

Loughborough University
Institutional Repository

*An application of demand
profiling and optimisation of
staffing levels within
Leicestershire Police Force*

This item was submitted to Loughborough University's Institutional Repository by the/an author.

Citation: EDLESTON, O.S.S.T. and BARTLETT, L.M., 2006. An application of demand profiling and optimisation of staffing levels within Leicestershire Police Force. Proceedings of the Operational Research Society Simulation Workshop 2010 (SW10).

Additional Information:

- This conference paper was presented at the Operational Research Society Simulation Workshop 2010: <http://www.theorsociety.com/>

Metadata Record: <https://dspace.lboro.ac.uk/2134/6144>

Version: Accepted for publication

Publisher: © Operational Research Society

Please cite the published version.

This item was submitted to Loughborough's Institutional Repository (<https://dspace.lboro.ac.uk/>) by the author and is made available under the following Creative Commons Licence conditions.



CC creative commons
COMMONS DEED

Attribution-NonCommercial-NoDerivs 2.5

You are free:

- to copy, distribute, display, and perform the work

Under the following conditions:

 **Attribution.** You must attribute the work in the manner specified by the author or licensor.

 **Noncommercial.** You may not use this work for commercial purposes.

 **No Derivative Works.** You may not alter, transform, or build upon this work.

- For any reuse or distribution, you must make clear to others the license terms of this work.
- Any of these conditions can be waived if you get permission from the copyright holder.

Your fair use and other rights are in no way affected by the above.

This is a human-readable summary of the [Legal Code \(the full license\)](#).

[Disclaimer](#) 

For the full text of this licence, please go to:
<http://creativecommons.org/licenses/by-nc-nd/2.5/>

An Application of Demand Profiling and Optimisation of Staffing Levels within Leicestershire Police Force

Mr. Oliver S.S.T. Edleston
Dr. Lisa M. Bartlett
Aeronautical and Automotive Engineering Dept.
Loughborough University
Loughborough, LE11 3TU
tosste@lboro.ac.uk, l.m.bartlett@lboro.ac.uk

ABSTRACT:

In modern society there is a necessity for industries to maximise their productivity but also to minimise their outlay, be that in the form of equipment or personnel. Motivated by interest within UK policing towards professionalization of their service, the dual problem of modelling the demands upon front-line Police officers and the optimisation of available resources is investigated. Demand borne from calls to service from the general public is considered and a method whereby this may be realistically quantified in a predictive model is shown. An optimisation technique is described that minimises the number of staff required in order to meet expected demand using a user input series of shift definitions.

A tool is created that contains both the demand profiling and optimisation functions described. The effectiveness of this tool is then shown by application to the shift allocation of Police personnel with results illustrating this provided.

Keywords: Police, Demand, Optimisation

1. INTRODUCTION

Currently for organisations to function not just adequately but most successfully the demand of the service or product supplied must be met. In an organisation or process there are a number of tasks to be performed and a limitation of the resources (machines or man power) available to carry out the tasks. Ultimately a balancing act is required, where a demand modelling approach, to understand the demand, and an optimisation approach to fulfil the demand are required.

Numerous techniques have been applied to the demand modelling problem both within the literature and through industry. These techniques widely consist of the modelling and solution of

problems for a variety of industry specific needs. For example there has been much work conducted in solution of flow shop problems [1], [2]; these being of the type where a set of machines are required to perform a series of tasks, with optimisation based on some objective function, most often the make span (this being the time taken in order to produce a single unit of a desired product). The work of Shumway [3] applies mathematical modelling to the supply and demand of resources within the agricultural industry of Texas; notably the work considers a multiproduct industry. Silver [4] provides an overview of the uses of operations research in the field of inventory management, detailing a series of problems and the practices that have been used in their solution. Meisels & Schaerf [5], in a more pure research application, define a general employee timetabling problem comparing the effectiveness of systematic and local search methods. The work of this paper relates to demand profiling in the Police domain and hence the most directly related problem is the well defined and thoroughly examined nurse rostering problem [6], [7]; this problem being one that also considers demands and the provision of staff within a public service business. The work conducted in this field generally focuses upon solution of pre-constructed staffing problems through use of optimisation and iterative search procedures. For example, the work conducted by Burke et al. [7] employs the tabu search algorithm for nurse roster optimisation whilst considering an objective function based upon staff satisfaction levels. In their work it is shown that tabu search is a highly suitable optimising tool for such applications and it is used in the development of a software package titled PLANE, for use by National Health Service hospitals in Belgium. This paper instead places emphasis on two areas, first of which is the creation of a method by which the demands on

