the simulation model was run in Simul8 for 40 runs over a period of a month, and results for key indicators were obtained. The indicators

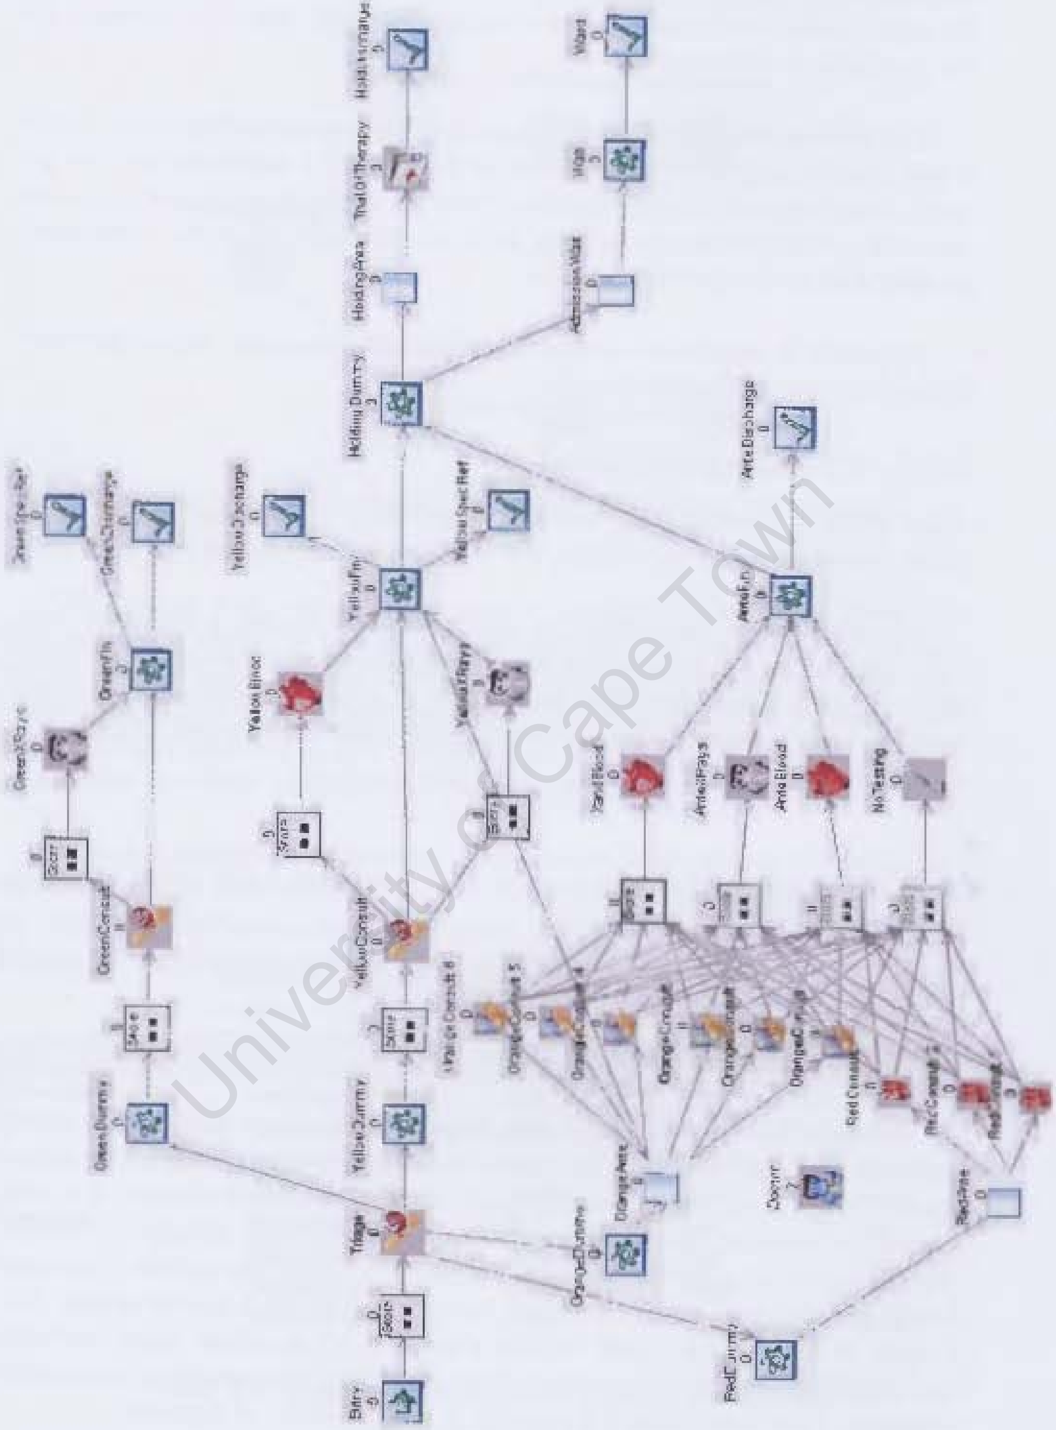

Figure 5.6: Simul8 screenshot of the complete simulation model

Table 5.8: Revised number of real patients in each treatment book over a month in the ED

Total	Green	Yellow	Ante	Holding
3717	951	1198	1568	1192

used were the total number of patients passing through key nodes. Referring to Figure 5.6, the nodes used were: *GreenDummy*, *YellowDummy*, *OrangeDummy*, *AnteDummy*, and *HoldingDummy*. Looking at Figure 5.6, all green coded patients pass through *GreenDummy*, and similarly for the other types of patients with the other nodes. The results are shown in Table 5.9.

		Low 95% Range	Average Result	High 95% Range
GreenDummy	Number Completed Jobs	958.19	965.95	973.71
YellowDummy	Number Completed Jobs	1173.95	1182.33	1190.70
OrangeDummy	Number Completed Jobs	1321.82	1330.83	1339.83
RedDummy	Number Completed Jobs	208.92	213.63	218.33
HoldingDummy	Number Completed Jobs	1414.69	1424.33	1433.96

Table 5.9: Simul8 results for key nodes for 40 runs of the simulation

Comparing Table 5.8 with Table 5.9, it can be seen that the average results given out by the model are very close to reality for the green and yellow coded patients. When the average figures for the orange and red coded patients are added up, we can see that the model is a close approximation to reality in this case as well.

The only figure where the model is substantially different to the reality is the one for the holding area. The actual number of 1192 patients passing through the holding area is substantially less than the average of 1424 given out by the model. In reality, around 32% of all patients in the ED pass through the holding area while, in the model, nearly 39% of patients pass through the holding area. In the model, as in reality, patients only enter the holding area from the yellow room or from the ante-room. In Section 4.4, an assumption was made whereby all patients leaving the ante-room for the holding area were grouped together with all patients leaving the ante-room to wait for a specialist. This assumption was made because most of the patients wait in the holding area for a specialist. However, there are some occasions in reality where patients will wait for a specialist in the ante-room, which would account for the higher percentage of patients passing through the holding area in the model. It would be too complicated for the purposes of this model to alter the patient split percentages to accommodate this minority of patients. As the difference is not too great, we will note this discrepancy, but leave it in the model.

Chapter 6

Sensitivity Testing

6.1 Brief overview

The simulation model described in the earlier chapters is not meant to be a completely accurate depiction of reality. The ED at Groote Schuur is very complicated, and any model cannot hope to capture all of this complexity. Further, due to lack of data, certain simplifications and adjustments were made as discussed in the previous two chapters. However, the model should give us a general overview of how the system operates. With this in mind, it is necessary to run some tests on the model, in order to assess where its strengths and weaknesses lie.

The methods used here fall under the realm of *model validation*. This is defined as "substantiation that a computerized model within its domain of applicability possesses a satisfactory range of accuracy consistent with the intended application of the model" [17]. The intended application of this model is not to be used as an accurate prediction tool, but rather as a helpful aid in order to identify factors influencing patients' waiting times. Therefore, it is not necessary to do extensive sensitivity analyses here. The basic tests performed on the model will include:

- Face Validity Tests
- Animation and Tracing Tests
- Internal Validity Tests
- Parameter Variability-Sensitivity Tests
- Extreme Condition Tests
- Model Structure Tests

Face Validity Tests are when people knowledgeable about the system are asked whether the model and its behaviour are reasonable. *Animation and Tracing Tests* are when the model's behaviour is examined at a slow pace through the graphical interface. More specifically, different types of entities are traced through the model in order to determine if the model's logic is correct. *Internal Validity Tests* are when several replications of a stochastic model are made to determine the amount of internal stochastic variability within the model. *Parameter Variability-Sensitivity Tests* consist of changing the values of certain parameters in the model to determine the effect on the output. *Extreme Condition Tests* test the model's reaction to extreme and unlikely reactions of levels of factors in the system. [17] *Model Structure Tests* are not really mentioned in the literature. However, for two specific parts of this model it is desirable to build different versions and then compare them in order to see which one is a more accurate depiction of reality.

The first two types of test mentioned cannot really be recorded here. For the face validity test, the animation model, along with the outputs of patient waiting times at various points in the system, were shown to the CLO and the Head of the ED at Groote Schuur. Both of them felt that, while the results were not an exact replication of reality, they were a reasonable approximation of the system. For the basic animation and tracing tests, various patients were assigned different colours within the model and these patients were followed visually. Again, the simulation seemed to be a reasonable approximation of reality.

Looking back at Chapter 5, we can see that internal validity tests have already actually been performed while the model was being built. The *Result Relative Consistency* indicator defined in Section 5.2.1 measures the internal stochastic variability within stochastic areas of the model. Therefore, the only tests which will be dealt with in this chapter will be parameter variability-sensitivity tests, extreme condition tests, and model structure tests.

6.2 Parameter variability-sensitivity tests

At the end of Section 4.4, there is a table which shows the percentage of patients getting each specific treatment (x-rays, bloodtests etc) for each of the consulting rooms (green, yellow, or ante). It is mentioned there that these patient splits are only estimates obtained after consultation with some of the ED doctors as actual records of that information do not exist. Therefore, in order to test whether this estimation will have any effect on the outcome of the model, we will use equal splits on all of the treatments, and note any differences in the outcomes. The criteria used to test this will be the average waiting times for each of the treatment paths, and the average overall time spent in the system for each patient type. Table 6.1 shows the original and altered splits used in the test, and the criteria results are given in Table 6.2.

Table 6.1: The original and altered splits

Room	Consult Only	X-Ray	Blood Test	X-Ray & BloodTest	Total
Green(Original)	60%	40%	-	-	100%
Green(Altered)	50%	50%	-	-	100%
Yellow(Original)	50%	15%	35%	-	100%
Yellow(Altered)	33.33%	33.33%	33.33%	-	100%
Ante(Original)	10%	20%	25%	45%	100%
Ante(Altered)	25%	25%	25%	25%	100%

Table 6.2: Criteria results for original and altered splits

Criterion	Original	Altered
Queue for GreenConsult	186.75	186.75
Queue for YellowConsult	213.39	213.39
OrangeAnte	118.23	119.70
RedAnte	15.37	14.89
Queue for GreenXRays	4.56	7.45
Queue for YellowXRays	0.99	2.55
Queue for YellowBlood	199.47	155.60
Queue for AnteXRays	3.54	4.90
Queue for XandBlood	239.12	31.34
Queue for AnteBlood	205.59	208.77
GreenDischarge	225.62	230.32
YellowDischarge	352.79	341.69
AnteDischarge	451.16	342.48

Looking at the results in Table 6.2, we can see that the difference in the patient splits does not really alter most of the waiting times within the model. The only times which are substantially affected are the wait for bloodtest results for yellow patients, and the wait for bloodtest and x-ray results for ante-room patients. In both cases, this wait is reduced when even patient splits are used. This result is understandable for the ante-room patients if the table listing the splits is examined. In the original model, 45% of ante-room patients wait for both bloodtest and x-ray results, and this is reduced to 25% in the alternative model. Obviously, less patients leads to shorter queues which leads to lower waiting times.

However, for the yellow coded patients' waiting for bloodtest results, the split is reduced by less than 2%, and there is still a significant reduction in the waiting time in the alternative model. It stands to reason then that this criterion is sensitive to changes in the patient splits within the model.

The reduction in waiting time for ante-room patients waiting for bloodtest and x-ray results also leads to a reduction in the criterion *AnteDischarge* which effectively measures the average time that an ante-room patient spends in the system. This is also logical, as a reduced waiting time in one part of the system for ante-room patients should lead to a reduction in their time in the system. There is also a slight decrease in the average time spent in the system for yellow coded patients, which again is explained by the reduction in the waiting time for bloodtest results for yellow coded patients. There is also a moderate increase in the average time spent in the system for green coded patients which can be explained by the slight increase in waiting time for green coded patients waiting for x-ray results.

6.3 Model structure tests

As mentioned in Chapter 5, the ante-room in the simulation model is constructed differently to its real-life counterpart. In reality, there is only one ante-room. In the model, there are nine different *consulting areas* in the ante-room, divided into six for orange coded patients and three for red coded patients, with the different patients queuing separately for these areas. The reasons for this were also given earlier. Additionally, the red areas have priority when it comes to using a doctor. i.e. when a patient comes into a red area, that area will get the next available doctor even if there was an orange coded patient waiting before them.

6.3.1 Ante-room 'layout'

In this sensitivity test, we will test whether this is the best number and combination of areas. The criteria used to judge this will be the average waiting time for orange and red coded patients. The results are given in the following table with the red coded patients' waiting times above the orange coded patients' waiting times:

'Areas'		Red		
		2	3	4
O	5	16.16	14.68	14.46
r		128.62	129.38	129.42
a	6	17.13	15.37	15.26
n		117.95	118.23	119.6
g	7	16.9	15.85	15.1
e		117.95	117.95	119.39

This should be able to be examined using a factorial experiment with the number of areas for each type of patient being the factors. However, seeing as there is only one observation per cell, it is difficult to estimate the two factor interaction. Using a test developed by Tukey, we were able to determine that there is not a significant interaction between these two factors [12]. Furthermore, any factorial experiment would only be able to tell us which of the factors was significant, which is not relevant to the problem here which is more subjective in nature. Instead, we are trying to determine the best representation of the ED at Groote Schuur. It is worth remembering here that, while orange coded patients are important, red coded patients are the priority in the ED as they are the most critical. Therefore, the combination which should be used should be the one with the lowest waiting time for red coded patients, provided it does not affect the orange coded patients in too dire a manner.

In practice, all of the combinations with more than five orange areas are adequate for the purposes of our model. However, there is a quite perceptible decrease in waiting time for red coded patients when the number of red areas is increased from two to three at all levels. After that, the gain becomes less apparent and there also seems to be a trade-off, with the waiting time for orange coded patients becoming worse. It also needs to be remembered that there are at most only six doctors in the ante-room at any one time, so having seven areas for orange coded patients would be a case of overkill. Therefore, the best combination would have to be the one with three red areas and six orange areas, which is the original one we had been using.

6.3.2 Simulation methods in ante-room

In this test, we will test whether it is not better to combine the red and orange coded patients into one queue, and have them share common consulting areas or service points. There would be a total of six areas, as that is the maximum number of doctors at any one time in the ante-room. Red and orange coded patients would then each be assigned a different label in the simulation in order to identify them, and the service point would apply a different distribution for the consultation depending on the type of patient. Red coded patients would be given priority and jump to the front of the queue. This is a neater way of doing the simulation, but is more complicated to set up. The results are given in Table 6.3.

Table 6.3: Criteria results for original and increased arrival patterns

Criterion	Original	Shared Ante-room
OrangeAnte	118.23	101.32
RedAnte	15.37	23.93

There is a quite perceptible drop in the time for orange coded patients. However, as discussed previously, the waiting time for red coded patients is our paramount concern, and this seems to increase by more than 8 minutes. This is not acceptable, and so we must conclude that the original method is preferable.

6.4 Extreme condition tests

The most important driver of the whole model is the patient arrivals. This is the critical unknown factor in any ED, and one over which they have absolutely no control. In Chapter 4, it is mentioned that the original data used to compile the arrival rates were based on a CLINICOM survey, and that these underestimate the true arrival rate by some 30%. Therefore, the arrival rates used when building the model were a 30% increase on these rates. In this extreme condition test, we will increase the arrival rates by 60% from the original CLINICOM statistics and see how the model reacts. The results are given in Table 6.4, where it becomes apparent that this increase in the arrival rates has a profound effect on the system as a whole. It is easier to begin with those criteria which are not badly affected by this change.

The queues for a consultation for both green and red coded patients do not seem to be badly affected. Similarly, the queues to wait for x-ray results by themselves for all patients also do not seem badly affected. Consequently, the total time spent in the system for green coded patients, as measured by *GreenDischarge* does not increase by a large amount.

