Relief Staff Rostering for the St John Ambulance Service

Hala Jasim
Department of Engineering Science
University of Auckland
New Zealand
jasims@ihug.co.nz

Abstract

In an organisation it is important to manage the work periods when staff members are on leave. Within St John, full time staff might take holiday leave or outside training at a known time. The shifts that the full time staff would have covered have to be allocated to other staff. Relief staff is the term given to the staff members who cover the planned leave that is being taken by other staff within the organisation.

In this fourth year project, a computerised system has been developed that makes use of optimisation techniques that allow for the efficient allocation of staff to roster periods while meeting shift requirements and staff objectives. The rostering problem is modelled as a set partitioning problem, which is a zero-one integer-programming problem. The optimisation methods used to solve this problem are the Revised Simplex Method and Branch and Bound.

A newly developed fatigue model is applied to the optimal solution, which is used to predict the impact of working hours on the fatigue that is experienced by the staff member. The fatigue model will assist in the design of effective rostering schedules that take into account the social, domestic and personal needs of employees.

1 Introduction

1.1 The St John Ambulance Service

St John provides nearly 90 per cent of ambulance services in New Zealand and each year it saves the lives of thousands of New Zealanders and visitors to the country [4].

The ambulance service employs 600 salaried ambulance officers and a further 2200 volunteers. Within the Auckland region, the service employs and rosters 214 ambulance officers and 110 voluntary staff. The organisation responds to 1540 calls per week or about 80,000 calls a year from within the Auckland region.

The St John Ambulance Service is committed to offering the highest possible standard of service to its patients. To help with the management of their resources, St John is working with the Operations Research Group in the Department of Engineering Science, where a simulation model of St John's operations in Auckland has been developed. This is part of a greater Better Ambulance Rostering Technology (BART) project that examines the problems of ambulance location and staff rostering.

1.2 The Project

Within St John full time staff might take holiday leave or outside training at a time known in advance. The shifts that the full time staff would have covered have to be allocated to other staff, termed relief staff.

The aim of this project is to investigate and commence the development of a computerised system that will assist with the rostering of relief staff. This project will build a general system that uses optimisation techniques to achieve the optimal allocation.

The use of optimisation techniques allows for efficient allocation of staff to shifts while meeting business requirements and staff objectives. The rostering problem is modelled as a set partitioning problem, which is a zero-one integer-programming problem. The optimisation methods used to solve this are the Revised Simplex Method and Branch and Bound.

This paper initially defines the problem and gives a brief description of how the shift data is extracted from weekly roster sheets. Later sections of the paper discuss the methods used in building the system and the techniques that are applied in generating the optimal allocation of staff to roster periods.

2 Problem Definition

2.1 Roster Periods

It is important to understand how the shift information is extracted from existing weekly rosters. The example below (Figure 1) demonstrates how the shifts that need to be allocated to relief staff are extracted from roster sheets provided by St John.

Staff

Line of Work

	Day 1	*Day 2*	*Day 3*	*Day 4*	*Day 5*	*Day 6*	*Day 7*
1	Day Albny	Day Albny		Night Central	Night Albny		Day Albny
2	Day Pkura		Day Mnkau	Day Pkura	Night Mnkau		Day Pkura
3		Night Nlynn	Night GI		Day Nlynn	Day Nlynn	Night Hwick

Figure 1. Example: - Line of Work of 3 full-time staff members.

Figure 1 above shows the sequence of shifts and days off – termed a line of work – for three full-time staff members. Each row of the table shows the line of work for one staff member. The lines of work shown are for 7 days. Empty cells indicate mandatory days off and the highlighted cells are the shifts (or roster periods) that the staff member has reported in advance that he/she will be absent for, and therefore a relief staff will have to be allocated to that roster period. From this table, a list can be made of all the roster periods that need to be allocated to relief staff.

For this example, the shifts to be allocated are given below. Each of these shifts has an associated start time, finish time and a location at one of the ambulance bases in Auckland. In addition, each shift has associated with it one or more skills that must be possessed by a staff member.

