Paper 1682-2014

Big Data Analysis for Resource-Constrained Surgical Scheduling

Elizabeth Rowse, Cardiff University; Paul Harper, Cardiff University

ABSTRACT

The scheduling of surgical operations in a hospital is a complex problem, with each surgical specialty having to satisfy their demand while competing for resources with other hospital departments. This project extends the construction of a weekly timetable, the Master Surgery Schedule, which assigns surgical specialties to operating theater sessions by taking into account the post-surgery resource requirements, primarily post-operative beds on hospital wards. Using real data from the largest teaching hospital in Wales, UK, this paper describes how SAS® has been used to analyze large data sets to investigate the relationship between the operating theater schedule and the demand for beds on wards in the hospital. By understanding this relationship, a more well-informed and robust operating theater schedule can be produced that delivers economic benefit to the hospital and a better experience for the patients by reducing the number of cancelled operations caused by the unavailability of beds on hospital wards.

INTRODUCTION

The scheduling of surgical operations in a hospital is a complex problem due to the stochastic nature of healthcare demand, the length of the surgical procedures, and the large number of resource constraints. This paper presents the analysis of large healthcare datasets using SAS® that was performed as part of a PhD research project. The availability of post-operative beds is critical to the scheduling of surgical procedures and the throughput of patients in a hospital. It is important to understand the relationship between the operating theatre schedule and the impact that it has on the demand for beds on hospital wards. Insights gained from the data analysis performed using SAS have been used within a novel robust optimisation approach to the formulation of the Master Surgery Schedule that uniquely takes into account the demand for post-operative beds. The impact of this research is potentially far-reaching and applicable to hospitals and healthcare providers generally.

OPERATING THEATRE SCHEDULING

Operating theatres are very expensive and resource-intensive facilities within modern hospitals, the efficiency of which has a significant impact on patient throughput and the patient experience. A large amount of work has been published in both Operational Research and medical journals on the challenging problem of constructing a weekly timetable, the Master Surgery Schedule (MSS), that assigns surgical specialties to operating theatre sessions whilst taking account of restrictions on resources (Cardoen et al, 2010). These resources can be surgeons, skilled nurses, specialist theatre equipment and the operating theatres themselves. Other aspects to consider while creating the MSS are the stochastic nature of the demand from hospital waiting lists and that of different surgical procedures; this includes both the occurrence of emergency patients and the duration of the procedures. The process of creating a theatre schedule can therefore be quite arduous and time-consuming, as often there is no systematic approach that exploits Operational Research techniques. An example of an operating theatre MSS is shown in Figure 1.

	Theatre 0	Theatre 1	Theatre 2	Theatre 3	Theatre 4	Theatre 5	Theatre 6	Theatre 7	Theatre 8	Theatre 9	Theatre 10	Theatre 11	Theatre 12	Theatre 14
MonAM	Trauma	Scoliosis	ENT	Renal	Oral	CEPOD	Urology	Colorectal	General	Thoracic	Cardiac	Cardiac	Neuro	Neuro
MonPM	Trauma	Scoliosis	ENT	Renal	Oral	CEPOD	Urology	Colorectal	General	Thoracic	Cardiac	Cardiac	Neuro	Neuro
TuesAM	Trauma	Vascular	Ophthal	Vascular	Paeds Gen	CEPOD	Urology	Colorectal	Renal	Thoracic	Cardiac	Cardiac	Neuro	Neuro
TuesPM	Trauma	Vascular	Ophthal	Vascular	Paeds Gen	CEPOD	Urology	Colorectal	Renal	Thoracic	Cardiac	Cardiac	Neuro	Neuro
WedAM	Trauma	Paeds Ortho	ENT	General	Paeds Gen	CEPOD	Urology	Colorectal	General	Thoracic	Cardiac	Cardiac	Neuro	Neuro
WedPM	Trauma	Paeds Ortho	ENT	General	Paeds Gen	CEPOD	Urology	Colorectal	General	Thoracic	Cardiac	Cardiac	Neuro	Neuro
ThurAM	Trauma	Trauma	Oral	Vascular	Paeds Gen	CEPOD	Urology	Colorectal	General	Thoracic	Cardiac	Cardiac	Neuro	Neuro
ThurPM	Trauma	Trauma	Oral	Vascular	Paeds Gen	CEPOD	Urology	Colorectal	General	Thoracic	Cardiac	Cardiac	Neuro	Neuro
FriAM	Trauma	Scoliosis	Paeds ENT	Vascular	Paeds Gen	CEPOD	Urology	Renal	Liver	Oral	Cardiac	Cardiac	Neuro	Neuro
FriPM	Trauma	Scoliosis	ENT	Vascular	Paeds Gen	CEPOD	Urology	Renal	Liver	Oral	Cardiac	Cardiac	Neuro	Neuro