On the other hand, the queues for consultations for both yellow and or-

Table 6.4: Criteria results for original and increased arrival patterns

Criterion	Original	Increased
Queue for GreenConsult	186.75	258.12
Queue for YellowConsult	213.39	3390.37
OrangeAnte	118.23	780.85
RedAnte	15.37	21.12
Queue for GreenXRays	4.56	4.69
Queue for YellowXRays	0.99	0.94
Queue for YellowBlood	199.47	378.27
Queue for AnteXRays	3.54	4.90
Queue for XandBlood	239.12	3609.73
Queue for AnteBlood	205.59	1443.83
HoldingArea	719.52	4682.88
AdmissionWait	594.1	4130.71
GreenDischarge	225.62	298.95
YellowDischarge	352.79	3591.08
AnteDischarge	451.16	3141.26
HoldDischarge	1221.04	7661.19
Ward	1095.50	7125.23

ange coded patients experience a huge increase in waiting times, as does the waiting time for bloodtest results for ante-room patients. Additionally, the waiting time for bloodtest results for yellow coded patients nearly doubles, but this increase is not as severe as the ones previously mentioned.

However, at three areas in the model, the increase in waiting times seem to be even worse than those previously mentioned. These areas are: the wait for bloodtest and x-ray results for ante-room patients; and both sets of waiting times for the holding area patients. These are measured by *Queue for XandBlood*, *HoldingArea*, and *AdmissionWait* respectively. This fact should be taken into account when designing any experiment around this system. As a direct consequence of these increased waiting times, the time spent in the system by yellow room, ante-room, and both types of holding area patients increases substantially.

Chapter 7

Model Experiments

7.1 Overview of experiments

Now that the model has been validated through the sensitivity analysis, it can begin to be used in various experiments. As explained in Chapter 6, this model is not an exact replication of the ED at Groote Schuur, but is rather an approximation of the system. Therefore, this model should not be used to make exact predictions about the system's behaviour. Nevertheless, such experiments still contribute to a broad understanding of the system's behaviour.

Seeing as there are essentially four different areas in the model, it was decided to do separate analyses of each of these areas as there are different driving factors in each of them. The criteria used to judge these analyses will be the total time spent in the system by patients in these areas.

After consultation with some of the ED doctors, and analysis of the model itself, several factors in each area were identified which were thought to be critical to affecting patient waiting time there. A factorial approach to experimentation was then adopted as it is reasonable to assume that there will be interactions amongst some of these factors. Full factorial experiments with three replications were conducted where it was possible, and in one instance a half factorial design was adopted. For the purposes of this research, it was decided that each factor would be tested at two levels. This is illustrative of what can be done in more detailed future studies. The factors chosen and their levels are listed in Table 7.1 (not all of them were used in each experiment).

For the factors which altered the number of doctors in each room, it was decided to use the existing level as the lower level for testing purposes. This is because it is not feasible to have less than one doctor in the yellow and

Table 7.1: The original and altered splits

Factor	Lower Level	Existing Level	Higher Level
No. Patients Arriving	-5%	+/-3691	+5%
No. of Doctors in Green Room	1	1	2
No. of Doctors in Yellow Room	1	1	2
No. of Doctors in Ante-room	2-4-6	2-4-6	3-5-7
Bloodtest Turnaround Time	180mins	290mins	400mins
Holding Area Waiting Time	7hrs	13hrs	19hrs

green rooms, and reducing the ante-room doctors by one per shift affects the model acutely, and is also not really feasible in practice.

It was noted in Chapter 6 that increasing the arrival rate had a dramatic effect on three areas in the model and that this was not an accurate depiction of the reality. These areas were: the two holding areas, and the wait for x-ray and bloodtest results in the ante-room. Therefore, for experiments where the arrival patterns were increased, the values for these areas were adjusted downwards in order to more accurately reflect the actual behaviour of the model. (Details of the adjustments to the parameters are included in Appendix A.)

7.2 Green room patients

For the green area of the model, there were only two factors which were felt to be driving the patient waiting time. These factors were:

- The arrival pattern
- The number of doctors in the green room

Seeing as there were only two factors, it was relatively simple to conduct a full factorial test as there would only need to be four runs of the simulation (multiplied by the three replications of course). Therefore a full 2^2 factorial design was able to be used (details of which can be found in Appendix C). The results of the ANOVA follow:

Variable	MS	F	p	Effect Size
Arrival Pattern(A)	1340.30	86843	0.000000	21.14
No of Doctors(B)	52658.20	3411975	0.000000	132.49
AB Interaction	323.00	20930	0.000000	10.38

All of the effects were significant at a very low level. This result is to be expected from a simulation as there is a lot of information available. Therefore, for the purposes of our test, it is a better exercise to look at the size of the effects. This gives a better sense of the size of the role that each factor plays in determining waiting time for green coded patients. A graph of these factors is presented in Figure 7.1.

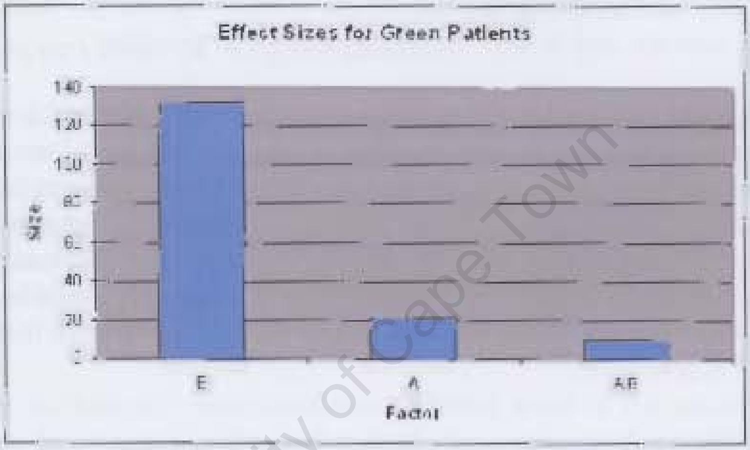

Figure 7.1: Graph of green factors

Looking at Figure 7.1, it becomes apparent that the number of doctors in the green room is clearly the over-riding factor in determining waiting time for green coded patients. In the extreme case, the model suggests a reduction of waiting time by nearly two hours following the introduction of an extra doctor in the green room. However, this should not be taken too literally; suffice it to say that this factor is the primary determinant of waiting time for green coded patients.

7.3 Yellow room patients

For the yellow area of the model, there were three factors which were felt to be driving the patient waiting time. These factors were:

- The arrival pattern
- The number of doctors in the yellow room
- The bloodtest turnaround time

Seeing as there were three factors this time, eight runs of each replication would be needed in order to estimate all of the effects. This would be a total of 24 runs. While this is twice the size of the previous experiment, it is still manageable, and it was felt that it would be better to do this as it would allow all of the interaction effects to be estimated as well as the main effects. Therefore a full 2^3 factorial design was able to be used (details of which can be found in Appendix C). The results of the ANOVA follow:

Variable	MS	F	p	Effect Size
Arrival Pattern(A)	42978	14795.90	0.000000	84.64
No of Doctors(B)	188101	64756.40	0.000000	177.06
Bloodtest Time(C)	37214	12811.40	0.000000	78.76
AB Interaction	11531	3969.60	0.000000	43.84
AC Interaction	3053	1051.00	0.000000	22.56
ABC Interaction	47	16.20	0.000981	2.80

The BC interaction was left out of the ANOVA table as this was the only effect which was not significant. For the same reasons cited earlier, most of the effects here were significant at a low level, and so it was decided to rather look at the effect sizes in order to get a better indication of the drivers of yellow coded patient waiting time. A graph of these effect sizes is displayed in Figure 7.2.

Looking at Figure 7.2, it becomes apparent that all of the main effects are fairly large, while the interaction effects are small enough to be ignored, at least for the purposes of understanding the primary system drivers. However, the addition or reduction of doctors in the yellow room has an effect which is more than twice that of the next largest effect. This indicates that, once again, the biggest driver of patient waiting times is the number of doctors in the area. Interestingly enough, this effect was actually noticed by the ED staff without the use of this model, as there were ongoing experiments with an additional doctor in the yellow room at the time of this research.

7.4 Ante-room patients

For the ante-room area of the model, there were three factors which were felt to be driving the patient waiting time. These factors were:

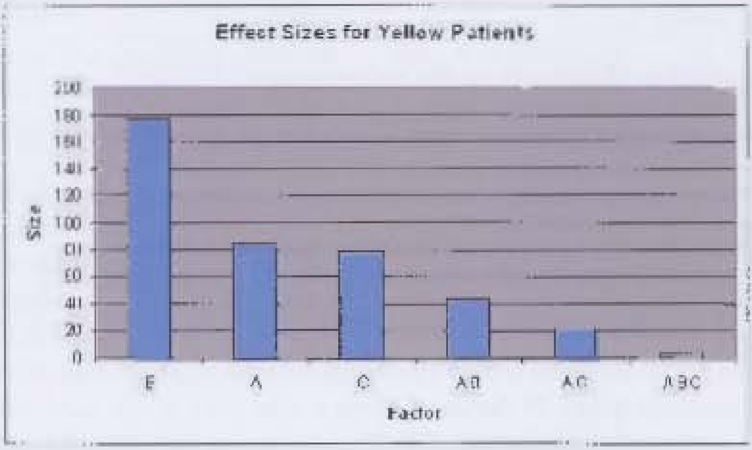

Figure 7.2: Graph of yellow factors

- The arrival pattern
- The number of doctors in the ante-room
- The bloodtest turnaround time

Similarly to the design for the experiments on the yellow area, it was felt that it would be best to use a full 2^3 factorial design here in order to estimate all the main and interaction effects present (details of which can be found in Appendix C). The ANOVA results for this experiment are as follows:

Variable	MS	F	p	Effect Size
Arrival Pattern(A)	339103	51400.80	0.000000	237.73
No of Doctors(B)	9135	1384.60	0.000000	39.02
Bloodtest Time(C)	135138	20484.10	0.000000	150.08
AB Interaction	13884	2104.50	0.000000	48.10
AC Interaction	140861	21351.60	0.000000	153.22
BC Interaction	13042	1977.00	0.000000	46.62
ABC Interaction	5112	774.80	0.000000	29.19

All of the effects here are significant. As explained before, this could be due to this being a simulation model and there being a lot of information available. Once again, it is more illuminating to examine the size of the effects, rather than simply observing the fact that they are significant. A graph of these sizes is depicted in Figure 7.3.

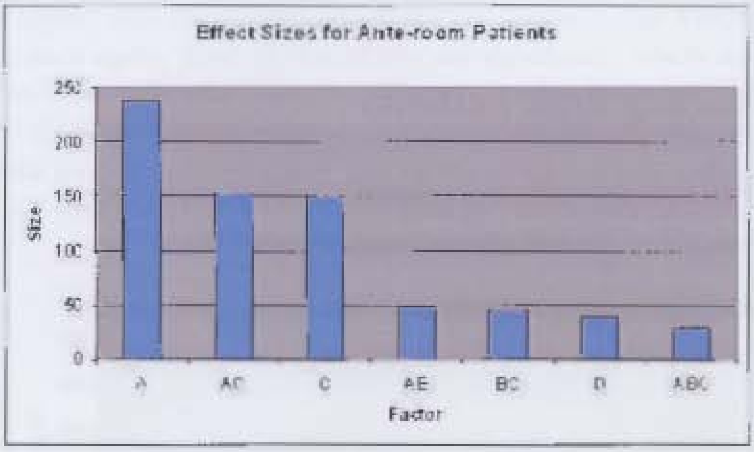

Figure 7.3: Graph of ante-room factors

Looking at Figure 7.3, it becomes clear that two of the factors (arrival pattern and bloodtest turnaround time) along with their interaction (which in this case is of the same order of magnitude as the main effects) drive the waiting time for ante-room patients. Interestingly enough, the only main effect which is not particularly large is that of the number of doctors in the ante-room, which is a completely different result from that of the previous two areas where this change had the biggest effect. A probable reason for this is that, for the majority of the day, there are between 4 and 6 doctors in the ante-room, while there is only 1 doctor in each of the green and yellow rooms. Therefore increasing the number of doctors in either the yellow or the green room by 1 amounts to a 100% increase in manpower, whereas increasing the number of doctors in the ante-room by 1 amounts to only about a 25% increase in manpower.

The main driver here is the arrival pattern into the ED. This does make sense, in that roughly 42% of ED patients go through the ante-room at some stage in their treatment. This is obviously a lot more than for the other rooms, and so any change in the arrival pattern should have the most effect here.

The bloodtest turnaround time also plays a big part in determining waiting time. Additionally, the interaction between these two factors is large. As noted in Chapter 6, increasing the arrival rate serves to further increase the turnaround time for bloodtest and x-ray results for ante-room patients, indicating that the interaction effect here is reinforcing. This was explicitly compensated for in designing this experiment, so this interaction may in fact have been dampened slightly by this, and could be larger than is depicted.

7.5 Holding area patients

For the holding area section of the model, there were five factors which were felt to be driving the patient waiting time. These factors were:

- The arrival pattern
- The number of doctors in the yellow room
- The number of doctors in the ante-room
- The bloodtest turnaround time
- The holding area waiting time

This time, there are five factors which need to be tested. Therefore, in order to measure all the possible main effects and interactions, there would need to be 32 runs of the model multiplied by three replications, which would be a total of 96 runs. This was felt to be a bit unnecessary, as significant interactions higher than 2nd level are not very likely based on the previous experiments. Therefore, it was decided to use a half factorial design as that would mean only doing 48 runs, and not losing too much critical information. The design used was a 2^{5-1} with defining relation I = ABCDE (where A, B, C, D, E are the five factors). This is a resolution V design which means that no main effect or two-factor interaction is aliased with any other main effect or two factor interaction (details of this can be found in Appendix C). The results of the ANOVA are as follows:

Variable	MS	F	p	Effect Size
Arrival Pattern(A)	3862905	5460.29	0.000000	567.37
No of Yellow Room Doctors(B)	28134	39.77	0.000000	48.42
No of Ante-room Doctors(C)	73255	103.55	0.000000	78.13
Bloodtest Time(D)	265117	374.75	0.000000	148.63
Holding Area Wait(E)	3031090	4284.50	0.000000	502.58
AD Interaction	58984	83.37	0.000000	70.11
AE Interaction	556968	787.29	0.000000	215.44
BC Interaction	5951	8.41	0.006690	22.27
DE Interaction	7319	10.35	0.002966	24.70

Once again, only the significant effects are shown in the ANOVA table. However, once again, most of the effects are significant, which means that we need to look at the effect sizes in order to get a better feel for the results. These are given in Figure 7.4.

Figure 7.4: Graph of holding area factors

The effects of two factors stand out here: the arrival pattern, and the holding area waiting time. Their interaction is the next biggest effect which indicates their influence here. This result makes sense, as obviously the holding area wait is going to have a big effect on the time in the system for holding area patients. Additionally, about 32% of all ED patients pass through the holding area, so the arrival pattern would have a big influence here as for the ante-room patients. It was already noted in Chapter 6 that increasing the arrival pattern had a marked effect on the holding area wait. Indeed, this was even compensated for when designing this test. Therefore, *as with the earlier test done on the ante-room patients, this compensation probably dampened the interaction effect, which was nevertheless still quite large. Therefore, it can safely be said that the patient arrival pattern and* the holding area wait have the biggest effect on the time spent in the system for holding area patients.

7.6 Overview of results

Looking at the results presented in this chapter, we can begin to gain some insights into the operation of the ED system. Firstly, it should be noted that the arrival pattern of the patients has the largest effect on patient waiting times in both the ante-room and holding areas, and is a large factor in the patient waiting times in both the yellow and green rooms. Therefore, any effort made which will limit the flow of patients into the ED should have a substantial effect in reducing the waiting time for all patients. This has already been recognised by the management of the ED, and experiments have been conducted with limiting the number of patients allowed into the unit. Obviously, this does create some issues, as no-one wants turn away a sick person. However, if it means that the unit can provide a better service to the existing patients, then it may be a necessary evil. This is definitely an area which warrants further research.

The next issue to emerge from the experiment results, is that increasing the number of doctors in the yellow and green rooms by one has a marked effect on reducing these patients' waiting time. In Section 7.3 it is briefly mentioned that experiments have already been conducted with adding an extra doctor into the yellow room. This doctor is normally transferred from out of the ante-room. Obviously, this is a contentious issue as the ante-room patients require the highest degree of care. However, if the ante-room can still operate effectively with the loss of one doctor, then it should greatly reduce the wait for yellow coded patients. It would be interesting to then look at transferring some of the green coded patients into the yellow room in order to also reduce the green room waiting time. One could also look at calculating the optimal times in which to do this transfer (i.e. finding the optimal compromise between times when the yellow room desperately needs a doctor, and times when the ante-room is able to spare a doctor). These are areas where further research can be performed.

The final insight to emerge from this chapter is that of interaction effects. For both holding area and ante-room patients, interaction between two main effects play a substantial role in the patients' waiting times. For holding area patients it is an interaction between the patient arrival pattern and the holding area wait, while for ante-room patients it is an interaction between the patient arrival pattern and the bloodtest turnaround time. In both cases, the effect is reinforcing, and in both cases the patient arrival pattern is one of the effects involved in the interaction. This means that by controlling patient arrivals, as discussed earlier in this section, we could obtain a more than expected reduction in patient waiting times for both holding area and ante-room patients.