1. A day shift in Albany
2. A day shift in Albany
3. A day shift in Papakura
4. A night shift in New Lynn
5. A night shift in Glen Innes

The staff member that needs to be allocated to these shifts must meet the skill requirements of that particular shift. The requirements of the staff member have to be considered in the allocation process. Each staff member has preferences on the locations he/she might want to work. There are also limits on the minimum and maximum hours the staff member can work in a given work period. Meeting both staff and shift requirements is an integral part of the allocation process

2.2 The Model

In building a model, the available roster periods are used to generate possible lines of work (sequences of shifts and days off) for each staff member. The problem involves the assignment of a line of work to each staff member, with an objective to maximise the total roster quality of the allocations.

The model that is used to represent this rostering problem is the set partitioning problem (SPP), which is a zero-one integer-programming problem.

The SPP is formulated as:

$$\begin{aligned} &\min && \boldsymbol{c}^T \boldsymbol{x} \\ &\text{st} && A_s\boldsymbol{x} = \boldsymbol{e} \\ & && A_c\boldsymbol{x} = \boldsymbol{b} \\ & && x_j = \{0,1\} \qquad j=1,\ldots, n_L \qquad n_L = \text{number of lines of work} \end{aligned}$$

Where $\boldsymbol{c}$ is the cost vector
$\boldsymbol{x}$ is the solution vector
A_s is an n_s by n_L constraint matrix, n_s = number of staff
A_c is an m by n_L a constraint matrix, m= number of shifts
$\boldsymbol{e}$ is a unit vector of ones
$\boldsymbol{b}$ is the right hand side vector

In this model, up to one shift per day can be allocated to each staff member. The first n_s constraints of the A matrix (given by A_s) represent the staff members and require each staff member to be assigned exactly one line of work (therefore the right hand side equals 1 for the first n_s rows of the A matrix). The remaining constraints (given by A_c) represent the number of staff required to work a particular shift. For this project, it was assumed that only one staff is required to work any given shift. Hence the right hand side vector $\boldsymbol{b}$ is unity. Figure 2 below shows the structure of the A matrix. Each column of the matrix represents a feasible line of work for a staff member. The cost $\boldsymbol{c}$ represents the quality of the line of work and is dependent on factors that reflect the real problem. A high cost indicates low quality for a line of work, hence our objective is to minimise the total cost of allocations.

	Staff 1				Staff 2				Staff 3				….	Staff n_s-1				Staff n_s				RHS $\boldsymbol{e}$ / $\boldsymbol{b}$
Staff 1	1	1	1	1																		1
Staff 2					1	1	1	1														1
:													: : :									1
Staff n_s-1														1	1	1	1					1
Staff n_s																		1	1	1	1	1
Shift 1	1			1	1		1				1			1			1			1	1	1
Shift 2		1			1			1			1				1			1			1	1
:			1	1	1		1	1	1			1		1		1		1				1
Shift m - 1	1	1		1	1		1	1		1	1			1	1		1	1		1	1	1
Shift m			1			1			1			1				1			1			1

Figure 2. Example: - A matrix structure for a general set partitioning problem.

3 Solution Process

3.1 Shift and Staff Information

As mentioned in section 2, shift information is gathered from existing St John roster sheets. This information contains the shift start time and finish time, the shift base location and the skill level that the shift requires the staff member to possess. The staff members have to specify their preference for the possible base locations, and the minimum and maximum hours they are willing to work in a given period. In addition, the skill level that the staff member possesses is recorded. The following section describes how this information is used to construct the columns of the *A* matrix.

3.2 Column Construction

3.2.1 Generating Lines of Work

We start initially by listing all the possible shifts a staff member can work. This was achieved by first considering the skill level of the staff member and ensuring it matched the skill required by the shift. Each staff member has a preference between 0 and 1 for all the possible base locations in which the shifts can be located. The staff member has to have a high base preference – greater than 50% or 0.5 - in order to work any shifts at that base.

Binary operators were used to match the skill of the staff member with the skill required by the shift. A simple algorithm was written to filter out all the shifts that a staff member is under-skilled for and only list the shifts that a staff member is able to work. The list of shifts had to be enumerated to generate all possible combinations of shifts and days off that a staff member can have in one week.

Staff members could have millions of possible lines of work depending on their skills and flexibility in working at different locations. These possible lines of work however need to be reduced, as not all combinations represent acceptable lines of work. Further constraints have to be applied in order to solve the problem for the legal lines of work.