staff members may be quantified; resulting in the creation of demand profiles. Given this representation of demand the second objective focuses on building a fast and suitable optimisation routine that allows for staffing level minimisation.

Though they are both involved with quantification of demands upon staff within a National service sector organisation, there are limits to the parallels seen between nurse rostering and the Police industry. In the health care industry the working abilities of individual staff are important as it is possible for some members to be assigned multiple types of duty; for example a head nurse may work in place of a regular nurse for the duration of a shift [2]. Within the police force however many tasks require specifically trained staff and accordingly the behaviour of covering shifts between differently trained personnel is much more uncommon. Further in Leicestershire, Police staffing levels are currently required to be planned far in advance of the date at which they are used; whereas it is fairly common in nursing to see a much more reactive approach to staffing with major adjustments accepted at short notice. Additionally it is common that demand affecting nursing staff in a hospital will all happen on hospital premises. The incidents that Police staff will respond to however will be spread across their entire policing region and as such the time travelling to and from these incidents is an additional layer of demand not considered in nurse rostering. These are the most notable differences and provide some distinct characteristics, resulting in the definition of a unique demand profiling problem in the Police industry.

The remainder of this paper is broken down as follows. Firstly the problem is described within Section 2, where an overview of relevant areas within the Police business and description of the demands upon staff that the profiler considers are given. Section 3 discusses development of a profiling tool designed for use by Police administration personnel in estimating the demands upon officer time. Section 4 covers the final addition of an optimisation process, where a guide for the minimum staff required to meet all demand using the defined shift pattern may be obtained. The final section of this paper discusses

the conclusions drawn from this work with regards to the suitability and capability of the created demand profiling tool.

2. PROBLEM DESCRIPTION

In the financial year running from 2006 to 2007 Leicestershire Police employed 2261 officers, of which approximately 700 were assigned as Local Police Officers (LPOs). These LPOs are the members of Police staff tasked with handling the majority of calls to service from the general public, in particular those being defined as emergency contacts in which an immediate risk of danger to life or property is present. Thus with this being an area with immediate and noticeable benefit to the general populace there is much emphasis placed on the efficiency with which these officers perform their duties. In order to allow for effective management of available resources it is first necessary to understand and quantify all of the demands that are present upon them.

2.1 Demands upon LPOs

The main way in which time of a Local Police Officer is spent is through the attendance and resolution of incidents; these most commonly being generated through contact from a member of the public reporting a criminal activity. Each of these incidents is manually recorded by Police personnel on the Operational Information System (OIS); every incident being logged through compilation of a series of data recording the time and place of occurrence in addition to important details describing the incident itself. These logs are updated as the incident resolution process progresses to give a full record of the event; importantly the time of each update is noted, allowing quantification of the amount of time spent undergoing each stage of resolution.

Some occurring incidents will result in the arrest of an individual or a group of offenders. In these cases the arresting officers will be those that are required in order to process arrestees through the full custody process. This begins with transport to the nearest available custody suite where, upon arrival, the involved officers and arrestees will be entered into a queue to await further action. The following series of actions can take up to 9 hours to fully resolve and involve recording details for

each arrestee in addition to performing any necessary tests (fingerprint, alcohol, etc.). It is possible that an officers shift may come to an end before the custody process is completed in which case either the officer will remain on duty and be paid overtime or the arrestees will be handed over to another officer to continue.