Chapter 8

Conclusion

8.1 Overview of research

The purpose of this research was to investigate efficiency in the ED at Groote Schuur hospital. It is not intended to be used as a definitive solution or to propose explicit recommendations for the running of the ED. It is rather intended to be regarded as a guide which shows the users likely areas where any improvements in efficiency should reduce patient waiting time.

This research is completed as part of a course in Operational Research in a development context. As such, a lot of emphasis is placed on applying traditional OR techniques in non-traditional (developmental) situations. Looking at Chapter 2, we can see that most of the past work done in this sort of field has been in a developed world context. However, this work was undertaken in a context which was very much third world.

Simply put, it was incredibly difficult to obtain any sort of accurate or consistent data. In actual fact, as mentioned in Chapter 1, the first project which was undertaken had to be abandoned due to lack of data. Another major hurdle was that it was very difficult to accurately map a patient path through the system. Because the unit is so over-utilised, the staff are forced to improvise where then is no space available. Indeed, the whole holding area within the unit arose because, when the main hospital was full, there was nowhere to put patients who were finished in the ED and who needed to be admitted into the main hospital. This sort of problem would never happen in a first world context, but is very common in a hospital such as Groote Schuur. Because of the difficulty of obtaining data, and the inconsistencies within the system, it was very difficult to build a completely accurate model. The model which was obtained in this research was sufficient for our purposes, but there is a lot of room for improvement. This will be looked at more closely in Section 8.4.

8.2 Review of methods

For the purposes of structuring the problem, we decided to use Soft Systems Methodology (SSM). This was the author's first experience in applying this method to a real-world problem. We did not attempt a full SSM intervention as this was not necessary for the purposes of this study. Instead, SSM was used to identify key stakeholders and to help choose an appropriate method for the actual problem-solving phase. In this regard, it was a very effective technique. The only issue which became a slight problem was that, since all of the stakeholders had a similar background and similar goals, their root definitions were in turn rather similar. However, seeing as SSM was only being used for a small part of the problem, this did not turn into an issue.

For the actual problem solving stage of the problem, it was decided to use Discrete Event Simulation (DES), and in particular the Simul8 DES package. Once again, it was the author's first experience at using the Simul8 package for a real-world problem. There are many specialist hospital DES packages, but Simul8 proved to be very adaptive and effective at modeling the ED for the purposes of this research.

DES is surely the only OR technique which could be effectively applied to this sort of problem. There are so many uncertainties and variables at work here, that to try and quantify all of them in another way would be an impossible task. Through DES, a working model of the ED was able to be built, and once that was available, it became a fairly simple task to identify areas which could be improved.

8.3 Review of results

Chapter 7 details all of the crucial areas within the ED where improvements need to be made to reduce patient waiting time, and the details do not need to be repeated here. However, at this point it is appropriate to examine the achievements of this research.

When we began this research, this was a largely undefined problem along the lines of "something not being right in the ED". Through hours spent on the floor in the unit, and extensive interviews with various staff members, we were able to gain an insightful understanding of the system as it currently operated. Using this knowledge, we were then able to translate this into a working simulation model which closely mimicked the system of the ED. We then used this model to perform statistically sound experiments on the system which enabled us to identify critical areas which needed attention. We could then combine this knowledge with the insight already gained through examining the system to look at ways in which it could be improved.

To put it simply, this research began with a largely unformulated problem. We structured that problem, finding out exactly what it was, and then found solutions for it in a statistically sound manner.

8.4 Possibilities for future research

This research was done at a fairly rudimentary level, and there is a lot of scope available for future research to be undertaken here. An obvious way in which to further the research is to improve the current simulation model. As mentioned throughout this paper, the model which was built is not an exact replication of the ED, and while it mimics the system to some degree, there is a lot of room for more accuracy to be built in.

Section 8.1 discusses the difficulty of obtaining accurate and consistent data for the ED. An obvious way to get around this is to spend a substantial period of time in the unit observing and recording data. In this way, accurate distributions could be built for the x-ray and bloodtest turnaround times, as well as for the doctors' service times. The anomalies in patient flow mentioned in Section 8.1 could also be more accurately represented in the model. As it stands, the behaviour of the doctors in the ante-room (as discussed in Section 5.3.7) is not a strict representation of reality, so this is another area which could be modified. Of course, the current model also ignores the fact that patient flow alters depending on the day of the week and the month of the year. If this were able to be built into the model, it would increase its power exponentially, and enable it to be used as a powerful predictive tool.

Section 7.6 mentions briefly that introducing an extra doctor into the yellow room, and altering the yellow and green coded patient split could result in reduced waiting times for both types of patients. Future research could possibly look at tweaking these sort of values in a revised model, and seeing what the results are. Another issue discussed in Section 7.6 was that of limiting patient flow into the ED. It should be possible to use the model to establish an optimal number of patients within the unit under certain (changeable) conditions. The model could also be used to calculate the optimal times to transfer staff between sections of the ED. Alternatively, the model could be modified and applied to a different part of the hospital, or even to a different hospital itself for comparative reasons.

In summary, the work done here is adequate for the purposes of this dissertation. However, it is not intended to have the final say in this matter, but is rather intended to be a base upon which others can build and improve.

Bibliography

[1] Kirtland A, Lockwood J, Poisker K, Stamp L, and Wolfe P. Simulating an emergency department "is as much fun as...". *Proceedings of the 1995 Winter Simulation Conference*, pages 1039–1042, 1995.

[2] Wijewickrama AKA and Takakuwa S. Simulation analysis of an outpatient department of internal medicine in a university hospital. *Proceedings of the 2006 Winter Simulation Conference*, pages 425–432, 2006.

[3] Lehaney B, Clarke SA, and Paul RJ. A case of an intervention in an outpatients department. *Journal of the Operational Research Society*, 50:877–891, 1999.

[4] P Checkland and J Poulter. *Learning for Action*. John Wiley and Sons, Chichester, 2006.

[5] Medeiros DJ, Swenson E, and DeFlitch C. Improving patient flow in a hospital emergency department. *Proceedings of the 2008 Winter Simulation Conference*, pages 1526–1531, 2008.

[6] Jun JB, Jacobson SH, and Swisher JR. Application of discrete-event simulation in health care clinics: A survey. *Journal of the Operational Research Society*, 50:109–123, 1999.

[7] Holm LB and Dahl FA. Simulating the effect of physician triage in the emrgency department of akershus university hospital. *Proceedings of the 2009 Winter Simulation Conference*, pages 453–458, 2009.

[8] Ahmed MA and Alkhamis TM. Simulation optimization for an emergency department healthcare unit in kuwait. *European Journal of Operational Research*, pages 936–942, 2009.

[9] Centeno MA, Giachetti R, Linn R, and Ismail AM. A simulation-ilp based tool for scheduling emergency room staff. *Proceedings of the 2003 Winter Simulation Conference*, pages 1930–1938, 2003.

[10] Garcia ML, Rivera C, Centeno MA, and DeCario N. Reducing time in an emergency room via a fast-track. *Proceedings of the 1995 Winter Simulation Conference*, pages 1048–1053, 1995.

[11] Gunal MM and Pidd M. Understanding accident and emergency department performance using simulation. *Proceedings of the 2006 Winter Simulation Conference*, pages 446–452, 2006.

[12] DC Montgomery. *Design and Analysis of Experiments.* John Wiley and Sons, Arizona, 2001.

[13] Bailey NTJ. Operational research in medicine. *Operational Research Quarterly*, 3:24–29, 1952.

[14] Harper PR and Gamlin HM. Reduced outpatient waiting times with improved appointment scheduling: a simulation modelling approach. *OR Spectrum*, 25:207–222, 2003.

[15] Ashton R, Hague L, Brandreth M, Worthington D, and Cropper S. A simulation-based study of a nhs walk-in centre. *Journal of the Operational Research Society*, pages 153–161, 2005.

[16] Davies R. 'see and treat' or 'see' and 'treat' in an emergency department. *Proceedings of the 2007 Winter Simulation Conference*, pages 1519–1522, 2007.

[17] Sargent RG. Verification, validation, and accreditation of simulation models. *Proceedings of the 2000 Winter Simulation Conference*, pages 50–58, 2000.

[18] JA Rice. *Mathematical Statistics and Data Analysis, 2nd Edition.* Duxbury Press, Belmont, California, 1995.

[19] J Rosenhead and J Mingers (editors). *Rational Analysis for a Problematic World Revisited, 2nd Edition.* John Wiley and Sons, Chichester, 2001.

[20] Samaha S, Armel WS, and Starks DW. The use of simulation to reduce the length of stay in an emrgency department. *Proceedings of the 2003 Winter Simulation Conference*, pages 1907–1911, 2008.

[21] Ruohonen T, Neittaanmaki P, and Teittinen J. Simulation model for improving the operation of the emergency department of special health care. *Proceedings of the 2006 Winter Simulation Conference*, pages 453–458, 2006.

Appendix A

Arrivals

In Section 4.3, it was said that for the patient arrivals, the null distribution of the test statistic could be shown to be approximated by the χ^2 distribution with degrees of freedom (df) being given by:

df =number of cells - number of independent parameters fitted - 1

The number of cells referred to in the above equation are the number of occurrences of each number of patients which arrived in the given hour. Looking at Figure A.1, this would be the 'Observed Frequency' column in the table at the bottom which shows a total of 8 cells. However, in order for the above approximation to be adequate, each cell should have at least 5 expected counts [18]. Therefore, in this figure, the first 2 and last 5 rows should be combined to give 3 cells in total, resulting in 1 degree of freedom. The rest of this section shows the output for all 24 values of λ, along with their cumulative distribution graphs and their degrees of freedom.

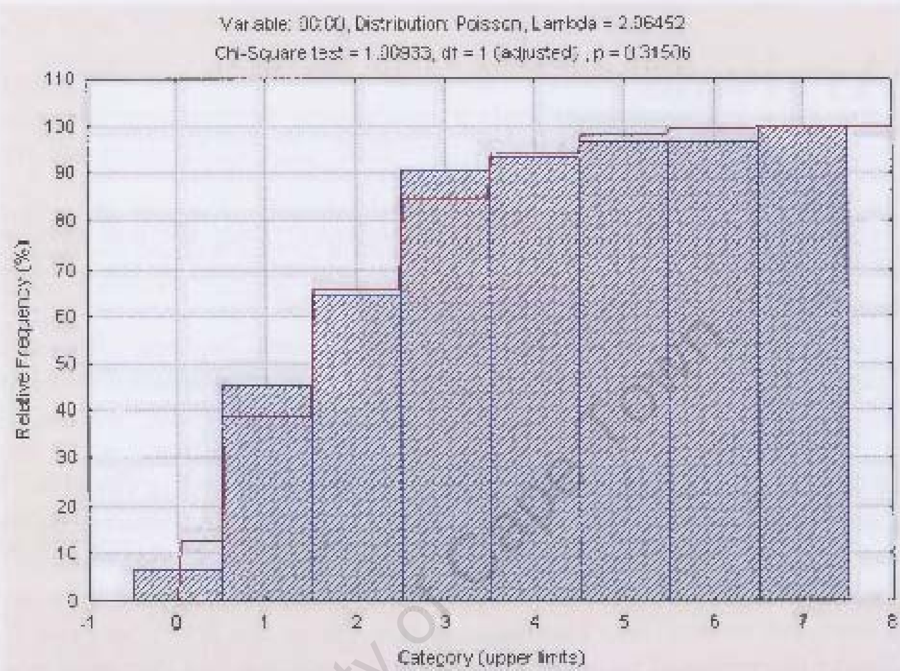

Category	Variable: 00:00, Distribution: Poisson, Lambda = 2.06 (05Arrivals) Chi-Square = 1.00933, df = 1 (adjusted) , p = 0.31506								
	Observed Frequency	Cumulative Observed	Percent Observed	Cumul. % Observed	Expected Frequency	Cumulative Expected	Percent Expected	Cumul. % Expected	Observed-Expected
<= 0.00000	2	2	6.45161	6.4516	3.933270	3.93327	12.68797	12.6880	-1.93327
1.00000	12	14	38.70968	45.1613	8.120300	12.05357	26.19452	38.8825	3.87970
2.00000	6	20	19.35484	64.5161	8.382244	20.43581	27.03950	65.9220	-2.38224
3.00000	8	28	25.80645	90.3226	5.768425	26.20424	18.60782	84.5298	2.23158
4.00000	1	29	3.22581	93.5484	2.977252	29.18149	9.60404	94.1338	-1.97725
5.00000	1	30	3.22581	96.7742	1.229317	30.41081	3.96554	98.0994	-0.22932
6.00000	0	30	0.00000	96.7742	0.422991	30.83380	1.36449	99.4639	-0.42299
< Infinity	1	31	3.22581	100.0000	0.166202	31.00000	0.53613	100.0000	0.83380

Figure A.1: The results for 00:00

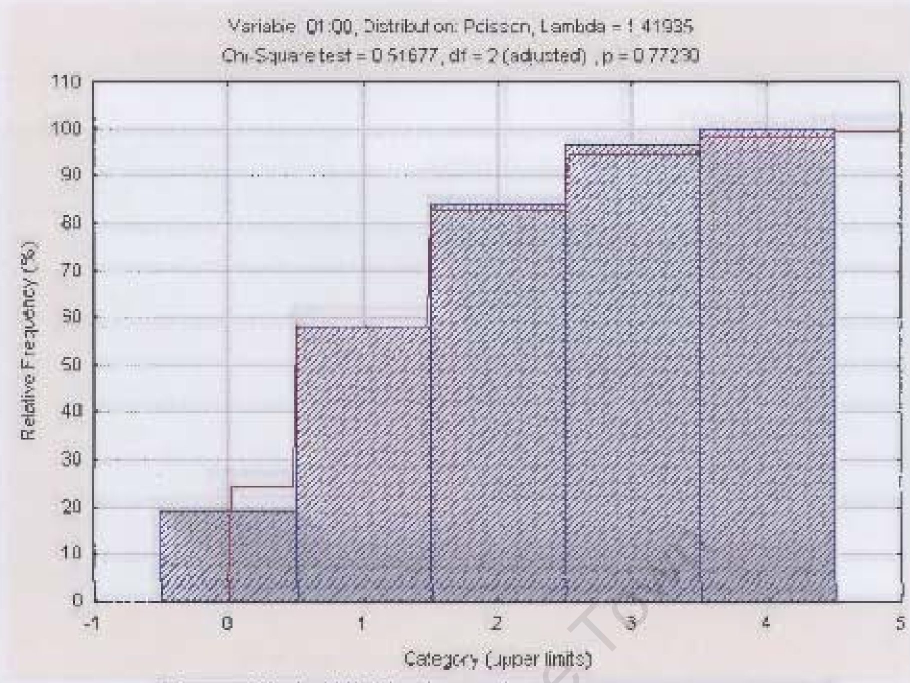

Category	Variable 01.00, Distribution: Poisson, Lambda = 1.41 (09Arrivals) Chi-Square = 0.51577, df = 2 (adjusted) , p = 0.77230 Expected Frequency	Cumulative Expected	Percent Expected	Cumul. % Expected	Observed-Expected
<= 0.00000	7.49797	7.49797	24.18700	24.1870	-1.49797
1.00000	10.64228	18.14025	34.32994	58.5169	1.35772
2.00000	7.55259	25.69284	24.36318	82.8801	0.44741
3.00000	3.57327	29.26610	11.52667	94.4068	0.42673
< Infinity	1.73390	31.00000	5.59321	100.0000	-0.73390