3.2.2 Output Reduction

The legal lines of work are those that reflect the rules and conditions of the real problem. Each staff member has minimum and maximum hours they can work. These limits have to be met in order to consider that line of work as legal. In addition, a specified minimum time off between shifts has to be met. It is not desirable for any staff member to work a shift that finishes at say 11 pm and work a shift that starts at 1 am the next morning.

Adding these constraints to test and possibly discard illegal columns during the column construction process reduced the size of the output dramatically. Using the final lines of work generated for each staff, the next stage is to calculate the costs for the legal lines of work.

3.2.3 Cost Calculation

The cost associated with each legal line of work reflects the quality of that line of work. The rules that were followed in order to apply costs to the lines of work are:

Sum of base location preferences

> As mentioned earlier each staff member has an associated preference for each of the bases a shift might be located at. A sum of all the preferences for the locations of the shifts in a line of work gave a portion of the total cost.

Change of base location penalty

> A penalty cost is applied whenever there is a change in the base location of shifts from one day to the next. Staff members prefer to work in the same location, but when they have to travel from one location to the next in order to work the next shift, a penalty has to be applied.

Start time change penalty

> A penalty is applied if there is a change in the start time of shifts on adjacent days. The penalty is increased if the shift starts earlier than that on the previous day. This reflects good rostering practice.

Day off penalty

> It is normal in a line of work to have two consecutive days off. However a penalty is applied if the staff member is allocated one day off instead of two days off consecutively.

The total cost of a line of work is the sum of all the costs mentioned above:
Preference cost + base penalty cost + start time penalty + day off penalty cost.

3.3 Solution Techniques

Once the problem had been set up, as described above, it was solved using the ZIP 4.0 (Zero-One Integer Programming) package, developed by Professor David M. Ryan. The ZIP package implements routines that utilise the Revised Simplex Method (both primal and dual) and Branch and Bound to solve the problems of set partitioning and set covering. In order for the solver to find an optimal solution the user has to input the data in an appropriate format that is recognised by the ZIP package. ZIP also allows the user to reset some of the default options to suit the problem being solved.

4 Graphical Output

Once ZIP is run with the appropriate user data, the optimal solution is output in a text format. A user-friendly interface was set up to read and then graphically display the optimal lines of work for each staff member.

4.1 Graphical Display of Rosters

The ZIP output provides the optimal lines of work for each staff member by listing all the shifts allocated to each individual. The information required to develop the graphical interface are the optimal allocations produced by ZIP and the individual shift properties. Shift properties include shift start time and finish time as well as the base location. This information can be displayed on a picture box on Visual Basic forms. A snapshot of this output is given below.