A common behaviour seen in active policing is double crewing wherein more than one officer will each share the same vehicle; effectively this then results in them acting as a single unit for the duration of their shifts. From a purely numerical standpoint this would appear to be an unwanted practice as it reduces the number of effective resources available for deployment; however it also provides an increase in officer safety, which is of high importance both during night shifts and in known troubled areas.

LPOs are often used to aid in prevention of crime through directed action policing. Most commonly this will occur through either reassurance patrols or officer attendance at planned events; for example a sporting fixture or major local event. This form of demand is highly variable as events in need of such a Police response will themselves vary a great deal; though importantly as it is from response to planned events or from planned Police action, this demand is also easily predicted and thus also easily managed.

A study of LPOs by Richard Ashton, a former technical management manager at the International Association of Chiefs of Police (IACP) Division of State and Provincial Police, has suggested that the average LPO will spend up to 40% of their shift partaking in administrative duties. This is however noted to include meal and toilet breaks in addition to what may be considered active working time such as attending briefings and carrying out paperwork. In addition to this it is concluded that 27% of an LPOs time will be spent travelling; though it is worth noting that this time will be composed for the most part of travel both to and from incident sites. Important to note is that incidents attended by LPOs recorded within the Police data system OIS will include not only the travelling time to an incident but indeed should also contain the majority of administrative paperwork associated with the same incident.

2.2 Summary of Demands the Profiler Represents

A base image of the volume of incidents expected to occur within the average calendar week is formed through consideration of a large amount of OIS incident records. The latest version of ProfInc to date makes use of data from all such recorded incidents in which LPOs were used in response between April 1st 2007 and June 30th 2009. Unfortunately data from earlier than April 1st 2007 was not readily available due to a major restructuring of the policing regions themselves that occurred prior to this date. Each incident record makes note of the time at which the incident is reported (this also being the time at which the incident is first created on OIS) and it is these values that allow an indication of incident occurrence volumes to be generated. In order to do this incident data are collected and segregated into categories based upon their time and day of occurrence, with each category spanning a 15 minute interval. For example then the first category covers the time slot between 00:00:00 and 00:14:59 on a Monday, and thus each incident happening inside this period across the 26 month would be assigned to this category. Upon completion of this then it is possible to obtain an average incident occurrence rate for every 15 minute period by simply dividing by the amount of times that period occurs throughout the entire dataset. So with 118 Mondays across the 26 month dataset the total for the previous period would be divided by 118 to get the average incidents historically occurring within that period. A 15 minute period length was selected for each of the categories as this allows for a suitably high resolution image of demand to be created whilst also ensuring ease of compatibility with the desired software; namely Microsoft Excel from 2003, chosen as a base for the profiling tool as it is a commonplace tool with suitable capability.

To accurately represent a base volume of demand borne from these incidents it is important to also take into account the duration of each incident. Indeed every incident will require the attendance of one or more members of Police staff for some period of time. The actual amount of time required in order to suitably resolve each incident will vary widely depending upon the type and severity of the incident in question. To accommodate this, average resolution times were

calculated for 15 categories each representing the pairing of one of 5 incident types with 3 levels of severity (or grade); an illustration of the types of average resolution times is displayed within Table 1. In this table are 5 types of incident by which all incidents will be categorised and these are Anti-Social Behaviour (ASB), Crime, Public Safety (PS), Road related and Other. As an example if a public safety incident of grade 1 is assumed to require, on average, 41 minutes in order to resolve and thus each incident of this type displayed in ProfInc will be set to be active and requiring officer attendance for this length of time. However considering that all data within ProfInc is segregated into 15 minute categories this will mean that the 41 minutes would have to be rounded up to the nearest multiple of such categories, in this case 45 minutes.