Figure A.2: The results for 01:00

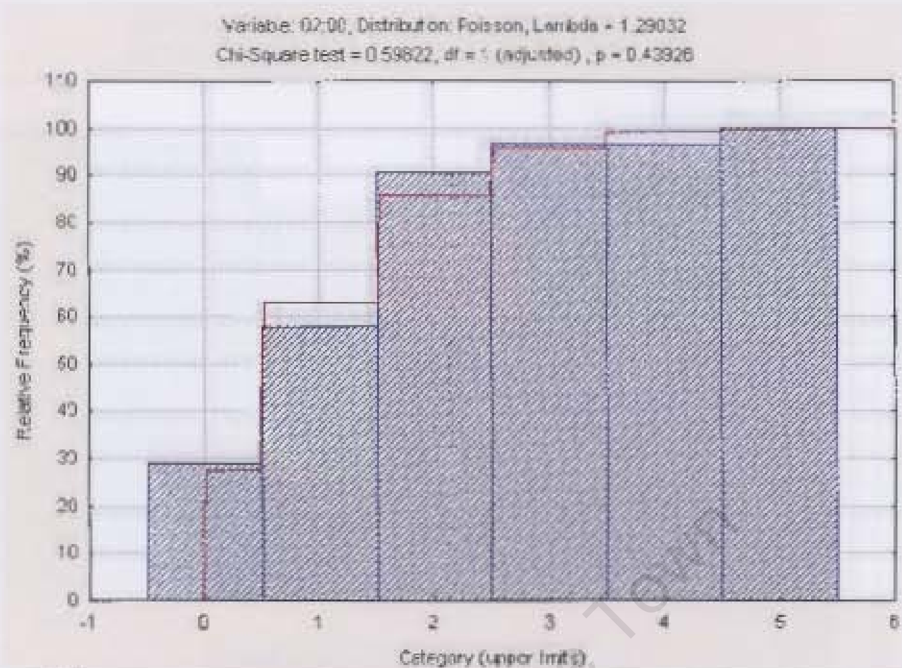

Category	Observed Frequency	Cumulative Observed	Percent Observed	Cumul. % Observed	Expected Frequency	Cumulative Expected	Percent Expected	Cumul. % Expected	Observed-Expected
	Variable: 02:00, Distribution: Poisson, Lambda = 1.29 (00Arrivals) Chi-Square = 0.59822, df = 1 (adjusted) , p = 0.43926								
<= 0.00000	9	9	29.03226	29.0323	8.53064	8.53064	27.51820	27.5182	0.46936
1.00000	9	18	29.03226	58.0645	11.00728	19.53792	35.50735	63.0256	-2.00728
2.00000	10	28	32.25806	90.3226	7.10147	26.63939	22.90797	85.9335	2.89853
3.00000	2	30	6.45161	96.7742	3.05440	29.69379	9.85289	95.7864	-1.05440
4.00000	0	30	0.00000	96.7742	0.98529	30.67908	3.17835	98.9648	-0.98529
< Infinity	1	31	3.22581	100.0000	0.32092	31.00000	1.03523	100.0000	0.67908

Figure A.3: The results for 02:00

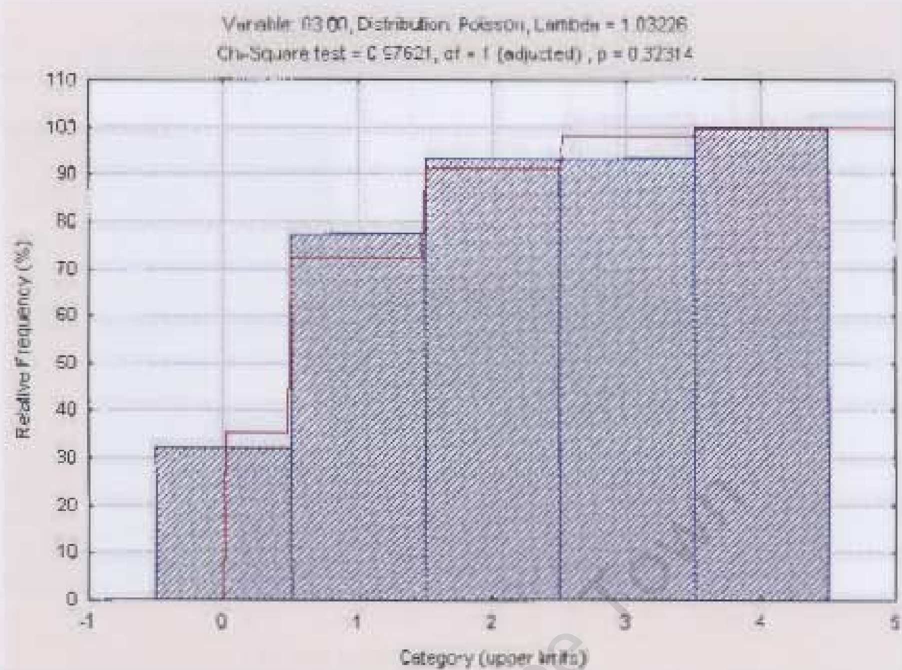

Category	Observed Frequency	Cumulative Observed	Percent Observed	Cumul. % Observed	Expected Frequency	Cumulative Expected	Percent Expected	Cumul. % Expected	Observed-Expected
Variable: 03:00, Distribution: Poisson, Lambda = 1.03 (09Arrivals) Chi-Square = 0.97621, df = 1 (adjusted) , p = 0.32314									
<= 0.00000	10	10	32.25806	32.2581	11.04225	11.04225	35.62017	35.6202	-1.04225
1.00000	14	24	45.16129	77.4194	11.39846	22.44071	36.76921	72.3894	2.60154
2.00000	5	29	16.12903	93.5484	5.88307	28.32378	18.97766	91.3670	-0.88307
3.00000	0	29	0.00000	93.5484	2.02428	30.34807	6.52995	97.8970	-2.02428
< Infinity	2	31	6.45161	100.0000	0.65193	31.00000	2.10301	100.0000	1.34807

Figure A.4: The results for 03:00

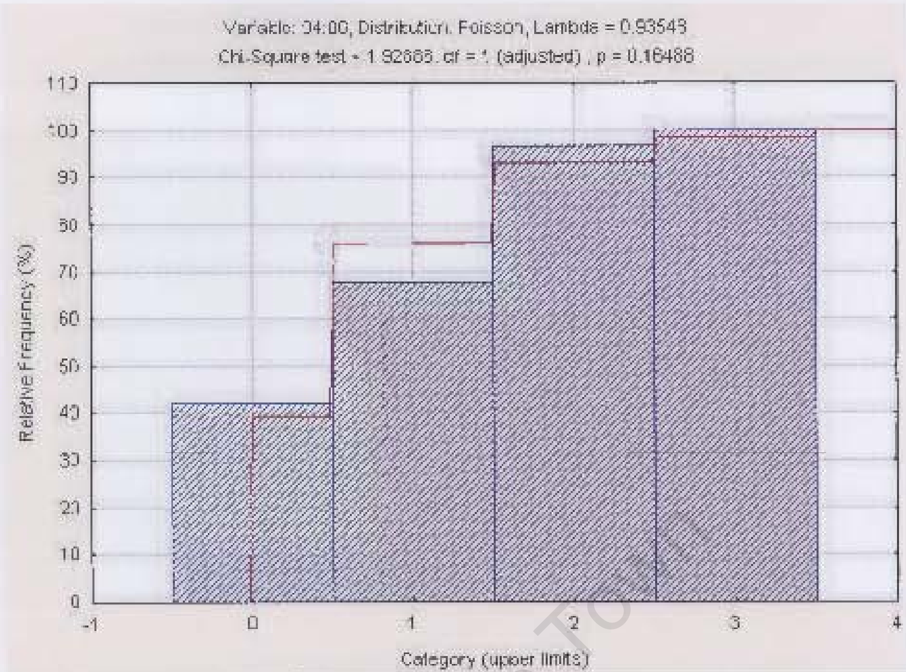

Category	Observed Frequency	Cumulative Observed	Percent Observed	Cumul. % Observed	Expected Frequency	Cumulative Expected	Percent Expected	Cumul. % Expected	Observed-Expected
	Variable: 04:00, Distribution: Poisson, Lambda = 0.93 (05Arrivals) Chi-Square = 1.92888, df = 1 (adjusted) , p = 0.16488								
<= 0.00000	13	13	41.93548	41.9355	12.16428	12.16428	39.23960	39.2396	0.83572
1.00000	8	21	25.80645	67.7419	11.37948	23.54376	36.70800	75.9476	-3.37948
2.00000	9	30	29.03226	96.7742	5.32266	28.86642	17.16987	93.1175	3.67734
< Infinity	1	31	3.22581	100.0000	2.13358	31.00000	6.88252	100.0000	-1.13358

Figure A.5: The results for 04:00

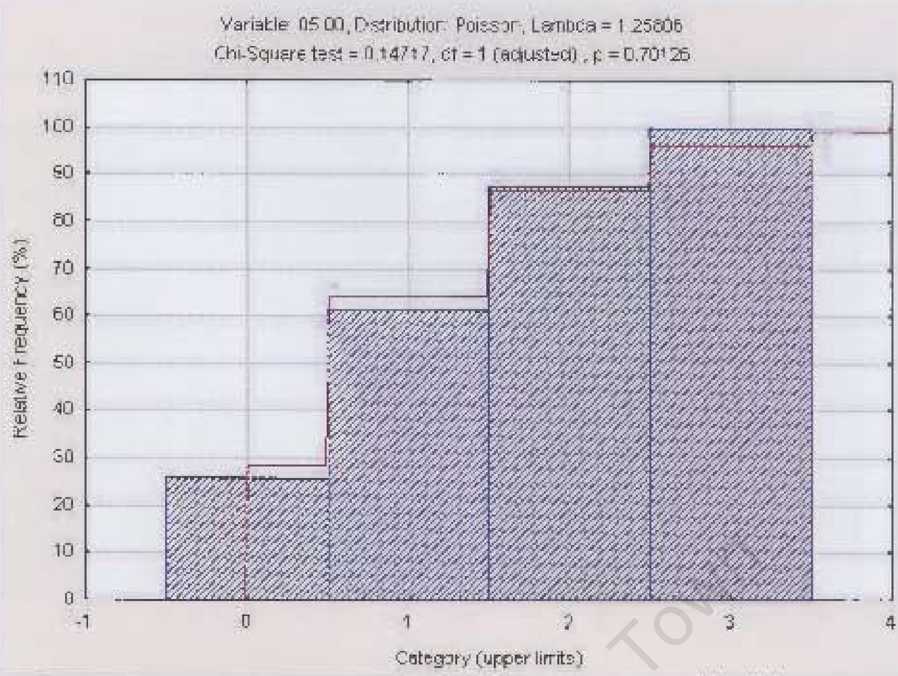

Category	Variable: 05:00, Distribution: Poisson, Lambda = 1.25 (05Arrivals) Chi-Square = 0.14717, df = 1 (adjusted) , p = 0.70126								
	Observed Frequency	Cumulative Observed	Percent Observed	Cumul. % Observed	Expected Frequency	Cumulative Expected	Percent Expected	Cumul. % Expected	Observed Expected
<= 0.00000	8	8	25.80645	25.8065	8.81031	8.81031	28.42036	28.4204	-0.810311
1.00000	11	19	35.48387	61.2903	11.08394	19.89425	35.75464	64.1750	-0.083939
2.00000	8	27	25.80645	87.0968	6.97216	26.86641	22.49082	86.6658	1.027845
< Infinity	4	31	12.90323	100.0000	4.13359	31.00000	13.33416	100.0000	-0.133595

Figure A.6: The results for 05:00

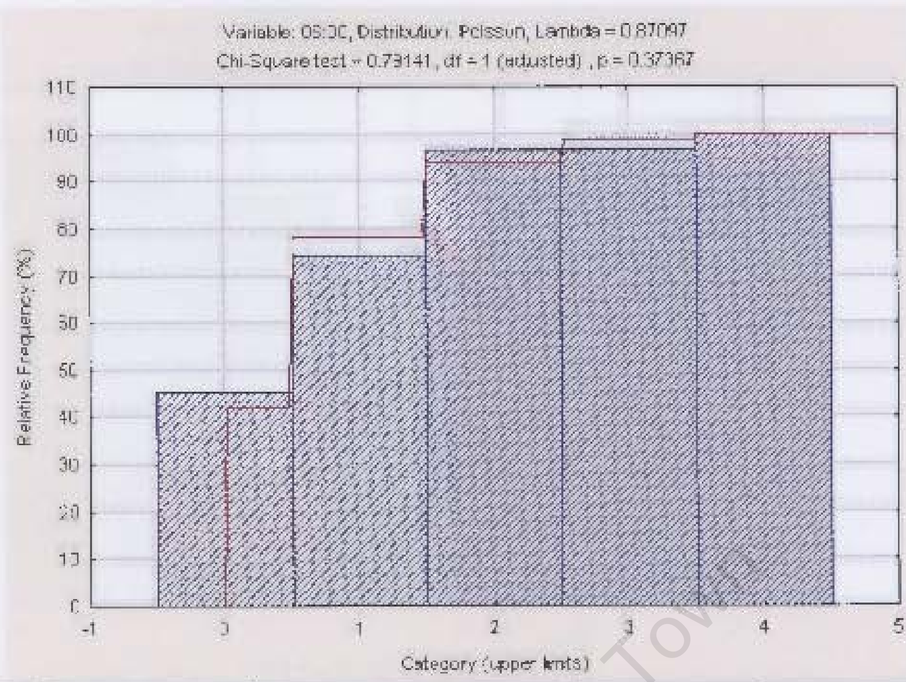

Category	Observed Frequency	Cumulative Observed	Percent Observed	Cumul. % Observed	Expected Frequency	Cumulative Expected	Percent Expected	Cumul. % Expected	Observed-Expected
	Variable: 06:00, Distribution: Poisson, Lambda = 0.87 (09Arrivals) Chi-Square = 0.79141, df = 1 (adjusted), p = 0.37367								
<= 0.00000	14	14	45.16129	45.1613	12.97484	12.97484	41.85463	41.8546	1.02506
1.00000	9	23	29.03226	74.1935	11.30075	24.27569	36.45403	78.3087	-2.30075
2.00000	7	30	22.58065	96.7742	4.92129	29.19698	15.87514	94.1838	2.07871
3.00000	0	30	0.00000	96.7742	1.42876	30.62574	4.60891	98.7927	-1.42876
< Infinity	1	31	3.22581	100.0000	0.37426	31.00000	1.20728	100.0000	0.62574

Figure A.7: The results for 06:00

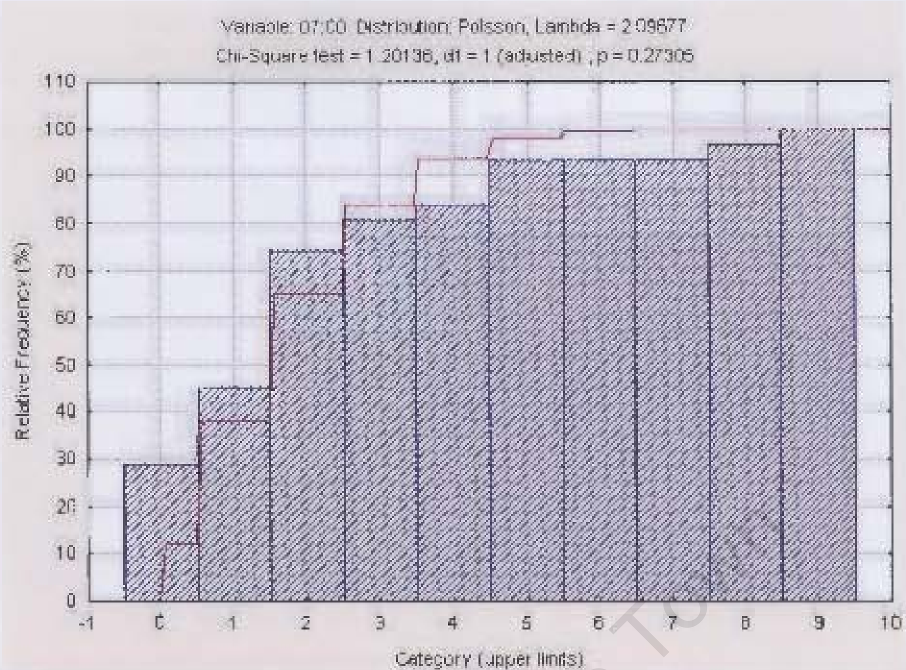

Category	Variable: 07.00, Distribution: Poisson, Lambda = 2.09 (09Arrivals) Chi-Square = 1.20136 df = 1 (adjusted) , p = 0.27305 Observed Frequency	Cumulative Observed	Percent Observed	Cumul. % Observed	Expected Frequency	Cumulative Expected	Percent Expected	Cumul. % Expected	Observed-Expected
<= 0.00000	9	9	29.03226	29.0323	3.808415	3.80841	12.28521	12.2852	5.19159
1.00000	5	14	16.12903	45.1613	7.985387	11.79380	25.75931	38.0445	-2.98539
2.00000	9	23	29.03226	74.1935	8.371775	20.16558	27.00573	65.0502	0.62822
3.00000	2	25	6.45161	80.6452	5.851239	26.01682	18.87497	83.9252	-3.85124
4.00000	1	26	3.22581	83.8710	3.067183	29.08400	9.89414	93.8194	-2.06718
5.00000	3	29	9.67742	93.5484	1.286238	30.37024	4.14915	97.9685	1.71376
6.00000	0	29	0.00000	93.5484	0.449492	30.81973	1.44997	99.4185	-0.44949
7.00000	0	29	0.00000	93.5484	0.134640	30.95437	0.43432	99.8528	-0.13464
8.00000	1	30	3.22581	96.7742	0.035289	30.98966	0.11383	99.9666	0.96471
< Infinity	1	31	3.22581	100.0000	0.010342	31.00000	0.03336	100.0000	0.98566