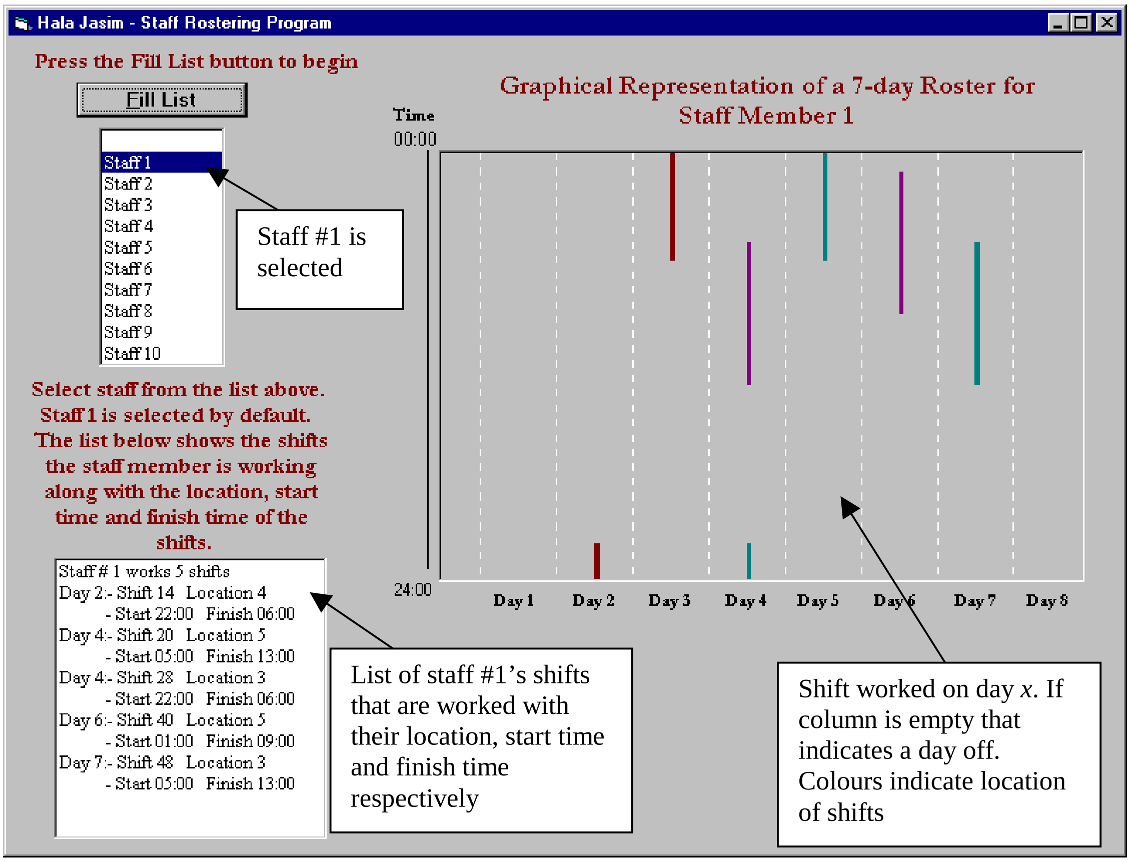

Figure 3. Visual Basic window showing the roster in graphical format.

This interface helps in visualising rosters and days off for each staff member. The second part of the interface is the application of a fatigue model to the optimal rosters and producing plots that represent the fatigue levels of the staff member when a shift is worked. The following section discusses the fatigue model in further detail.

4.2 Roster Fatigue

4.2.1 Background on Shiftwork Fatigue

A new development in rostering is the roster fatigue calculation, which has been developed by the Centre for Applied Behavioural Science (CABS). CABS is an Australian leader in shiftwork and sleep research. The centre is equipped with state of the art behavioural laboratories through which research programs are conducted.

Hours of paid work influence many aspects of an individual's life [3]. They determine a worker's social activities, financial status, family life, geographical location, health and community interaction. The consequences of non-standard working arrangements are reflected in the absenteeism rates of workers, accident and injury rates, decrease in working performance and health related problems.

4.2.2 Fatigue Modelling

Depending on the risks associated with a particular job or task, the impacts of sleepiness, alertness and fatigue can be severe. Models can be used to predict the impact of working hours on sleep, alertness, performance and fatigue.

Fatigue and accident risks have traditionally been viewed as being directly related to the number of hours that an individual has worked. However, the fatigue value of a work period varies as a function of both shift duration and the time of day covered by the shift. Indeed the time of day is as important as the number of hours an individual works [3].

Twenty-four hour cycles, or circadian rhythms determine a wide range of factors such as our drive to wake up during the day and go to sleep at night. Circadian rhythms control a variety of measures related to fatigue including mental efficiency, computational performance and logical reasoning.

Due to circadian rhythms the performance capacity is at the lowest during the early hours of the morning (approximately 2am to 7am) and is also reduced in the early afternoon (approximately 2pm to 4pm) [3].

Mathematical Description of the Fatigue Model

The underlying fatigue of the model is the measure of how much fatigue is accumulated in hour *i* of the day. The fatigue model described by Dawson [3] uses simple linear weighted averages of past fatigue levels for the underlying fatigue. The weights are determined by the time of day of work and breaks, duration of work and breaks, work history in the preceding 7 days and the biological limits on recovery and sleep.

The model used to determine the fatigue levels in this project is given below.

Let

$$\text{Circadian}(i) = \text{circadian factor for hour } i,\ i = 0,\ldots,23$$

$$\text{WorkFatigue}(i) = \text{working fatigue in hour } i \text{, defined by}$$

$$= \begin{cases} \text{Circadian}(i) & \text{if staff is working hour } i \\ 0 & \text{otherwise} \end{cases}$$

$$\text{UnderlyingFatigue}(i+1) = 0.98 * \text{UnderlyingFatigue}(i) + 0.02 * \text{WorkFatigue}(i)$$

$$\text{ObservedFatigue}(i) = \text{UnderlyingFatigue}(i) * \text{Circadian}(i)$$

This model is different to the one given by Dawson [3] in that weights decrease in a non-linear fashion with past fatigue levels. As in Dawson’s model, fatigue accumulates differently for different hours of the day. Accumulation is greater if working across the night and early in the morning than across the day and afternoon. For the purposes of this task the circadian factor was assumed to take on an approximately sinusoidal behaviour. The sinusoidal behaviour of the circadian fatigue factors indicates a smooth pattern of fatigue throughout the day.