Table 1 Illustrative values for incident resolution times for each pairing of incident type and grade

| | | Incident Grade | | |
|---------------|-------|----------------|----|----|
| | | 1 | 2 | 3 |
| Incident Type | ASB | 33 | 33 | 71 |
| | Crime | 90 | 88 | 92 |
| | Other | 48 | 61 | 71 |
| | PS | 41 | 52 | 68 |
| | Road | 85 | 57 | 63 |

The effects of double crewing are simulated through application of a multiplicative factor that, for each 15 minute time segment, represents the probability of incident attendance by a double or single crewed unit. This factor is described by equation 1.

$$1 + \text{double} / (\text{double} + \text{single}) \quad (1)$$

With 'double' representing the number of double crewed units on duty for a segment of time and 'single' the number of single crewed units on duty in the same segment. Multiplying the amount of demand at any time by the corresponding double crewing factor then gives the effective demand level; this being the amount of staff required in order to meet the given demand using the specified double and single crewing ratio.

The factor of directed action policing was decided to be not strictly necessary for inclusion within ProfInc as part of the expected demand level. This is due to the fact that all such policing is planned for in advance and accordingly all

arrangements for suitable staff will be made well before the event.

The occurrence of custodial incidents is considered in much the same way as the base incident data. A series of average rate of occurrence figures, one for each 15 minute time segment, are obtained through consideration of historical data spanning the same 26 month period. These are then used to estimate the amount of such incidents occurring within each time segment and each such occurring incident is assigned a specific length of time requiring the attendance of an LPO in order to resolve. The average time required over the entire 26 month dataset was seen to be approximately 7 hours; thus for the demand estimates provided in ProfInc custodial incidents were set to take this amount of time to resolve. This is however entirely adjustable by the user in order to allow any length of custody time to be simulated.

A common occurrence in any field of employment is that of abstractions. Though these are not strictly a form of demand upon the time of staff members, they are certainly a factor that needs to be considered in evaluating suitable staffing provisions. An internationally assumed standard for the total amount of abstractions for any given organisation is 30%. For example, in the case that an employer wanted to ensure that 7 staff members were available on a particular shift, it would be advised to schedule 10 and assume that 3 of these would be abstracted. It is important to note that such an abstraction rate is only truly valid when both long and short term abstractions are taken into consideration; organisations scheduling staff at short notice will not be affected during their planning by as many of the longer term abstractions. The work conducted in this paper however has been primarily focussed upon the minimum staffing levels used by Leicestershire Police; these levels being the minimum number of staff members that must be on duty at any given time. Accordingly then abstraction rates play no part in the view of demand as experienced by LPOs.

To summarise, the profiler considers each of the following when constructing a profile of the demand upon LPOs:

- Historical mean incident occurrence rates for the many small time intervals covering the entire period of time of interest.
- Incident resolution time variation based upon both type and grade of incident.
- Effects of double crewing upon the effective amount of available units.
- Administrative duties that occur in addition to that dealing with incidents.
- Occurrence of custodial incidents and the time required in their resolution.

3. CONSTRUCTION OF A DEMAND PROFILE

All the considerations detailed within Section 2 are used in creating a piece of software named ProfInc, which provides a user with the ability to swiftly render the demands upon resources generated from expected incident levels. Further, it allows the consideration of a user defined staffing roster to cover the same period; this allows for immediate comparison to be made between resources and demand, highlighting areas of resource deficiency and surplus. This section details how a representation of demand is created within the profiling tool and will be followed in Section 4 by the description of the optimising tool created for it.

3.1 Generating the Profile

In order to construct a profile of the incident based demands upon LPO staff, each of the pertinent criteria mentioned within Section 2.2 must be quantified. To do this, firstly a base demand value is created for each time interval within the period of interest; this time interval is set to 15 minutes in length. These base demand values are precisely the average amount of incidents seen to occur historically within it, adjusted to take account of the amount of time required to resolve an incident. Illustrative values for use in the examples presented in this paper for the time spent at each incident are shown in Table 1, describing average resolution times broken down by incident severity (or grade) and type. For example, an incident with resolution time of

45 minutes would then be shown as active for the quarter hour interval in which it occurs and also within the following 2 intervals. This is illustrated in Table 2 wherein it is seen that the total occurring incidents for any time interval is the sum of the amount of incidents that initialise within the interval itself (New Incs) and those initialising in the preceding 2 intervals; for example the total incidents at 8.15am (11.59) is the sum of the newly occurring incidents at 7.45am, 8am and 8.15am (3.33, 3.75 and 4.51 respectively).