Figure 1. An Example of an Operating Theatre Timetable at the Case Study Hospital

The work covered in this paper includes, for the first time, the constraints of post-operative resource requirements, primarily beds on hospital wards, when devising a practicable and efficient MSS. Every inpatient that has an operation in hospital requires a bed on a ward after surgery, so it is obvious that the theatre schedule will affect the

demand for the beds on hospital wards as scheduling different surgical specialties at different times will result in a different amount of inpatients requiring a bed. Beds are a limited resource in a hospital, especially for speciality-specific nursing. The situation is further complicated by the un-predictable occurrence of emergency patients. These emergency patients will often have priority over elective inpatients due to the severity and urgency of their medical conditions. Hence it is not always certain how many beds are available to elective inpatients, sometimes causing the cancellation of elective operations. This is a situation that needs to be avoided as it can be distressing for patients and costly in terms of surgical resource deferral and the time for the administration to reschedule elective operations.

A novel robust optimisation approach has been taken to include the effect of the availability of post-surgery beds when constructing a MSS. The intended outcome of which is that there will be fewer cancelled operations and better bed management on the hospital wards. A sophisticated robust optimisation and simulation model has been developed to construct the operating theatre timetable when the demand for post-operative beds is considered. The focus of this paper is the analysis of the relationship between the theatre schedule and the demand for beds on wards in the hospital. SAS has been used to analyse large datasets from a case study hospital as described in the next section. By understanding the relationships between the theatre schedule and the impact it has on bed usage, a more well informed and robust theatre schedule can be produced that delivers economic benefit to the hospital and a better experience for the patients.

CASE STUDY

The largest teaching hospital in Wales, UK, has provided data for use in this project. Approximately 25,000 surgical procedures are performed at this hospital each year, with around 12,000 of these being elective inpatient operations that are complex procedures and require the patient to be in a bed for post-surgery recovery. These inpatient operations are performed in a suite of 14 operating theatres that are available for use by 18 surgical specialties which are shown in Figure 1.

A common occurrence in hospitals in the UK is that a large proportion of elective operations are cancelled due to a variety of reasons. Some reasons, such as the deterioration of the medical condition of the patient, are beyond the hospital’s control; however the hospital can affect the occurrence of some cancellations, such as those due to the unavailability of resources. At the case study hospital, approximately 18% of scheduled operations are cancelled annually with an average of 32% of these cancellations being caused by lack of bed availability on hospital wards. This situation clearly has a significant economic and personal impact.

DATA PREPARATION

Two large datasets were provided by the case study hospital; one containing details of each surgical operation performed in the hospital, and the other with information on the admission and discharge of every patient in the hospital. The datasets contain records for every patient admission, operation and discharge from April 2009 to March 2013. The theatre dataset contains 153,790 records, and the admissions dataset contains 438,383 records.

SAS has been used to efficiently analyse these large NHS (National Health Service) datasets. The datasets are data-rich and are mined using SAS to extract meaningful statistics for use in the robust optimisation model that produces the MSS. An initial search for outliers was performed on both datasets, removing any unreasonable outliers; the outliers removed were clearly outliers caused by errors in data input. For example, records that showed operations taking several days to be performed were clearly an error in data input so were removed, however operations that took up to around 30 hours were deemed to be feasible and so were retained.