Figure A.8: The results for 07:00

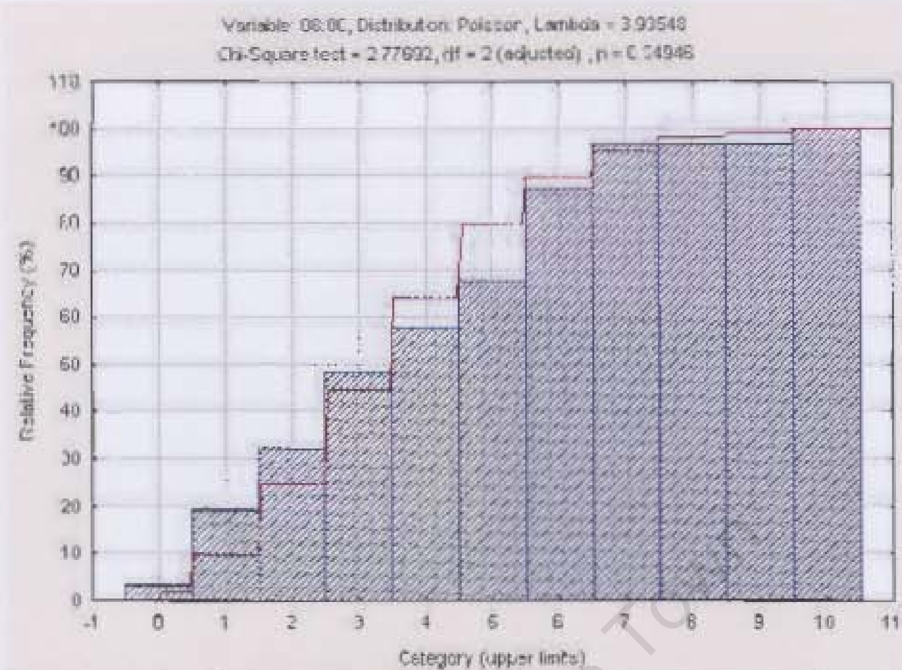

Category	Variable: 08:00, Distribution: Poisson, Lambda = 3.93 (08Arrivals) Chi Square = 2.77692, df = 2 (adjusted) , p = 0.24946 Observed Frequency	Cumulative Observed	Percent Observed	Cumul. % Observed	Expected Frequency	Cumulative Expected	Percent Expected	Cumul. % Expected	Observed-Expected
<= 0.00000	1	1	3.22581	3.2258	0.605624	0.60562	1.95362	1.9536	0.39436
1.00000	5	6	16.12903	19.3548	2.383422	2.98905	7.68846	9.6421	2.61658
2.00000	4	10	12.90323	32.2581	4.689965	7.67901	15.12892	24.7710	-0.68997
3.00000	5	15	16.12903	48.3871	6.152414	13.83142	19.84650	44.6175	-1.15241
4.00000	3	18	9.67742	58.0645	6.053187	19.88461	19.52641	64.1439	-3.05319
5.00000	3	21	9.67742	67.7419	4.764444	24.64906	15.36917	79.5131	-1.76444
6.00000	6	27	19.35484	87.0968	3.125064	27.77412	10.08085	89.5939	2.87494
7.00000	3	30	9.67742	96.7742	1.756949	29.53107	5.66758	95.2615	1.24305
8.00000	0	30	0.00000	96.7742	0.864305	30.39537	2.78808	98.0496	-0.86431
9.00000	0	30	0.00000	96.7742	0.377940	30.77331	1.21916	99.2588	-0.37794
< Infinity	1	31	3.22581	100.0000	0.226686	31.00000	0.73124	100.0000	0.77331

Figure A.9: The results for 08:00

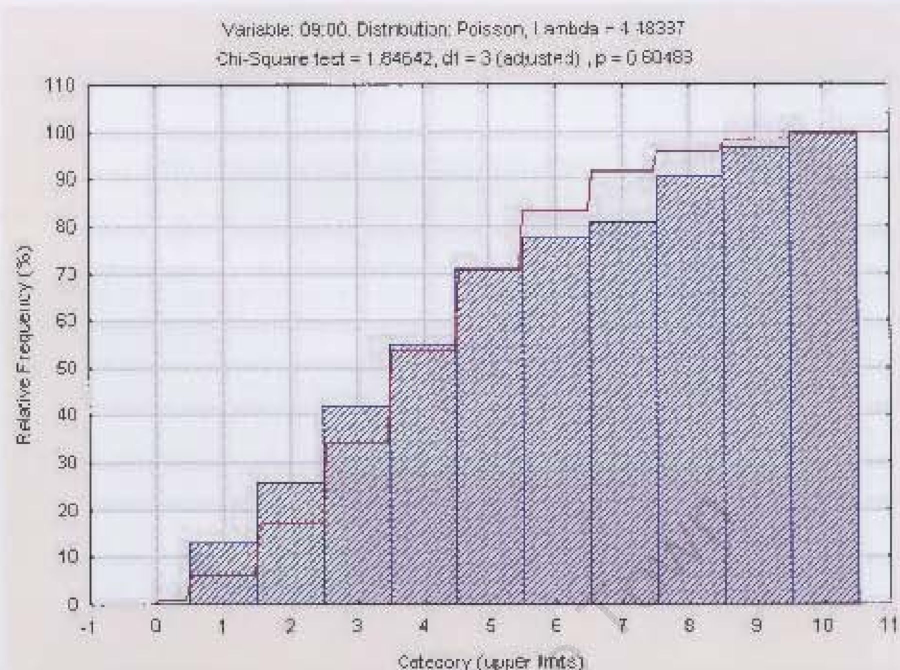

Category	Observed Frequency	Cumulative Observed	Percent Observed	Cumul. % Observed	Expected Frequency	Cumulative Expected	Percent Expected	Cumul. % Expected	Observed-Expected
	Variable: 09:00, Distribution: Poisson, Lambda = 4.48 (09Arrivals) Chi-Square = 1.84642, df = 3 (adjusted) , p = 0.60483								
<= 0.00000	0	0	0.00000	0.0000	0.349978	0.34998	1.12896	1.1290	-0.34998
1.00000	4	4	12.90323	12.9032	1.569258	1.91924	5.06212	6.1911	2.43074
2.00000	4	8	12.90323	25.8065	3.518175	5.43741	11.34895	17.5400	0.48182
3.00000	5	13	16.12903	41.9355	5.258350	10.69576	16.96242	34.5025	-0.25835
4.00000	4	17	12.90323	54.8387	5.894439	16.59020	19.01432	53.5168	-1.89444
5.00000	5	22	16.12903	70.9677	5.285981	21.87618	17.05155	70.5683	-0.28598
6.00000	2	24	6.45161	77.4194	3.950274	25.82646	12.74282	83.3111	-1.95027
7.00000	1	25	3.22581	80.6452	2.530361	28.35682	8.16245	91.4736	-1.53036
8.00000	3	28	9.67742	90.3226	1.418226	29.77504	4.57492	96.0485	1.58177
9.00000	2	30	6.45161	96.7742	0.706572	30.48161	2.27926	98.3276	1.29343
< Infinity	1	31	3.22581	100.0000	0.518385	31.00000	1.67221	100.0000	0.48161

Figure A.10: The results for 09:00

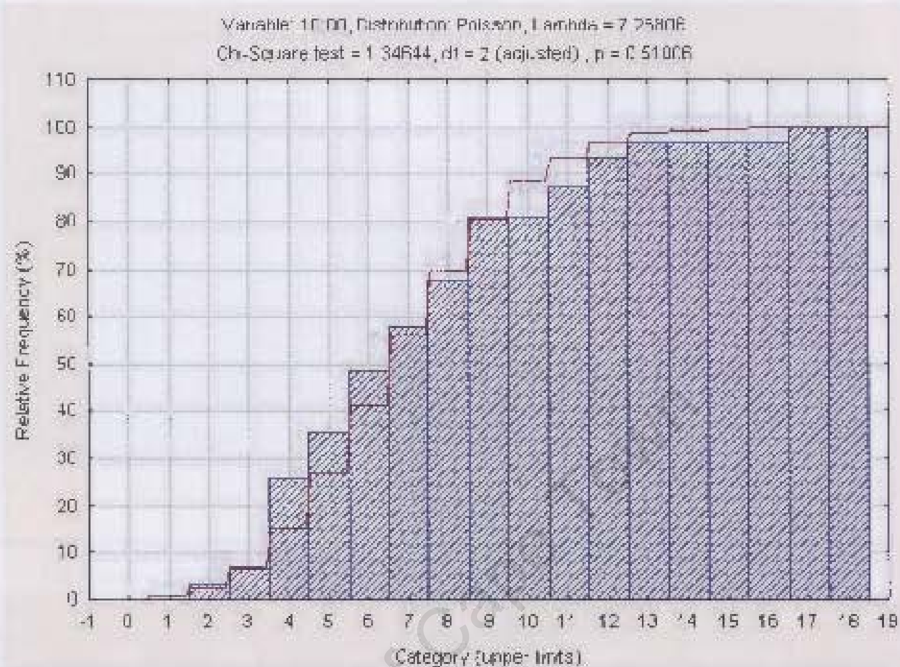

Category	Observed Frequency	Cumulative Observed	Percent Observed	Cumul. % Observed	Expected Frequency	Cumulative Expected	Percent Expected	Cumul. % Expected	Observed-Expected
	Variable 10:00, Distribution: Poisson, Lambda = 7.25 (09/intervals) Chi-Square = 1.34644 df = 2 (adjusted) , p = 0.51006								
<= 0.00000	0	0	0.00000	0.0000	0.021839	0.02'84	0.07045	0.0704	-0.02184
1.00000	0	0	0.00000	0.0000	0.153506	0.18034	0.51131	0.5818	-0.1585'
2.00000	1	1	3.22581	3.2258	0.575223	0.75557	1.85556	2.4373	0.42478
3.00000	1	2	3.22581	6.4516	1.391668	2.14723	4.48925	6.9266	-0.39167
4.00000	6	8	19.35484	25.8065	2.525203	4.67244	8.14582	15.0724	3.47480
5.00000	3	11	9.67742	35.4839	3.665616	8.33806	11.82457	26.8970	-0.66562
6.00000	4	15	12.90323	48.3871	4.434222	12.77228	14.30394	41.2009	-0.43422
7.00000	3	18	9.67742	58.0645	4.597684	17.36996	14.83124	56.0321	-1.59768
8.00000	3	21	9.67742	67.7419	4.171290	21.54125	13.45577	69.4879	-1.17129
9.00000	4	25	12.90323	80.6452	3.363944	24.90520	10.85143	80.3393	0.63606
10.00000	0	25	0.00000	80.6452	2.441571	27.34677	7.87603	88.2154	-2.44157
11.00000	2	27	6.45161	87.0968	1.611008	28.95777	5.19680	93.4122	0.38999
12.00000	2	29	6.45161	93.5484	0.974400	29.93217	3.14322	96.5554	1.02560
13.00000	1	30	3.22581	96.7742	0.544020	30.47619	1.75490	98.3103	0.45598
14.00000	0	30	0.00000	96.7742	0.282038	30.75823	0.90980	99.2201	-0.28204
15.00000	0	30	0.00000	96.7742	0.136470	30.89470	0.44023	99.6603	-0.13647
16.00000	0	30	0.00000	96.7742	0.061907	30.95661	0.19970	99.8600	-0.06191
17.00000	1	31	3.22581	100.0000	0.026431	30.98304	0.08526	99.9453	0.97357
< Infinity	0	31	0.00000	100.0000	0.016961	31.00000	0.05471	100.0000	-0.01696

Figure A.11: The results for 10:00

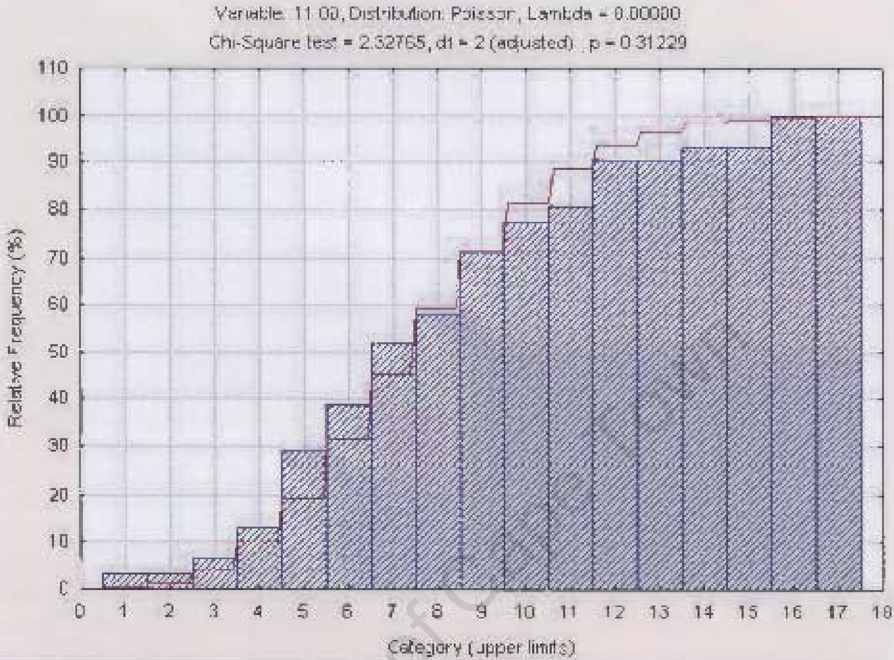

Category	Observed Frequency	Cumulative Observed	Percent Observed	Cumul. % Observed	Expected Frequency	Cumulative Expected	Percent Expected	Cumul. % Expected	Observed-Expected
	Variable: 11.00, Distribution: Poisson, Lambda = 3.00 (09Arrivals) Chi-Squars = 2.32765, df = 2 (adjusted), p = 0.31229								
<= 1.00000	1	1	3.22581	3.2258	0.093594	0.09359	0.30192	0.3019	0.90641
2.00000	0	1	0.00000	3.2258	0.332779	0.42637	1.07348	1.3754	-0.33278
3.00000	1	2	3.22581	6.4516	0.887410	1.31378	2.86261	4.2380	0.11259
4.00000	2	4	6.45161	12.9032	1.774821	3.08860	5.72523	9.9632	0.22518
5.00000	5	9	16.12903	29.0323	2.839714	5.92832	9.16037	19.1236	2.16029
6.00000	3	12	9.67742	38.7097	3.786285	9.71460	12.21382	31.3374	-0.78628
7.00000	4	16	12.90323	51.6129	4.327186	14.04179	13.95866	45.2961	-0.32719
8.00000	2	18	6.45161	58.0645	4.327182	18.36897	13.95865	59.2547	-2.32718
9.00000	4	22	12.90323	70.9677	3.846384	22.21536	12.40769	71.6624	0.15362
10.00000	2	24	6.45161	77.4194	3.077105	25.29246	9.92615	81.5886	-1.07711
11.00000	1	25	3.22581	80.6452	2.237896	27.53036	7.21902	88.8076	-1.23790
12.00000	3	28	9.67742	90.3226	1.491930	29.02229	4.81268	93.6203	1.50807
13.00000	0	28	0.00000	90.3226	0.918111	29.94040	2.96165	96.5819	-0.91811
14.00000	1	29	3.22581	93.5484	0.524635	30.46503	1.69237	98.2743	0.47537
15.00000	0	29	0.00000	93.5484	0.279805	30.74484	0.90260	99.1769	-0.27981
16.00000	2	31	6.45161	100.0000	0.139903	30.88474	0.45130	99.6282	1.86010
< Infinity	0	31	0.00000	100.0000	0.115259	31.00000	0.37180	100.0000	-0.11526