Custodial incident occurrence rates are then added in a similar fashion and are assigned a resolution time defined by the user. Thus an average historical rate of custodial incident occurrence is calculated for each 15 minute interval within the entire time period and then these are adjusted to account for the time spent resolving the custodial process. This time is by default set at 7 hours, the average resolution time of a custodial incident seen over the past 2 years. All demand is then multiplied by a double crewing factor (as defined in Section 2.2), that effectively represents the chance that a double crewed unit will attend an incident as opposed to a single crewed one, scaling demand appropriately. The effects of both these factors are demonstrated within Table 3 where, for example, the total effective incidents seen at 7.15am (11.91) is given by the sum of the amount of live incidents (Total Incs, 5.94) and the number of custodial incidents (Custody, 2) which is then multiplied by the double crewing factor (1.5).

Table 2 Sample base and total incident demand levels over a 2 hour time period

| | 07:00 | 07:15 | 07:30 | 07:45 | 08:00 | 08:15 | 08:30 | 08:45 |
|------------|-------|-------|-------|-------|-------|-------|-------|-------|
| New Incs | 2.01 | 2.43 | 2.75 | 3.33 | 3.75 | 4.51 | 4.91 | 5.08 |
| Total Incs | 5.01 | 5.94 | 7.19 | 8.51 | 9.83 | 11.59 | 13.17 | 14.50 |

Table 3 Effects of custodial incidents and double crewing in ProfInc

| | 07:00 | 07:15 | 07:30 | 07:45 | 08:00 | 08:15 | 08:30 | 08:45 |
|----------------|-------|-------|-------|-------|-------|-------|-------|-------|
| Total Incs | 5.01 | 5.94 | 7.19 | 8.51 | 9.82 | 11.59 | 13.17 | 14.51 |
| Custody | 2 | 2 | 0 | 0 | 0 | 1 | 1 | 0 |
| Double Crewing | 1.5 | 1.5 | 1.5 | 1 | 1 | 1 | 1 | 1 |
| Effective Incs | 10.51 | 11.91 | 10.79 | 8.51 | 9.82 | 12.59 | 14.17 | 14.51 |

This demand is summarised by the formula shown in Equation 2. For each time interval, t , the total demand, D_t , is equal to the sum of the incident occurrence rate for that interval I_t , the custodial incidents for that interval C_t and the incidents of each type still live from previous time intervals (i_t and c_t respectively) all multiplied by a double crewing factor, F_t . An additional factor, A_t , is included to take account of any user defined additional demands in each time interval; for example this could be known demand from a planned event, or proposed administrative time such as lunch breaks.

$$D_t = (I_t + C_t + i_t + c_t + A_t)F_t \quad (2)$$

Performing this calculation for each time interval will then provide a profile representative of the number of LPO staff members required on duty within each interval in order to meet expected levels of incident based demand. In this way the compilation of demand may be seen to be completely deterministic, with resulting demand levels being calculated based solely on historical and user entered data.

3.2 Output from the Profiler

With a well defined profile of demand, the incident profiler couples the generated profile of demand with a user specified staffing roster which provides a view similar to that illustrated in Figure 1. A shaded region is shown which indicates the amount of officers necessary in order to meet the incident based demand at the corresponding time along the x-axis and a thick line is plotted to represent the staff scheduled at the same points in time.

It can be seen that there are several regions where staffing levels are in excess of the expected demand and indeed some regions where understaffing occurs. The next requirement from

the profiler is to match these two lines, minimising the staffing requirement; hence the use of optimisation techniques are considered.