The graph (Figure 4) below shows how the fatigue factor fluctuates between hours of the day.

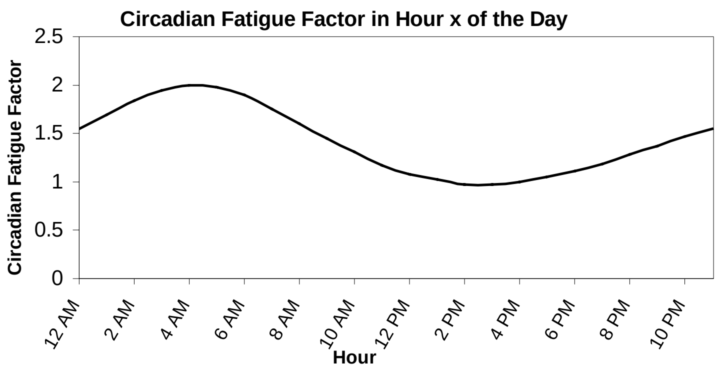

Figure 4. Circadian fatigue factor in hour x of the day

The person's working times determine the observed fatigue profile. Under normal conditions a person with no shifts worked would experience the same fatigue levels throughout the week. Figure 5 below shows the fatigue profile for a staff member who is not working any shifts. The figure shows the constant sinusoidal pattern of the fatigue throughout the week. Figures 6 and 7 indicate that there are disruptions to the fatigue profile when shifts are worked during the week. The fatigue experienced is higher when night shifts are worked as shown in figure 7 due to the high circadian fatigue when working during the hours of the night.

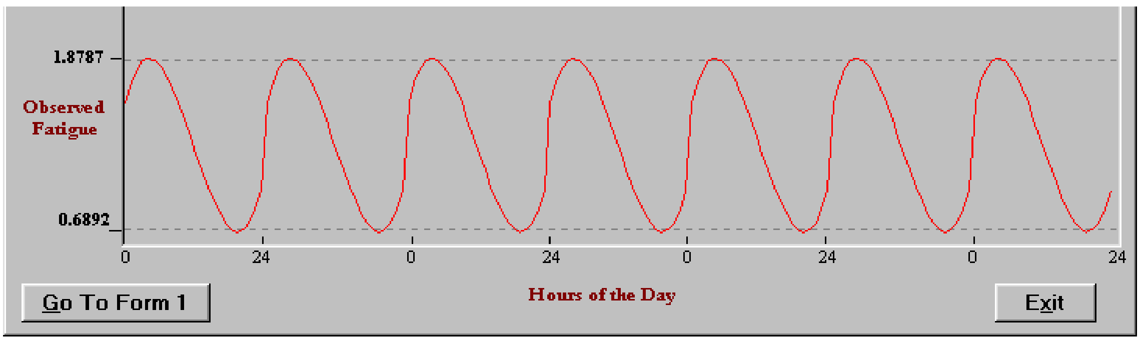

Figure 5. Fatigue plot for a staff member with no shifts worked in a roster week

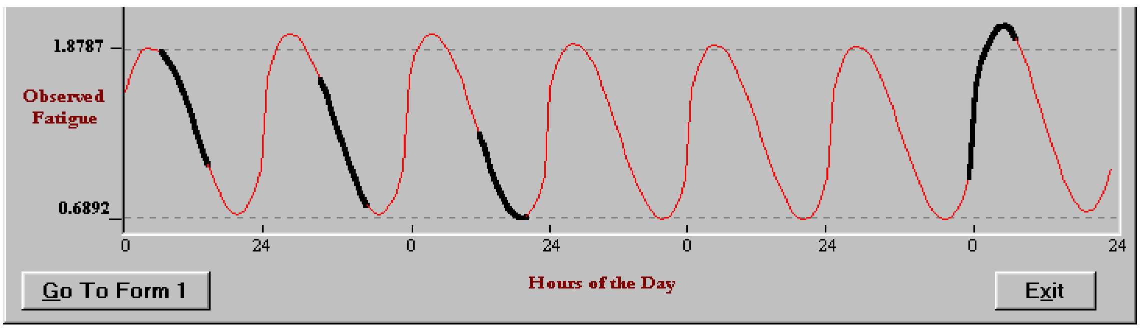

Figure 6. Fatigue plot for a staff member who is working 4 day-shifts in a roster week (shifts worked are shown by bold lines)

