Maintenance Operation Centre Rostering Problem

Dominic Fok
Department of Engineering Science Part IV Project 2010
University of Auckland
In affiliation with Air New Zealand
New Zealand
cfok004@aucklanduni.ac.nz

Abstract

This report details the work done on the Maintenance Operations Centre (MOC) Rostering Solution Project. MOC looks after the short-term maintenance of Air New Zealand and other airlines. In order to provide sufficient service levels MOC has specific daily staff requirements. Each staff member has a set of skills that apply to specific aircraft systems. The objective of the project is to produce a method of generating high quality rosters for MOC that meet daily staffing requirements and maximize the skills available during any time period.

The current method MOC uses to generate their roster is by hand. The problem is characterized as a Rostering Problem. Information on how the objective and staffing requirements are placed into the Rostering Problem is given.

The modeling language used for this project is AMPL, which in conjunction with CPLEX is used to solve the problem. The solution rosters generated provide rosters that meet the daily staffing requirements and provide improved skill and service coverage. However, the rosters cannot be currently used in the operation of MOC, as they do not allocate the shifts fairly. It is expected that equal shift allocation will still allow high quality rosters to be generated.

There are also plenty of areas where additional improvements to the model can be made to improve solution quality. The most prominent improvement would be the addition of employee holiday inputs. This would allow the model to consider the effects of an employee going on leave and adjust solutions accordingly.

Once the equal allocation of shifts is achieved and employee vacation information is implemented, this model will be able to produce solutions that MOC can use. The solutions used will provide higher skill and shift coverage during MOC's operation.

Key words: MOC, skills, roster, shifts

1 Introduction

1.1 Background

MOC is operated by Air New Zealand. It provides short-term and response maintenance for Air New Zealand aircraft, as well as aircraft of other airlines which have contracts with Air New Zealand. An example where MOC would come into action would be: an inbound flight notices its wing has a slight malfunction. MOC would arrange for its repair before the aircraft would fly out again. MOC operates on a 24/7 basis to provide service.

MOC has three different staff tiers to facilitate its operation. The first tier is made up of work crews, who do the actual maintenance work required. The second tier is made up of sub-managers, who direct the work crews and provide aircraft specific technical expertise. The third tier consists of the managers, they ensure the daily operation of the MOC goes unhindered and also provide technical expertise when necessary. MOC has asked that only the manager and sub-manager tiers be investigated for this project.

1.2 Shifts

In order to provide adequate services, the daily operation of MOC has specific staffing requirements. To meet these requirements, MOC uses ten different shift types. In addition to the on-duty shifts available, a staff member can be placed on a call shift. While on a call shift, staff members are not required to be at the work place, but must be available for contact and able to work if asked to. On call employees are only called in when a currently or soon to be on-duty employee is unable to work. Managers and sub-managers work different shift types. Each shift type is required to have an employee assigned to that shift each day. The combination of shift allocations in a roster for an employee is called a work-line.

1.5 Skills

MOC services several different types of aircraft. Each of these aircraft is different and has its own technical systems. Because of this, employees need to be certified that they are able to work on a particular aircraft system. This certification means each employee has a skill set that allows them to work on specific aircraft systems. For example, if a Boeing 747 has mechanical issues, an employee with Boeing 747 Mechanical certification is needed to oversee its repair. Due to the operation requirements of MOC, an employee must be on location to provide a skill, this means employees on call shifts must be ignored when looking at the available skills for a time period.

1.6 Project Objective

The objective of this project is to produce a method to generate rosters for MOC employees that meet all daily shift requirements as well as any other requirements and maximize the skills available for each hour of the day.

2 Method

2.1 Restrictions

In generating a roster, there are rules on the shift allocations allowed for an individual. Some examples of these rules include; no more than 60 hours of work in a week, no more than 6 consecutive days with night shifts, shift allocation must obey Circadian Rhythm and others. The inclusion of these rules in roster generation reduces the number of variables in the problem. However, the inclusion of these rules alone is insufficient to limit the number of possibilities to a level where it can be solved. Thus, further steps are taken to reduce the size of the problem.