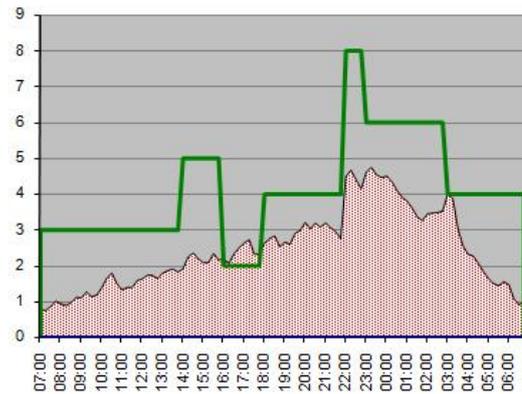


Figure 1 A sample demand profile and staffing roster rendered in ProfInc

4. OPTIMISATION

4.1 Overview

To enhance the capability of the profiler, an optimisation approach was added to make use of the user defined shifts in order to calculate a minimum staffing level such that all demand is met over the entire time period. This program was written using the Visual Basic components included as part of the Microsoft Excel software package; this approach being selected due to the fact that Excel is a freely available software package across the Leicestershire Police organisation. The greatest benefit this provides is the ability to read directly any data that is stored within the workbook and further to automatically take account of any changes that occur to the data. The final optimisation process that was created and is used within the current version of ProfInc is summarised in the flow chart shown in Figure 2.

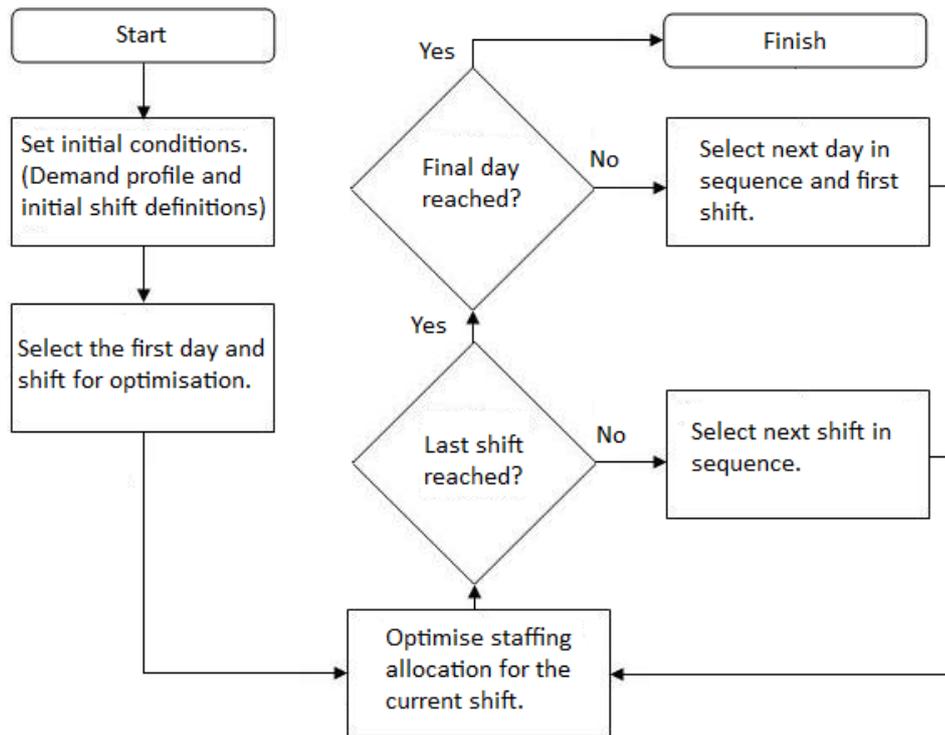


Figure 2 Flow chart describing the optimisation process

4.2. Initialisation

The process works individually upon each day of the week and as such a loop is required such that the full optimisation procedure will occur 7 times accordingly. The first stage within this loop is to initialise all of the user defined shifts to some high value; doing this ensures that later stages will find a better fit to the demand profile, and thus minimise the suggested staffing levels.