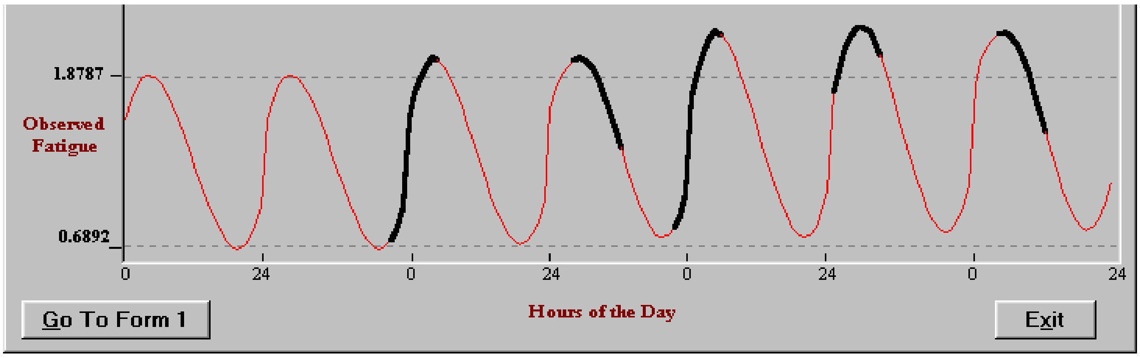

Figure 7. Fatigue plot for a staff member who is working 5 night-shifts in a roster week (shifts worked are shown by bold lines)

Fatigue studies provide employees and managers with information that will assist in the design of effective rostering schedules that take into account the social, domestic and personal needs of employees.

5 Conclusions

In order to achieve the optimum allocation of staff members to roster periods it is necessary to meet both staff and shift requirements.

An optimisation system has been developed for rostering relief staff that generated over a million legal lines of work for all the staff members considered. The legal lines of work that were produced by the system were solved to optimality using ZIP 4.0 (Zero-One Integer Programming) package. The solution produced was effective in rostering all available staff with maximum quality.

A graphical user interface has been developed that is helpful in displaying weekly rosters. The staff member can easily see the distribution of days off and shifts worked over the week. The interface also displays the shift's location, start time and finish time.

Fatigue modelling presents a representation of the fatigue experienced in a week, given a staff member is working a number of shifts. A new fatigue model has been developed and implemented on the optimal solution. Ongoing work will examine how this model should be tuned to match observed fatigue for an individual, which can then be combined with the optimisation system to improve the overall quality of the rosters being generated.

Acknowledgements

I would first like to thank my supervisor Dr Andrew Mason for his continued feedback and flow of new ideas for this fourth year project. It has been a pleasure working with a supervisor who knows so much about his field.

I would also like to thank my friends, some of whom have helped me with the programming problems I was facing.

Last but not least, my family for putting up with the late nights and offering support and showing interest in my progress throughout the project.

References

[1] Lupis, N. (1996) “Maintenance, Operating, Rostering and Optimisation”, A year Four Project, Department of Engineering Science, School of Engineering.

[2] Smith, J.J. “Estimating Ambulance Requirements for the St John Ambulance Service”, Year 4 – Industrial Mathematics Project.

[3] *The Centre for Sleep Research* (2001, March – Last update) [Homepage of The University of South Australia] [Online]. Available: http://www.unisa.edu.au/sleep [2001, September 23].

[4] *The Order Of St John In New Zealand* [Homepage of The St John Ambulance Service in New Zealand] [Online]. Available: http://www.stjohn.org.nz [2001, September 7].

[5] Thompson, R. (1996) “Rostering at the Sheraton”, A year Four Project, Department of Engineering Science, School of Engineering.

[6] *Articles / Press Cuttings* [Homepage of Shiftwork Services], [Online]. Available: http://www.shiftwork.co.nz/ [2001, September 30].