Figure A.12: The results for 11:00

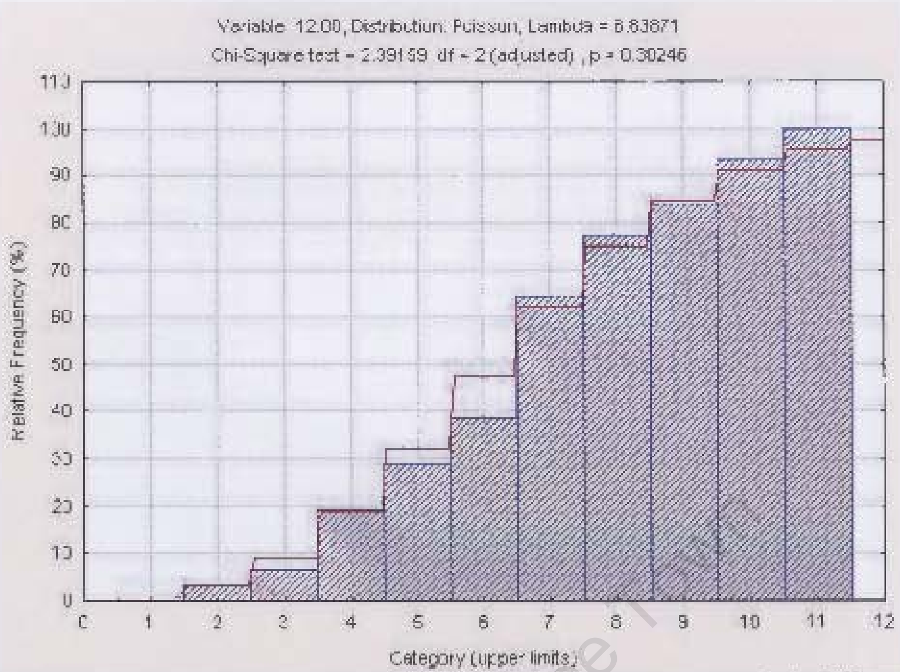

Category	Observed Frequency	Cumulative Observed	Percent Observed	Cumul. % Observed	Expected Frequency	Cumulative Expected	Percent Expected	Cumul. % Expected	Observed-Expected
	Variable: 12:00, Distribution: Poisson, Lambda = 6.83 (09Arrivals) Chi-Square = 2.39159, df = 2 (adjusted) , p = 0.30245								
<= 1.00000	0	0	0.00000	0.0000	0.260371	0.26037	0.83991	0.8399	-0.26037
2.00000	1	1	3.22581	3.2258	0.776723	1.03709	2.50556	3.3455	0.22328
3.00000	1	2	3.22581	6.4516	1.770594	2.80769	5.71159	9.0571	-0.77059
4.00000	4	6	12.90323	19.3548	3.027145	5.83483	9.76498	18.8220	0.97285
5.00000	3	9	9.67742	29.0323	4.140361	9.97519	13.35600	32.1780	-1.14036
6.00000	3	12	9.67742	38.7097	4.719107	14.69430	15.22292	47.4010	-1.71911
7.00000	8	20	25.80645	64.5161	4.610377	19.30468	14.87218	62.2732	3.38962
8.00000	4	24	12.90323	77.4194	3.941129	23.24581	12.71332	74.9865	0.05887
9.00000	2	26	6.45161	83.8710	2.994691	26.24050	9.66029	84.6468	-0.99469
10.00000	3	29	9.67742	93.5484	2.047984	28.28848	6.60640	91.2532	0.95202
< Infinity	2	31	6.45161	100.0000	2.711516	31.00000	8.74683	100.0000	-0.71152

Figure A.13: The results for 12:00

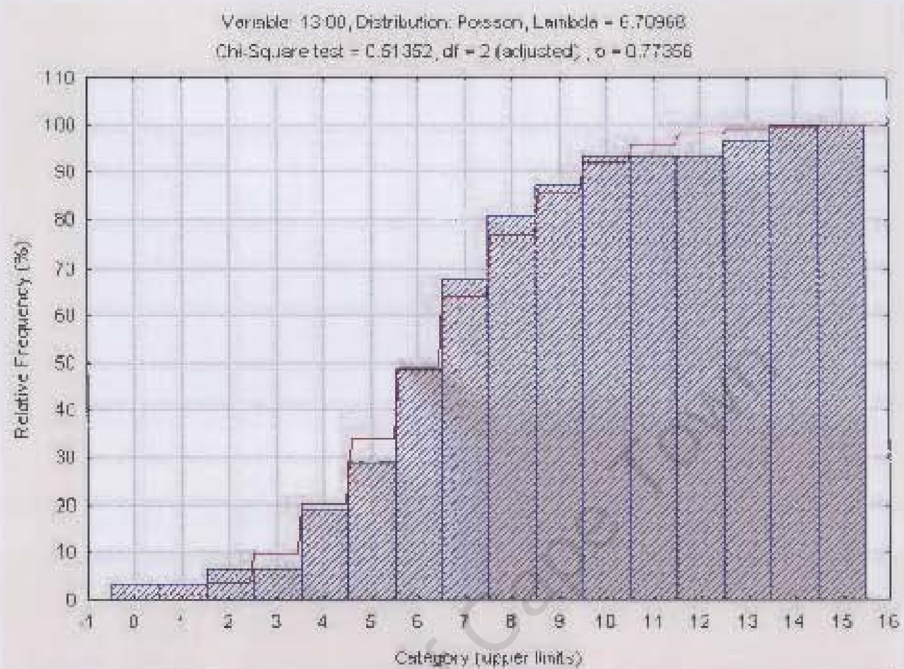

Category	Observed Frequency	Cumulative Observed	Percent Observed	Cumul. % Observed	Expected Frequency	Cumulative Expected	Percent Expected	Cumul. % Expected	Observed-Expected
	Variable: 13:00, Distribution: Poisson, Lambda = 6.70 (00Arrivals) Chi-Square = 0.51352, df = 2 (adjusted) ; p = 0.77356								
<= 0.00000	1	1	3.22581	3.2258	0.037791	0.03779	0.12191	0.1219	0.96221
1.00000	0	1	0.00000	3.2258	0.253564	0.29135	0.81795	0.9399	-0.25356
2.00000	1	2	3.22581	6.4516	0.850666	1.14202	2.74408	3.6839	0.14933
3.00000	0	2	0.00000	6.4516	1.902565	3.04459	6.13731	9.8212	-1.90256
4.00000	4	6	12.90323	19.3548	3.191399	6.23598	10.29484	20.1161	0.80860
5.00000	3	9	9.67742	29.0323	4.282658	10.51864	13.81502	33.9311	-1.28266
6.00000	6	15	19.35484	48.3871	4.789202	15.30784	15.44904	49.3801	1.21080
7.00000	6	21	19.35484	67.7419	4.590567	19.89841	14.80828	64.1884	1.40943
8.00000	4	25	12.90323	80.6452	3.850157	23.74857	12.41986	76.6083	0.14984
9.00000	2	27	6.45161	87.0968	2.870368	26.61894	9.25925	85.8675	-0.87037
10.00000	2	29	6.45161	93.5484	1.925923	28.54486	6.21266	92.0802	0.07408
11.00000	0	29	0.00000	93.5484	1.174757	29.71962	3.78954	95.8697	-1.17476
12.00000	0	29	0.00000	93.5484	0.656853	30.37647	2.11888	97.9886	-0.65685
13.00000	1	30	3.22581	96.7742	0.339021	30.71549	1.09362	99.0822	0.66098
14.00000	1	31	3.22581	100.0000	0.162480	30.87797	0.52413	99.6064	0.83752
< Infinity	0	31	0.00000	100.0000	0.122028	31.00000	0.39364	100.0000	-0.12203

Figure A.14: The results for 13:00

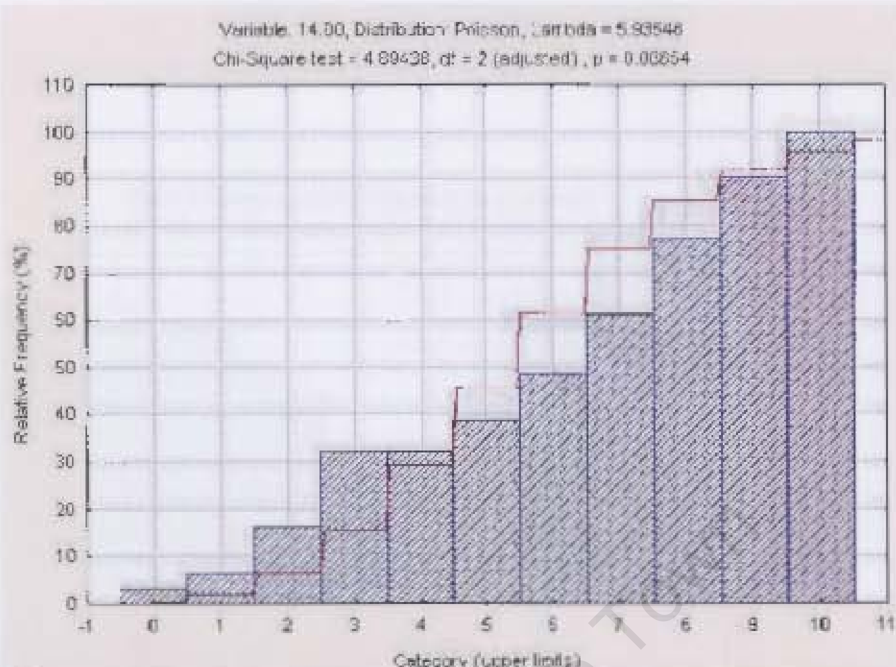

Category	Observed Frequency	Cumulative Observed	Percent Observed	Cumul % Observed	Expected Frequency	Cumulative Expected	Percent Expected	Cumul % Expected	Observed-Expected
Variable: 14:00, Distribution: Poisson, Lambda = 5.93 (09Arrivals) Chi-Square = 4.69438, df = 2 (adjusted), p = 0.08654									
<= 0.00000	1	1	3.22581	3.2258	0.081962	0.08196	0.26439	0.2644	0.91804
1.00000	1	2	3.22581	6.4516	0.486486	0.56845	1.56931	1.8337	0.51351
2.00000	3	5	9.67742	16.1290	1.443764	2.01221	4.65730	6.4910	1.55624
3.00000	5	10	16.12903	32.2581	2.856478	4.86869	9.21445	15.7055	2.14352
4.00000	0	10	0.00000	32.2581	4.236648	9.10734	13.67306	29.3785	-4.23665
5.00000	2	12	6.45161	38.7097	5.031682	14.13902	16.23123	45.6097	-3.03168
6.00000	3	15	9.67742	48.3871	4.977578	19.11660	16.05670	61.6664	-1.97758
7.00000	4	19	12.90323	61.2903	4.220619	23.33722	13.61490	75.2813	-0.22062
8.00000	5	24	16.12903	77.4194	3.131425	26.46864	10.10137	85.3827	1.86857
9.00000	4	28	12.90323	90.3225	2.065171	28.53381	6.66184	92.0446	1.93483
< Infinity	3	31	9.67742	100.0000	2.466167	31.00000	7.95544	100.0000	0.53381

Figure A.15: The results for 14:00

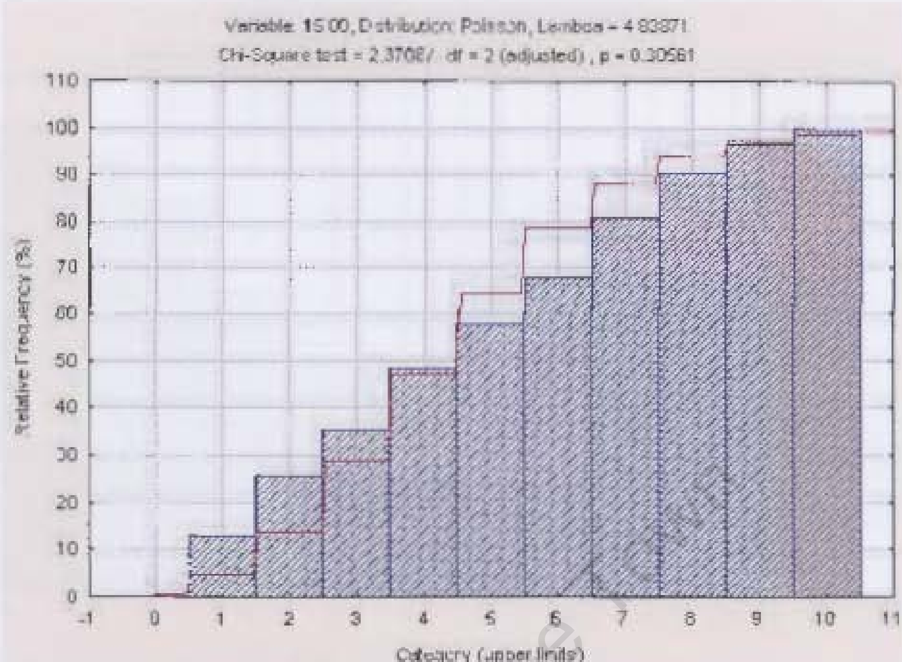

Category	Variable: 15.00, Distribution: Poisson, Lambda = 4.83 (OSArrivals) Chi-Square = 2.37087, df = 2 (adjusted) , p = 0.30661								
	Observed Frequency	Cumulative Observed	Percent Observed	Cumul. % Observed	Expected Frequency	Cumulative Expected	Percent Expected	Cumul. % Expected	Observed-Expected
<= 0.00000	0	0	0.00000	0.0000	0.245435	0.24544	0.79173	0.7917	-0.24544
1.00000	4	4	12.90323	12.9032	1.187589	1.43302	3.83093	4.6227	2.81241
2.00000	4	8	12.90323	25.8065	2.873200	4.30623	9.26839	13.8910	1.12680
3.00000	3	11	9.67742	35.4839	4.634201	8.94043	14.94903	28.8401	-1.63420
4.00000	4	15	12.90323	48.3871	5.605875	14.54630	18.08347	46.9236	-1.60588
5.00000	3	18	9.67742	58.0645	5.425045	19.97135	17.50015	64.4237	-2.42505
6.00000	3	21	9.67742	67.7419	4.375036	24.34638	14.11302	78.5367	-1.37504
7.00000	4	25	12.90323	80.6452	3.024218	27.37060	9.75554	88.2923	0.97578
8.00000	3	28	9.67742	90.3226	1.829165	29.19976	5.90053	94.1928	1.17084
9.00000	2	30	6.45161	96.7742	0.983422	30.18319	3.17233	97.3651	1.01658
< Infinity	1	31	3.22581	100.0000	0.816814	31.00000	2.63488	100.0000	0.18319

Figure A.16: The results for 15:00

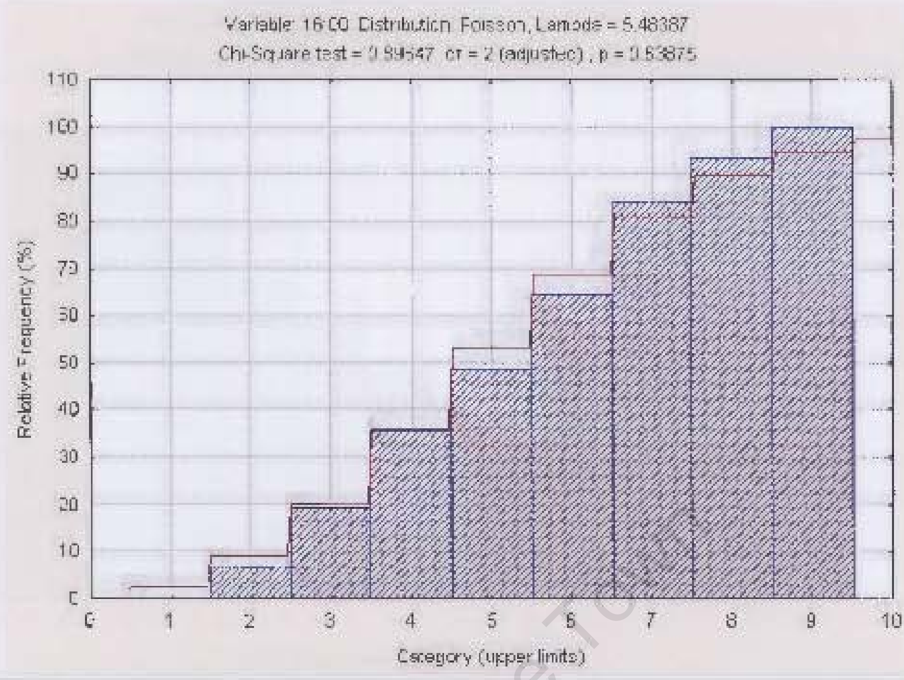

Category	Variable 16:00, Distribution: Poisson, Lambda = 5.48 (09Arrivals) Chi-Square = 0.89347 df = 2 (adjusted) , p = 0.63875								
	Observed Frequency	Cumulative Observed	Percent Observed	Cumul. % Observed	Expected Frequency	Cumulative Expected	Percent Expected	Cumul. % Expected	Observed-Expected
<= 1.00000	0	0	0.00000	0.0000	0.834798	0.83480	2.59290	2.6928	0.83480
2.00000	2	2	6.45161	6.4516	1.935937	2.77073	6.24496	8.9379	0.06406
3.00000	4	6	12.90323	19.3548	3.538810	6.30954	11.41552	20.3534	0.46119
4.00000	5	11	16.12903	35.4839	4.851597	11.16114	15.65031	36.0037	0.14840
5.00000	4	15	12.90323	48.3871	5.321103	16.48224	17.16485	53.1685	-1.32110
6.00000	5	20	16.12903	64.5161	4.863374	21.34562	15.68830	68.8568	0.13663
7.00000	6	26	19.35484	83.8710	3.810014	25.15563	12.29037	81.1472	2.18999
8.00000	3	29	9.67742	93.5484	2.611705	27.76734	8.42485	89.5721	0.38830
< Infinity	2	31	6.45161	100.0000	3.232662	31.00000	10.42794	100.0000	-1.23266