One feature of SAS that was of particular importance to this project was the ability to merge datasets based on a single variable by using a simple DATA step. The theatre and admissions datasets needed to be merged on a patient identifier variable in order to obtain a master dataset that can be used to analyse all aspects of a patient’s stay in hospital.

The SQL procedure was used to perform a full join on the two sorted datasets, recognising the possibility of multiple records corresponding to the same patient, as shown in the code sample below.

```
/* Using SQL to merge theatre and admissions datasets */
proc sql;
    create table M201213.MergedHospAdmitWard as
    select *
    from M201213.AdmissionsData A FULL JOIN M201213.TheatreData B
    on (A.Patient = B.Patient) and (A.Episode_Start_Date <= B.SessionDate) and
            (A.EpisodeEndDate >= B.SessionDate);
    order by Patient, SurgeryDate;
quit;
```

A conditional merge is used to ensure that the correct operation record is matched with the correct patient admission record. To do this, the datasets were merged on a patient identifier variable and on the condition that the operation date in the theatre dataset lies between the admission and discharge dates in the admissions dataset. The resulting master dataset can then be used for analysis on post-operative length of stay as discussed later. A full join was chosen to be the most suitable type of join because patients might have records in the dataset for more than one operation and/or more than one episode in hospital. Hence the full join ensured that all combinations of operation date and admissions were considered and then filtered out using the conditional statements.

DATA ANALYSIS

Surgical procedures can vary in duration between the different surgical specialties, and within each surgical specialty depending on the complexity of the procedure. It is not always known during the planning process how long each procedure will take. Even if an estimate is used, the unpredictable nature of surgical operations can soon negate any estimates. For each surgical specialty, different numbers of operations are carried out in each operating theatre session. This is an important aspect needed for capacity planning, as the expected number of patients that will require a post-operative bed will need to be estimated. This information is used as an input to the MSS construction model. The box and whisker plots shown in Figure 2 summarise the planned number of operations per session for each specialty in the year 2012/2013.