4.3 Optimisation

To begin with the latest starting shift within the current day is selected. The high staffing level assigned in the initialisation stage (a value of 99) is now reduced to a reasonable point from where the search may progress. In the code the current staffing levels for the selected shift is firstly set to the value 10, then an if statement is used to decide whether this value should remain at 10 or whether it should be reduced to 1. The reason for this is that for the next stage of the search it is important that there be some demand surplus present at some point in time over the duration of the shift, thus if assigning 10 staff will provide no such surplus then they are reduced to 1. Initialising the if statement, makes reference to a cell within the workbook that contains the highest

demand surplus for the current shift; this is then tested to see if it is negative (implying demand surplus), and if so then staffing levels remain for now at 10. Now for the current shift, the search process will steadily increase the amount of assigned staff until all demand surpluses shown within that shift have been met. This should provide the least amount of officers required for that shift to meet all of the predicted demand for its duration.

Overall the addition of this mechanic to the incident profiler is highly useful as it will allow a user to obtain a series of benchmark staffing levels that will meet predicted demand for any defined pattern of shifts they desire. In this way it is capable of providing at the least a starting point from which staff planners may then work in order to not only meet the predicted incident demand but also to meet other demands upon the Police service.

5. RESULTS FROM OPTIMISATION

Upon completion, the optimisation process will provide a staffing roster based on the user defined series of shifts that is able to meet all expected demand levels with the least number of staff. Figure 3 illustrates the effects of this process

through comparison of a demand profile with 2 separate staffing rosters, one before optimisation and the other after. In this figure, the dashed line is representative of the staffing level provided from the initial roster and the whole line shows that given by the post optimisation roster.

This result is further detailed within Table 4a where the staffing rosters both before and after optimisation are displayed. It can be seen in this example that the initial (base) roster and the resulting (end) roster both contain the same total number of staff. However as Figure 3 demonstrates, the distribution of these staff across the available shifts in the post optimisation roster allows all expected demands to be met; whereas the initial roster suffered from understaffing

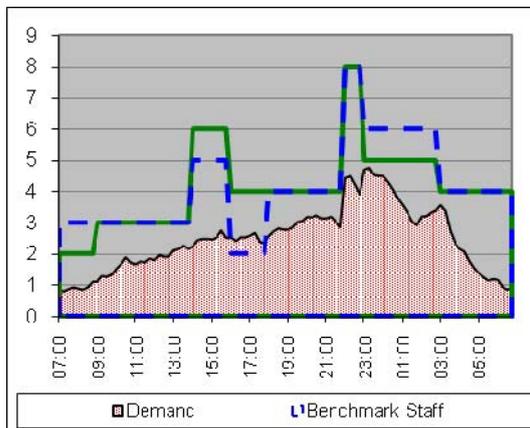


Figure 3 Results of optimisation process applied to a sample problem

between the hours of 16:00 and 18:00. Table 4b shows the results of optimisation upon a problem where the initial staffing roster was seen to have many more staff than were necessary in order to meet the expected levels of incident based demand; indeed the final roster shows all such demand may be met with a reduction of 11 officers.

It is noted that a closer match between the numbers of staff and expected levels of demand could be met through the introduction of additional shift starting times. This would however pose problems due to restrictions imposed upon suitable staff working times within the European Working Time Directive [8].

6. CONCLUSIONS

The software resulting from this work, ProfInc, is well suited to creating realistic estimates of the staffing levels required in order to meet the specified types of incident based demand. Indeed, though this is a relatively simple demand profiling tool, it has been well received and shown to be of use to Police staff in assessing incident based demand. Currently a review is underway within Leicestershire of the uses that have been made of the profiling tool by Police personnel and of the effectiveness and capability of the tool in meeting these uses.