Figure A.17: The results for 16:00

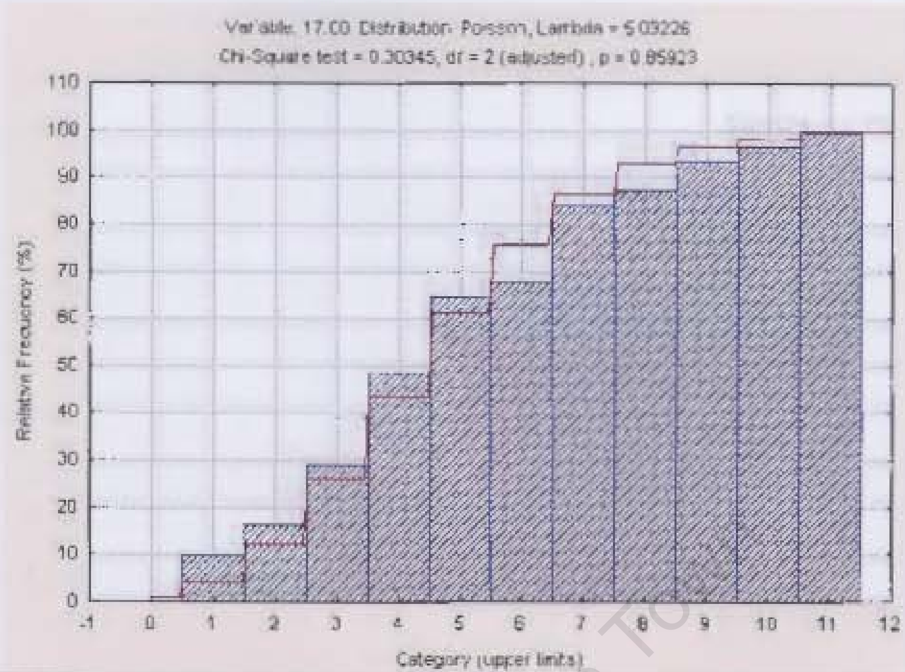

Variable: 17:00 Distribution: Poisson, Lambda = 5.03 (C9Arrivals)
Chi-Square = 0.30345, df = 2 (adjusted) , p = 0.85923

Category	Observed Frequency	Cumulative Observed	Percent Observed	Cumul. % Observed	Expected Frequency	Cumulative Expected	Percent Expected	Cumul. % Expected	Observed-Expected
<= 0.00000	0	0	0.00000	0.0000	0.202246	0.20225	0.65241	0.6524	-0.20225
1.00000	3	3	9.67742	9.6774	1.017754	1.22000	3.28308	3.9355	1.90225
2.00000	2	5	6.45161	16.1290	2.560800	3.78080	8.26064	12.1961	-0.56080
3.00000	4	9	12.90323	29.0323	4.295535	8.07633	13.85658	26.0527	-0.29553
4.00000	6	15	19.35484	48.3871	5.404063	13.48040	17.43246	43.4852	0.59594
5.00000	5	20	16.12903	64.5161	5.438925	18.91932	17.54492	61.0301	-0.43892
6.00000	1	21	3.22581	67.7419	4.561677	23.48100	14.71509	75.7452	-3.56168
7.00000	5	26	16.12903	83.8710	3.278364	26.76036	10.57659	86.3237	1.72164
8.00000	1	27	3.22581	87.0968	2.062026	28.82319	6.65428	92.9780	-1.06203
9.00000	2	29	6.45161	93.5484	1.153407	29.97660	3.72067	96.6987	0.84659
10.00000	1	30	3.22581	96.7742	0.580424	30.55702	1.87234	98.5710	0.41958
< Infinity	1	31	3.22581	100.0000	0.442980	31.00000	1.42897	100.0000	0.55702

Figure A.18: The results for 17:00

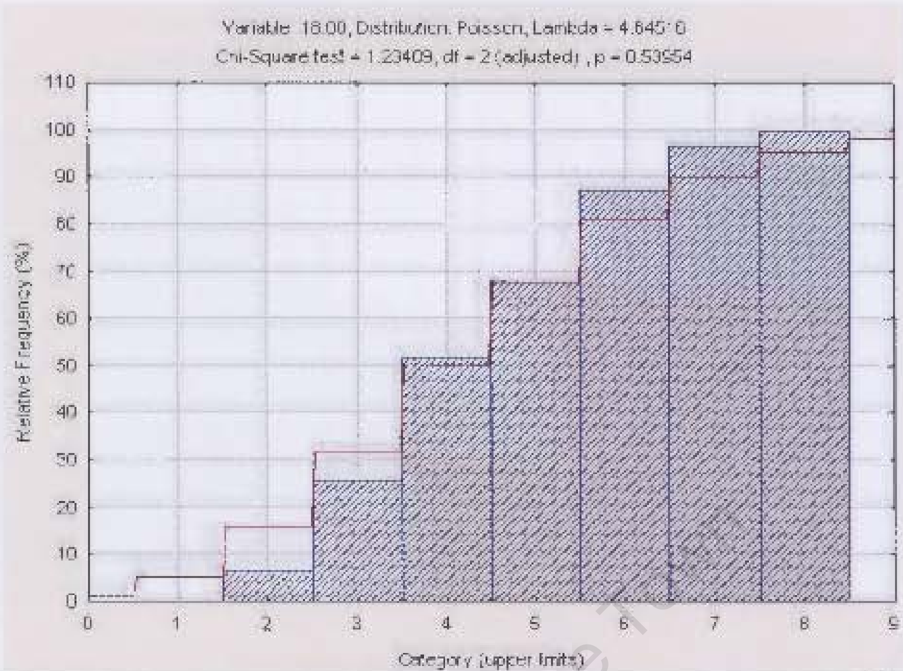

Category	Variable: 18:00, Distribution: Poisson, Lambda = 4.64 (09Arrivals) Chi-Square = 1.23409, df = 2 (adjusted), p = 0.53954 Observed Frequency	Cumulative Observed	Percent Observed	Cumul. % Observed	Expected Frequency	Cumulative Expected	Percent Expected	Cumul. % Expected	Observed-Expected
<= 1.00000	0	0	0.00000	0.0000	1.681396	1.68140	5.42386	5.4239	-1.68140
2.00000	2	2	6.45161	6.4516	3.213404	4.89480	10.36582	15.7897	-1.21340
3.00000	6	8	19.35484	25.8065	4.975597	9.87040	16.05031	31.8400	1.02440
4.00000	8	16	25.80645	51.6129	5.778109	15.64851	18.63906	50.4791	2.22189
5.00000	5	21	16.12903	67.7419	5.368047	21.01655	17.31628	67.7953	-0.36805
6.00000	6	27	19.35484	87.0968	4.155909	25.17246	13.40616	81.2015	1.84409
7.00000	3	30	9.67742	96.7742	2.757838	27.93030	8.89625	90.0977	0.24216
< Infinity	1	31	3.22581	100.0000	3.069958	31.00000	9.90225	100.0000	-2.06970

Figure A.19: The results for 18:00

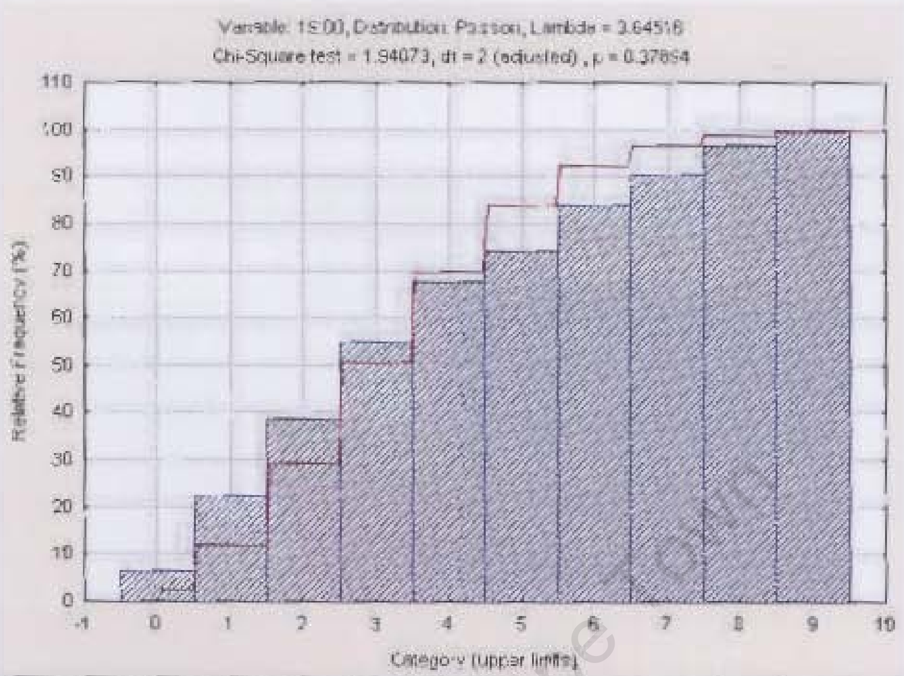

Category	Variable: 19:00, Distribution: Poisson, Lambda = 3.64 (09Arrivals) Chi-Square = 1.94073, df = 2 (adjusted) , p = 0.37894								
	Observed Frequency	Cumulative Observed	Percent Observed	Cumul % Observed	Expected Frequency	Cumulative Expected	Percent Expected	Cumul % Expected	Observed-Expected
<= 0.00000	2	2	6.45161	6.4516	0.809633	0.00963	2.61172	2.6117	1.19037
1.00000	5	7	16.12903	22.5806	2.951243	3.76088	9.52014	12.1319	2.04876
2.00000	5	12	16.12903	38.7097	5.378881	9.13976	17.35123	29.4831	-0.37888
3.00000	5	17	16.12903	54.8387	6.535627	15.67538	21.08267	50.5658	-1.53563
4.00000	4	21	12.90323	67.7419	5.955854	21.63124	19.21243	69.7782	-1.95585
5.00000	2	23	6.45161	74.1935	4.342008	25.97325	14.00643	83.7847	-2.34201
6.00000	3	26	9.67742	83.8710	2.637887	28.61113	8.50931	92.2940	0.36211
7.00000	2	28	6.45161	90.3226	1.373646	29.98478	4.43112	96.7251	0.62635
8.00000	2	30	6.45161	96.7742	0.625895	30.61067	2.01902	98.7441	1.37410
< Infinity	1	31	3.22581	100.0000	0.389325	31.00000	1.25589	100.0000	0.61067

Figure A.20: The results for 19:00

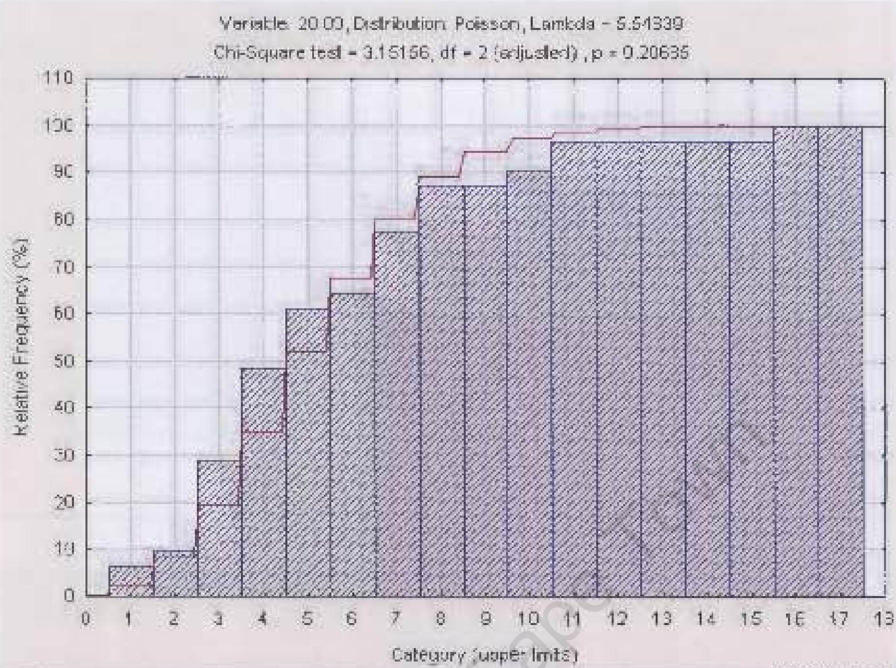

Category	Observed Frequency	Cumulative Observed	Percent Observed	Cumul. % Observed	Expected Frequency	Cumulative Expected	Percent Expected	Cumul. % Expected	Observed-Expected
	Variable: 20:00, Distribution: Poisson, Lambda = 5.54 (D9Arrivals) Chi-Square = 3.15156, df = 2 (adjusted) , p = 0.20635								
<= 1.00000	2	2	6.45161	6.4516	0.790428	0.79043	2.54977	2.5498	1.20957
2.00000	1	3	3.22581	9.6774	1.857938	2.64837	5.99335	8.5431	-0.85794
3.00000	6	9	19.35484	29.0323	3.436187	6.08455	11.08447	19.6276	2.56381
4.00000	6	15	19.35484	48.3871	4.766327	10.85088	15.37525	35.0028	1.23367
5.00000	4	19	12.90323	61.2903	5.289082	16.13996	17.06156	52.0644	-1.28908
6.00000	1	20	3.22581	64.5161	4.890977	21.03094	15.77734	67.8417	-3.89098
7.00000	4	24	12.90323	77.4194	3.876721	24.90766	12.50555	80.3473	0.12328
8.00000	3	27	9.67742	87.0968	2.688693	27.59635	8.67320	89.0205	0.31131
9.00000	0	27	0.00000	87.0968	1.657545	29.25390	5.34692	94.3674	-1.65755
10.00000	1	28	3.22581	90.3226	0.919671	30.17357	2.96668	97.3341	0.08033
11.00000	2	30	6.45161	96.7742	0.463881	30.63745	1.49639	98.8305	1.53612
12.00000	0	30	0.00000	96.7742	0.214483	30.85193	0.69188	99.5224	-0.21448
13.00000	0	30	0.00000	96.7742	0.091541	30.94347	0.29529	99.8177	-0.09154
14.00000	0	30	0.00000	96.7742	0.036279	30.97975	0.11703	99.9347	-0.03628
15.00000	0	30	0.00000	96.7742	0.013419	30.99317	0.04329	99.9780	-0.01342
16.00000	1	31	3.22581	100.0000	0.004653	30.99783	0.01501	99.9930	0.99535
< Infinity	0	31	0.00000	100.0000	0.002175	31.00000	0.00702	100.0000	-0.00217

Figure A.21: The results for 20:00

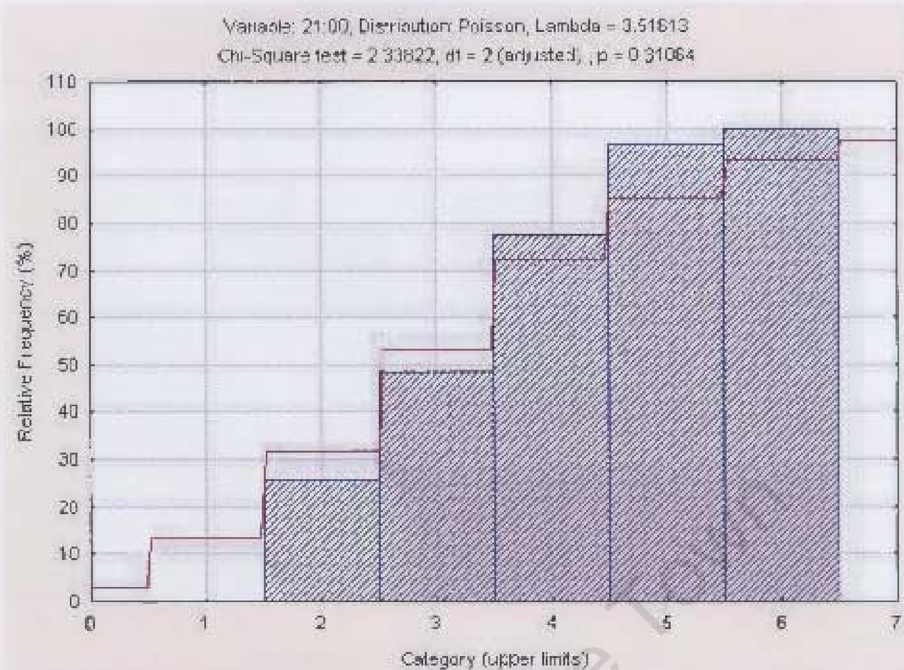

Category	Variable: 21.00, Distribution: Poisson, Lambda = 3.51 (09Arrivals) Chi-Square = 2.33822, df = 2 (adjusted) , p = 0.31064								
	Observed Frequency	Cumulative Observed	Percent Observed	Cumul. % Observed	Expected Frequency	Cumulative Expected	Percent Expected	Cumul. % Expected	Observed-Expected
<= 1.00000	0	0	0.00000	0.0000	4.159993	4.15999	13.41933	13.4193	-4.15999
2.00000	8	8	25.80645	25.8065	5.694111	9.85410	18.36810	31.7874	2.30589
3.00000	7	15	22.58065	48.3871	6.673742	16.52785	21.52820	53.3156	0.32626
4.00000	9	24	29.03226	77.4194	5.866435	22.39428	18.92398	72.2396	3.13357
5.00000	6	30	19.35484	96.7742	4.125427	26.51971	13.30783	85.5474	1.87457
< Infinity	1	31	3.22581	100.0000	4.480292	31.00000	14.45265	100.0000	-3.48029