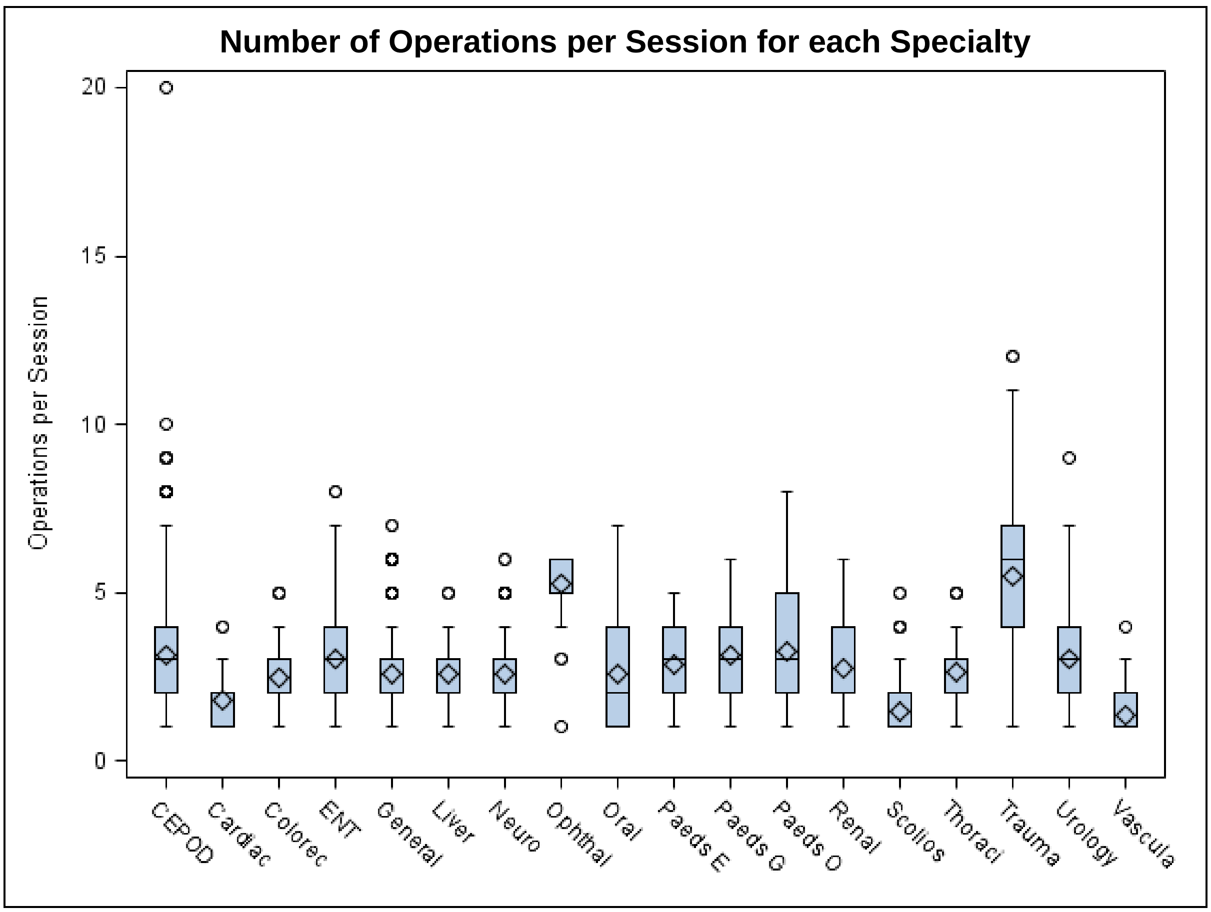

Figure 2. Box and Whisker Plots of the Number of Operations per Session

As expected, specialties that generally involve more complex, and hence longer, operations have fewer operations per session than other specialties that perform less complex operations. For example, more ophthalmology operations are able to be performed in an operating session than cardiac operations. It can also be seen that there is great variability in the number of operations performed in each session.

Post-operative length of stay distributions are typically skewed to the right with a long tail towards longer length of stays, and there are often outliers in the data that are vastly longer than the majority of length of stay values (Marazzi et al, 1998). The distribution of post-operative length of stay for all surgical specialties is shown in Figure 3.