Table 4a & 4b Sample rosters before and after optimisation process

| Base Roster | | End Roster | |
|--------------|-----------|--------------|-----------|
| Start | Staff | Start | Staff |
| 07:00 | 3 | 07:00 | 2 |
| 09:00 | 0 | 09:00 | 1 |
| 14:00 | 2 | 14:00 | 3 |
| 18:00 | 2 | 18:00 | 1 |
| 22:00 | 4 | 22:00 | 4 |
| Total | 11 | Total | 11 |

(a)

| Base Roster | | End Roster | |
|--------------|-----------|--------------|-----------|
| Start | Staff | Start | Staff |
| 07:00 | 16 | 07:00 | 5 |
| 09:00 | 0 | 09:00 | 5 |
| 14:00 | 18 | 14:00 | 6 |
| 18:00 | 0 | 18:00 | 7 |
| 22:00 | 16 | 22:00 | 16 |
| Total | 50 | Total | 39 |

(b)

Coupled with a staffing roster describing actual staffing levels it has been shown possible to begin highlighting problem areas in which it would be expected for demand to regularly outstrip supply. Application of minimisation method has allowed identification of suitable and scientifically supported figures estimating the minimum staffing commitment necessary in order to meet expected levels of incident based demand on a daily basis. However there is still scope for improvement through additional research and development; in particular by considering optimisation that allows adjustment to the shift structure itself in order to find even better uses of available resources. Indeed research is already being conducted to this end with further enhancements from application of more powerful optimisation techniques such as tabu search and genetic algorithms in finding ways to best assign available staff based on estimated demand levels.

Though Excel has proven a reliable and well suited tool for the work presented in this research, to enhance the functionality further more powerful alternatives, in particular application of the potentially computationally demanding tabu search and genetic algorithms, are to be sought. With this added computational flexibility there is potential to consider objectives to make best use of a given maximum amount of staff expected in meeting demand, ultimately considering levels of surplus demand. Indeed work has already been conducted by the authors in investigation of tabu search optimisation algorithms designed to tackle this problem.

ACKNOWLEDGEMENTS

The authors would like to acknowledge the assistance provided by Leicestershire Constabulary throughout this research. Notably the assistance of Ch. Supt. Jason Masters, Sgt. Dave Hall and Mrs Carolyn Steptoe in managing Police aspects of the research and of Mr Jonathan White in the provision of essential data.

REFERENCES

- [1] Reeves, C.R. A Genetic Algorithm for Flowshop Sequencing. *Computers Ops Res.* Vol. 22 (1995) pp. 5-13
- [2] Grabowski, J. & Pempera, J. The Permutation Flow Shop Problem with Blocking. *A Tabu*

Search Approach. *The International Journal of Management Science* Vol. 35 (2007) pp. 302-311

- [3] Shumway, R. Supply, Demand, and Technology in a Multiproduct Industry: Texas Field Crops. *American Journal of Agricultural Economics* Vol. 65 (1983) pp. 748-760
- [4] Silver, E.A. Operations Research in Inventory Management: A Review and Critique. *Operations Research* Vol. 29 (1981) pp. 628-645
- [5] Meisels, A. & Schaerf, A. A Modelling and Solving Employee Timetabling Problems. *Annals of Mathematical and Artificial Intelligence* Vol. 39 (2003) pp. 41-59
- [6] Burke, E.K. Causmaecker, P.D. Berghe, G.V. & Landeghem, H.V The State of the Art of Nurse Rostering. *Journal of Scheduling* Vol. 7 (2004) pp. 441-499
- [7] Burke, E.K. Causmaecker, P.D. & Berghe, G.V. A Hybrid Tabu Search Algorithm for the Nurse Rostering Problem. *Lecture Notes in Computer Science* Vol. 1585 (1999) pp. 187-194
- [8] Working Time Directive of the European Union. Council Directive 93/104/EC. 23-11-1993

AUTHOR BIOGRAPHIES

OLIVER EDLESTON received an MMath (Hons) Mathematics from Loughborough University in 2007. He is currently a PhD research student in the Department of Aeronautical and Automotive Engineering at Loughborough University.

LISA BARTLETT is a lecturer within the Risk and Reliability research group in the department of Aeronautical and Automotive Engineering at Loughborough University. Her research interests are in reliability assessment methods, demand modelling and route planning.