A set roster length of 27 days is introduced. Study of the current roster shows a prevalence of 6 days on-duty and 3 days off-duty for employees. Thus, the 27 days is made up of 3 9-day blocks. Each of these 9-day blocks may be chosen from a list of preset shift combinations. One example is 3 morning shifts, followed by 3 afternoon shifts and then 3days off. The exact composition of these 9-day blocks was chosen so that the number of 9-day blocks is as small as possible, while ensuring feasible solution can still be found. The composition of these 9-day blocks follows the rules specified above and not in this paper, as well as some introduced to reduce the number of combinations. As a result there are only 1500 possible choices for an employee when deciding their shift allocations over a 27 day period.

2.2 Cost of Skill Slack Variables

As these are the only variables with a coefficient in the objective function, the values for the cost of a lack of skill for a time period must be carefully scrutinized. The minimization of the presence of active slack variables is what drives the optimization problem. For the MOC problem, each skill corresponds to an aircraft. By looking at the fleet of aircraft Air New Zealand has, a weight can be associated with each skill, giving a cost. The air fleet compositions of airlines other than Air New Zealand has been ignored for this project, as Air New Zealand is the primary customer and there is insufficient information on other airlines. The skill weight is based on the number of aircraft that requires that skill.

2.4 History

As 27 days is an insufficient length of time for the roster it is necessary to include past information. By considering the last 6 days of a roster, this information can be used to generate constraints for the first 6 days of the next roster, so the connection of one roster to another is seamless. An implication of this is that a complete roster of any length can be generated. The reason only 6 days of data is required is that only 6 consecutive days need to be defined to choose a work line without breaking any rules.

2.5 Restrictions on Number of Shifts for an Employee

Because of the different skills each employee has, it is possible that an employee with less valuable skills may be repeatedly assigned call shifts and a employee with valuable skills on duty shifts. As well as being unequal in treatment, the annualization of hours means there must a method of ensuring employees do an equal amount of the different shift types. This is a requirement by MOC and must be added as a constraint. Currently, a limitation method is used to force equality. The limit is designated individually for each shift type and each employee. If an employee has more than their fair allocation of a shift, they will be limited in how many times that shift can appear during roster generation. This will force the shift to be allocated to other employees. The restriction method assigns each employee a value for each shift type.

The 27-day work line assignments may not assign shifts to employees that cause this value to be exceeded. The current method of deciding the restriction numbers is arbitrarily done by hand as an effective algorithm has not been investigated. In general, if an employee has had more than the average number of a particular shift type allocated to them, they will have a lower limit that will force them to take other shifts. After each 27 day step has been generated and the solution generated so far has been analyzed, the results of the analysis decide the restriction limits to be applied to the next 27 day step.

2.6 Formulation and Application

The MOC problem is formulated as a Rostering problem. The constraints have been expanded to include skills as well as staffing requirements. For each of these skill constraints a slack variable has been introduced, which will become active if that skill is not covered during that time period. Some of the constraints also have inequalities rather than equalities, this is because there is no penalty with the same skill being available multiple times during a time period. Other elements of the constraint vector may be larger than one, such as the requirement of two sub-managers on a night shift per day.

Positive costs are attached to the slack variables, so the solver will try to minimize the number of periods where a skill is not available. The cost for work lines is zero as this project currently ignores work line specific employee preferences. This is due to the annualization of hours, where employees receive an equal salary regardless of the exact number of shifts worked. To maintain equality of shifts the solver will try to assign shifts evenly to each person, this will be detailed later. An example of the MOC problem cost vector is shown in Figure 2.

Each column of the A matrix corresponds to either the skill and staffing requirement effects of an employee being assigned that roster, or to a slack variable for a skill during a specific day and time period. Where the roster would meet or help satisfy the staffing/skill requirement for that day and day period, the column has an entry of one. If it does not contribute to that day's staffing or skill requirement, the entry would be zero. The final part of the model is the binary variable vector x. The binary requirement is necessary, as multiples of the same employee is not possible. Furthermore, the constraints corresponding to the slack variables only require one instance to meet that constraint, so multiples of slack variables would not occur. Should a particular work line for an employee be chosen, the x variable will be 1 in the corresponding position and similarly for the slack variables. The combination of the c, b, x vectors and the A matrix allow the problem to be solved, and to be passed to a solver.

3 Results

Results for two different model states are to be shown. This is to compare how solution effectiveness changes according to whether or not there is fair allocation of shifts.