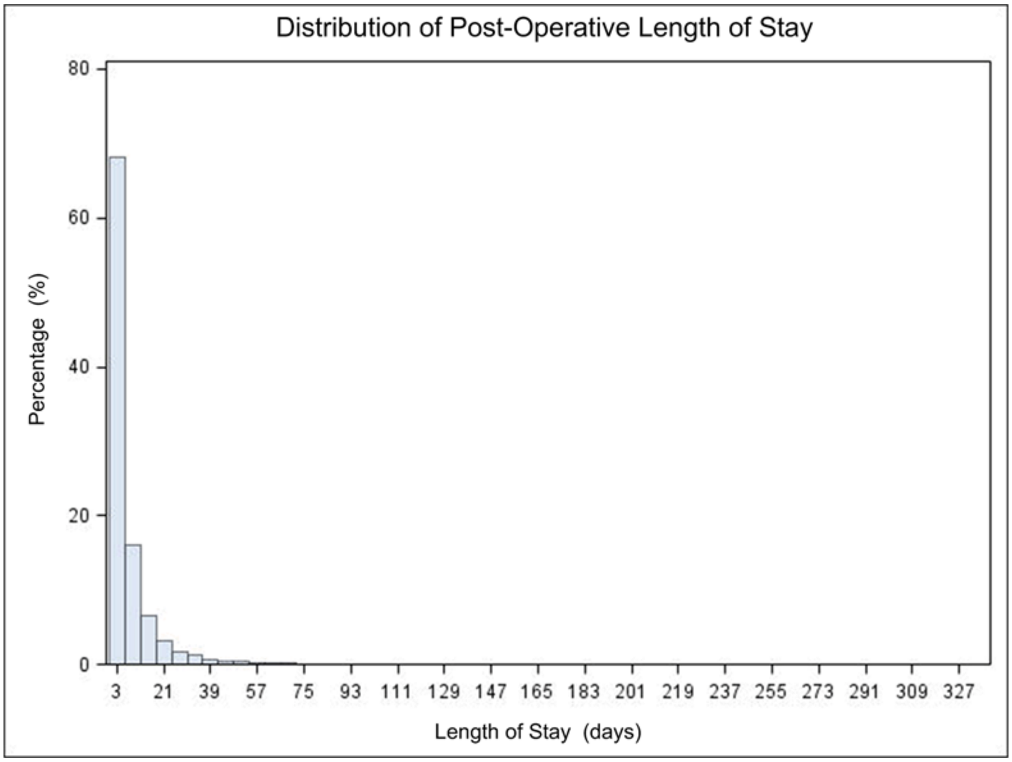

Figure 3. Graph to Show the Distribution of Post-op Length of Stay

For each surgical specialty, the Anderson-Darling and Kolmogorov-Smirnov goodness of fit tests were performed using SAS to test whether a continuous distribution, such as the Lognormal, Weibull and Gamma distributions, could be fitted to the empirical data. Following analysis of all surgical specialties in turn, it was found that the post-op length of stay for all specialties could not be modelled using a known statistical distribution (at the 5% significance level). This implies that Monte Carlo sampling from the empirical distributions of length of stay will have to be performed for use within the optimisation model to construct the MSS.

From analysis of the data, it has also been found that 31% of patients who have an operation are moved to a post-operative bed on a ward that is not directly related to their speciality. This means that the ward is not their surgical specialty's ward and so possibly may not have specifically skilled nurses or specialist equipment. Many ethical issues arise from patients being on different wards (Audit Commission, 2003), however these will not be discussed here. Instead, it is of interest to determine if being on the non-specialist ward affects the post-op length of stay of surgical patients. The working hypothesis was that the length of stay would be longer on non-specialist wards due to not having sufficient access to specifically skilled nurses or specialist equipment. A series of Mann-Whitney tests were performed for each specialty using the NPAR1WAY procedure. The analysis showed that the post-op length of stay is no different on the non-specialist ward than on the specialist ward for the majority of surgical specialties. This is contradictory to the original hypothesis, however, by taking these findings back to key stakeholders in the hospital, it was possible to understand this result; surgeons are more likely to put a fitter patient on a non-specialist ward if no beds are available on their specialist ward, hence the length of stays on both wards tend to be similar.

Of particular interest to hospital managers was an investigation of the relationship, if any, between the total time of an operation spent in theatre and the post-operative length of stay. It was hypothesised that the longer the time in theatre, the longer the stay in a hospital bed afterwards as the more complex surgeries often take longer to perform and require longer to recover due to the severity of the medical condition. It was found that, across all surgical specialties in the case study hospital, the time spent in theatre was statistically significantly weakly positively correlated with post-operative length of stay ($r = 0.16$, $p\text{-}value < 0.0001$). This confirms the hypothesis that they are

both correlated, however not as strongly as suspected. Further work will be required to investigate this relationship at specialty level.

CONCLUSION

There is a need for a smarter way of scheduling operations so that the impact on the demand for post-operative beds on hospital wards is considered and integrated into the scheduling process. In order to fully understand these closely linked resources, an in depth analysis of data is required. Making use of SAS procedures and sophisticated statistical modelling for the analysis of large healthcare datasets has made this process very efficient and productive.

REFERENCES

- Audit Commission. 2003. *Bed Management: Review of National Findings*: Audit Commission.
- Cardoen B., Demeulemeester E., and Belien J., 2010. “Operating room planning and scheduling: A literature review”. *European Journal of Operational Research*, 201, pp. 921-932.
- Marazzi A., Paccaud F., Ruffieux C. and Beguin C., 1998. “Fitting the distributions of length of stay by parametric models”. *Medical Care*, 36(6), pp. 915-27.

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the author at:

Elizabeth Rowse
School of Mathematics, Cardiff University
Senghennydd Road, Cardiff, CF11 9LQ, UK
RowseEL@cardiff.ac.uk
uk.linkedin.com/in/elizabethrowse/