Figure A.22: The results for 21:00

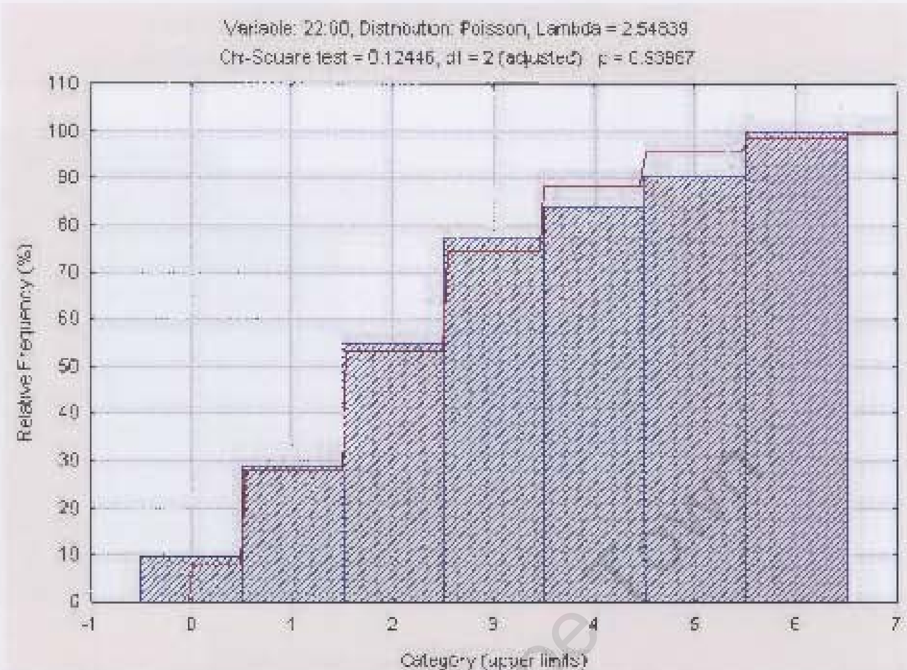

	Variable: 22:00, Distribution: Poisson, Lambda = 2.54 (09Arrivals) Chi-Square = 0.12446, df = 2 (adjusted) , p = 0.93967								
Category	Observed Frequency	Cumulative Observed	Percent Observed	Cumul. % Observed	Expected Frequency	Cumulative Expected	Percent Expected	Cumul. % Expected	Observed-Expected
<= 0.00000	3	3	9.67742	9.6774	2.424439	2.42444	7.82077	7.8208	0.57556
1.00000	6	9	19.35484	29.0323	6.178411	8.60285	19.93036	27.7511	-0.17841
2.00000	8	17	25.80645	54.8387	7.872489	16.47534	25.39512	53.1463	0.12751
3.00000	7	24	22.58065	77.4194	6.687380	23.16272	21.57220	74.7184	0.31262
4.00000	2	26	6.45161	83.8710	4.260510	27.42323	13.74358	88.4620	-2.26051
5.00000	2	28	6.45161	90.3226	2.171485	29.59471	7.00479	95.4668	-0.17149
< Infinity	3	31	9.67742	100.0000	1.405285	31.00000	4.53316	100.0000	1.59471

Figure A.23: The results for 22:00

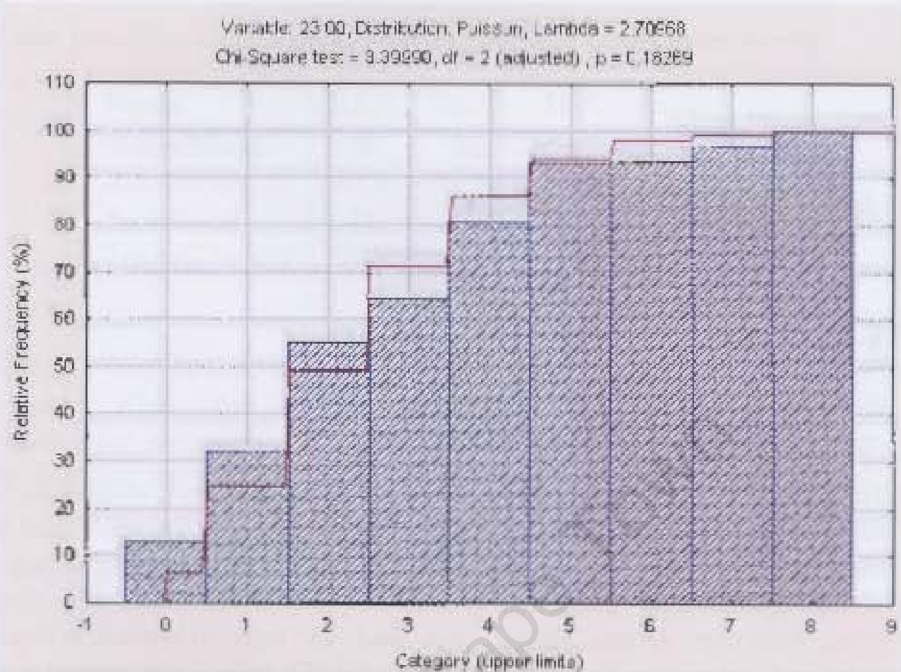

Category	Observed Frequency	Cumulative Observed	Percent Observed	Cumul. % Observed	Expected Frequency	Cumulative Expected	Percent Expected	Cumul. % Expected	Observed-Expected
	Variable: 23:00, Distribution: Poisson, Lambda = 2.70 (OSArrvals) Chi-Square = 3.39990, df = 2 (adjusted) , p = 0.18269								
<= 0.00000	4	4	12.90323	12.9032	2.063306	2.06331	6.65583	6.6558	1.93669
1.00000	6	10	19.35484	32.2581	5.590896	7.65420	18.03515	24.6910	0.40910
2.00000	7	17	22.58065	54.8387	7.574761	15.22896	24.43471	49.1257	-0.57476
3.00000	3	20	9.67742	64.5161	6.841720	22.07068	22.07006	71.1956	-3.84172
4.00000	5	25	16.12903	80.6452	4.634712	26.70540	14.95069	86.1464	0.36529
5.00000	4	29	12.90323	93.5484	2.511716	29.21711	8.10231	94.2487	1.48828
6.00000	0	29	0.00000	93.5484	1.134323	30.35143	3.65911	97.9079	-1.13432
7.00000	1	30	3.22581	96.7742	0.439093	30.79053	1.41643	99.3243	0.56091
< Infinity	1	31	3.22581	100.0000	0.209473	31.00000	0.67572	100.0000	0.79053

Figure A.24: The results for 23:00

Appendix B

Parameter values

Table B.1 lists all of the values used on all of the work centres in the final model. At the beginning of Chapter 7, it was mentioned that it was necessary to adjust some of the parameter values when conducting some of the experiments. Table B.2 lists those parameters which were changed, along with their altered values. 'NA' and 'IA' refer to normal arrivals and increased arrivals respectively, and 'I' and 'D' refer to increased and decreased values respectively. So 'NA/I' would mean that the normal arrival pattern was used, and the parameter for that particular work centre was being increased. As can be seen in the table, it was not necessary to perform this for all of the values.

Table B.1: The parameter values used in the final model

Node	Value	Distribution
GreenConsult	10 - 30	Uniform
GreenXRays	25	Erlang 15
YellowConsult	15 - 40	Uniform
YellowBlood	86	Erlang 15
YellowXRays	30	Erlang 15
OrangeConsult	30 - 150	Uniform
RedConsult	30 - 180	Uniform
XandBlood	51	Erlang 15
AnteXRays	27	Erlang 15
AnteBlood	83.75	Erlang 15
TrialofTherapy	61.25	Erlang 15
Wait	60.25	Erlang 15

Table B.2: The altered parameter values used in the experiments

Node	NA/I	NA/D	IA/I	IA/D
YellowBlood	76	92	-	-
XandBlood	47.5	52.75	-	51.9
AnteBlood	74	89	-	-
TrialofTherapy	58.5	62.6	56.25	59.75

Table B.3 lists the final inter-arrival times for each hour which were used in the model. Along with this, it also lists the inter-arrival times for each hour which were used when conducting the experiments in Chapter 7 (when the arrivals were increased and decreased).

Table B.3: The inter-arrival times for each hour

Hour	Original	Decreased	Increased
00:00	22.36	41.52	18.16
01:00	32.52	60.39	26.12
02:00	35.77	66.43	29.06
03:00	44.71	83.04	36.33
04:00	49.34	91.63	40.09
05:00	36.69	68.13	29.81
06:00	52.99	98.41	43.06
07:00	22.01	40.88	17.88
08:00	11.73	21.78	9.53
09:00	10.29	19.12	8.36
10:00	6.36	11.81	5.17
11:00	5.77	10.71	4.69
12:00	6.75	12.53	5.48
13:00	6.88	12.77	5.59
14:00	7.78	14.44	6.32
15:00	9.54	17.71	7.75
16:00	8.42	15.63	6.84
17:00	9.17	17.03	7.45
18:00	9.91	18.45	8.07
19:00	12.66	23.51	10.29
20:00	8.32	15.45	6.76
21:00	13.13	24.38	10.67
22:00	18.11	33.63	14.72
23:00	17.03	31.63	13.84

Appendix C

Experimental designs

This section contains the experimental designs including the results for each run of each replication for the experiments conducted in Chapter 7. More details of these designs can be found in [12].

	1 Replicate	2 Arrivals	3 Doctors	4 GreenDischarge
1	1	-1	-1	208.96
2	1	-1	1	87.17
3	1	1	-1	240.48
4	1	1	1	97.73
5	2	-1	-1	209.06
6	2	-1	1	86.92
7	2	1	-1	240.68
8	2	1	1	97.75
9	3	-1	-1	209.25
10	3	-1	1	86.85
11	3	1	-1	240.65
12	3	1	1	97.74

Figure C.1: The experimental design for the green area

	1 Replicate	2 Arrivals	3 Doctors	4 BloodTime	5 YellowDischarge
1	1	-1.00000	-1.00000	1.00000	337.21
2	1	1.00000	-1.00000	-1.00000	383.85
3	1	-1.00000	1.00000	-1.00000	146.22
4	1	1.00000	1.00000	1.00000	265.55
5	1	-1.00000	-1.00000	-1.00000	260.47
6	1	1.00000	-1.00000	1.00000	488.3
7	1	-1.00000	1.00000	1.00000	201.91
8	1	1.00000	1.00000	-1.00000	165.5
9	2	-1.00000	-1.00000	1.00000	335.65
10	2	1.00000	-1.00000	-1.00000	383.7
11	2	-1.00000	1.00000	-1.00000	145.27
12	2	1.00000	1.00000	1.00000	264.47
13	2	-1.00000	-1.00000	-1.00000	281.12
14	2	1.00000	-1.00000	1.00000	491.09
15	2	-1.00000	1.00000	1.00000	203.59
16	2	1.00000	1.00000	-1.00000	166.16
17	3	-1.00000	-1.00000	1.00000	331.63
18	3	1.00000	-1.00000	-1.00000	382.48
19	3	-1.00000	1.00000	-1.00000	144.74
20	3	1.00000	1.00000	1.00000	261.25
21	3	-1.00000	-1.00000	-1.00000	279.09
22	3	1.00000	-1.00000	1.00000	486.59
23	3	-1.00000	1.00000	1.00000	204.11
24	3	1.00000	1.00000	-1.00000	167.69

Figure C.2: The experimental design for the yellow area

	1 Replicate	2 Arrivals	3 Doctors	4 BloodTime	5 AnteDischarge
1	1	-1.00000	-1.00000	1.00000	256.93
2	1	1.00000	-1.00000	-1.00000	441.48
3	1	-1.00000	1.00000	-1.00000	263.43
4	1	1.00000	1.00000	1.00000	659.55
5	1	-1.00000	-1.00000	-1.00000	337.3
6	1	1.00000	-1.00000	1.00000	728.58
7	1	-1.00000	1.00000	1.00000	342.62
8	1	1.00000	1.00000	-1.00000	334.96
9	2	-1.00000	-1.00000	1.00000	257.53
10	2	1.00000	-1.00000	-1.00000	439.14
11	2	-1.00000	1.00000	-1.00000	269.83
12	2	1.00000	1.00000	1.00000	656.24
13	2	-1.00000	-1.00000	-1.00000	336.57
14	2	1.00000	-1.00000	1.00000	721.1
15	2	-1.00000	1.00000	1.00000	341.67
16	2	1.00000	1.00000	-1.00000	334.3
17	3	-1.00000	-1.00000	1.00000	255.6
18	3	1.00000	-1.00000	-1.00000	436.59
19	3	-1.00000	1.00000	-1.00000	267.49
20	3	1.00000	1.00000	1.00000	649.95
21	3	-1.00000	-1.00000	-1.00000	333.06
22	3	1.00000	-1.00000	1.00000	725.12
23	3	-1.00000	1.00000	1.00000	340.46
24	3	1.00000	1.00000	-1.00000	334.28

Figure C.3: The experimental design for the ante-room area

	1 Replicate	2 Arrivals	3 YellowDoctors	4 AnteDoctors	5 BloodTime	6 HoldingWait	7 HoldDischarge
1	1	-1.00000	-1.00000	-1.00000	-1.00000	1.00000	891.64
2	1	1.00000	-1.00000	-1.00000	-1.00000	-1.00000	892.56
3	1	-1.00000	1.00000	-1.00000	-1.00000	-1.00000	572.53
4	1	1.00000	1.00000	-1.00000	-1.00000	1.00000	1587.72
5	1	-1.00000	-1.00000	1.00000	-1.00000	-1.00000	537.41
6	1	1.00000	-1.00000	1.00000	-1.00000	1.00000	1550.32
7	1	-1.00000	1.00000	1.00000	-1.00000	1.00000	786.42
8	1	1.00000	1.00000	1.00000	-1.00000	-1.00000	742.81
9	1	-1.00000	-1.00000	-1.00000	1.00000	-1.00000	687.6
10	1	1.00000	-1.00000	-1.00000	1.00000	1.00000	1819.17
11	1	-1.00000	1.00000	-1.00000	1.00000	1.00000	933.25
12	1	1.00000	1.00000	-1.00000	1.00000	-1.00000	1091.01
13	1	-1.00000	-1.00000	1.00000	1.00000	1.00000	904.17
14	1	1.00000	-1.00000	1.00000	1.00000	-1.00000	1078.58
15	1	-1.00000	1.00000	1.00000	1.00000	-1.00000	588.62
16	1	1.00000	1.00000	1.00000	1.00000	1.00000	1660.1
17	2	-1.00000	-1.00000	-1.00000	-1.00000	1.00000	927.57
18	2	1.00000	-1.00000	-1.00000	-1.00000	-1.00000	911.57
19	2	-1.00000	1.00000	-1.00000	-1.00000	-1.00000	580.79
20	2	1.00000	1.00000	-1.00000	-1.00000	1.00000	1648.98
21	2	-1.00000	-1.00000	1.00000	-1.00000	-1.00000	546.31
22	2	1.00000	-1.00000	1.00000	-1.00000	1.00000	1607.17
23	2	-1.00000	1.00000	1.00000	-1.00000	1.00000	823.79
24	2	1.00000	1.00000	1.00000	-1.00000	-1.00000	759.83
25	2	-1.00000	-1.00000	-1.00000	1.00000	-1.00000	694.11
26	2	1.00000	-1.00000	-1.00000	1.00000	1.00000	1872.03
27	2	-1.00000	1.00000	-1.00000	1.00000	1.00000	966.24
28	2	1.00000	1.00000	-1.00000	1.00000	-1.00000	1107.59
29	2	-1.00000	-1.00000	1.00000	1.00000	1.00000	939.34
30	2	1.00000	-1.00000	1.00000	1.00000	-1.00000	1093.91
31	2	-1.00000	1.00000	1.00000	1.00000	-1.00000	593.43
32	2	1.00000	1.00000	1.00000	1.00000	1.00000	1721.56
33	3	-1.00000	-1.00000	-1.00000	-1.00000	1.00000	874.11
34	3	1.00000	-1.00000	-1.00000	-1.00000	-1.00000	871.55
35	3	-1.00000	1.00000	-1.00000	-1.00000	1.00000	568.41
36	3	1.00000	1.00000	-1.00000	-1.00000	1.00000	1567.42
37	3	-1.00000	-1.00000	1.00000	-1.00000	-1.00000	535.7
38	3	1.00000	-1.00000	1.00000	-1.00000	1.00000	1526.78
39	3	-1.00000	1.00000	1.00000	-1.00000	1.00000	775.05
40	3	1.00000	1.00000	1.00000	-1.00000	-1.00000	720.34
41	3	-1.00000	-1.00000	-1.00000	1.00000	-1.00000	682.09
42	3	1.00000	-1.00000	-1.00000	1.00000	1.00000	1796.81

Figure C.4: The experimental design for the holding area