3.1 No Shift Restrictions

For this model state the slack variable weights are included, but not the shift restriction. This is to see the effect of adding weights to the skills. The objective function value is calculated by summing the coverage ratios for each skill over all periods and then multiplying that value by the skill weighting, these new values are then added together to get the objective function. A solution was obtained for the weighted roster problem. An analysis of the original roster and Solution Roster 1 was performed. The objective function for the original roster is 740.68, for Solution Roster 1 874.88.

3.2 Full Model

For this model state the full model is used by including shift restriction. This is to see the effect of restricting number of specific shift types an employee may do. The objective function value is calculated as before. A comparison is also made with rosters without the restriction constraints for both objective function and fairness of allocation.

A solution was obtained for the restricted roster problem. An analysis of the original roster, Solution Roster 1 and Solution Roster 2 was performed. The objective function for the original roster is 740.68, for Solution Roster 1 it is 874.88 and for Solution Roster 2 it is 869.58. A comparison of their average standard deviations and the average difference of the maximum and mean for all shift type allocations can be found in Figure 1.

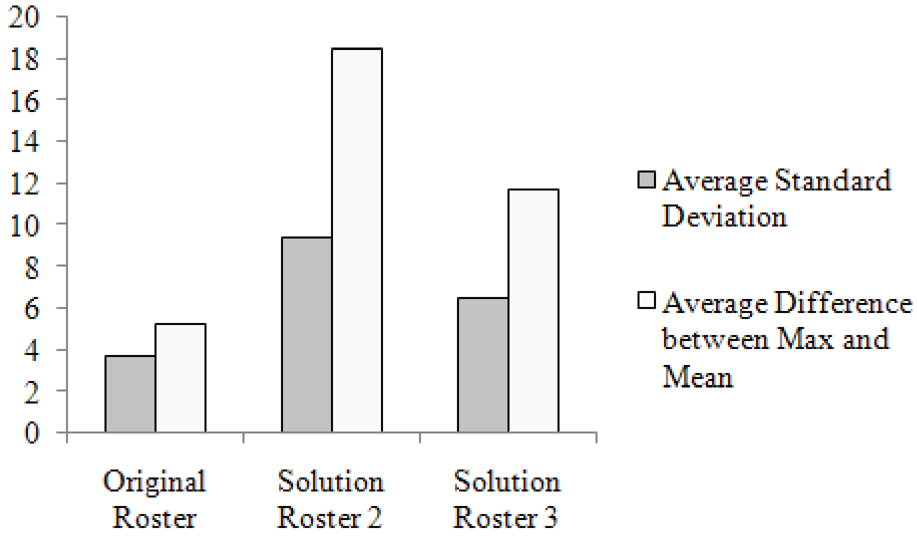

Figure 1 Comparison of Roster Occurrence Statistics

4 Discussion

It is clear that the use of a Rostering Problem optimization produces solutions that have a clear advantage over the original roster for the maximization of skills in a work period. However, the solutions produced so far have issues with the equality of shift assignments although, the shift restriction area of this problem has not been investigated in depth. It is expected that a solution that provides equality as well as high skill coverage is will be found. The issue of shift allocation equality needs to be addressed before the generation of solutions can be used in operation of MOC.

Another advantage of the solutions, is that one of the call shifts that was rarely filled in the original roster, is available single and even twice daily in the solution rosters. This extra coverage allows MOC some flexibility in dealing with the call shift, as they have an option allow a staff member to instead use the extra call shift when two people are assigned on the same day as an off day.

5 Limitations and Future Work

The current formulation has several limitations, one such is the lack of employee input. Employees may prefer have a say as to when they assigned certain shifts. Currently, it is possible to address this by explicitly adding employee requests as part of the history constraints, although if too many employees have requests this can result in an infeasible problem.

The formulation also ignores the length of each time period. This can be introduced such that each time period has a weighting equal to its length of time. Alternatively, data about the aircraft, or the peak arrival/departure times could be placed in the model to give the time periods weights.

The model currently does not have a method to consider aircraft flight information, such data could allow the model to prioritize skill coverage to certain time periods.

The optimization does not consider the possibility that MOC may prefer a skill be available at least twice in a time period. However, it is reasonably simple to adjust the constraint vector b to allow for such a requirement.

The model takes not take into account that employees may wish to take leave or require training. If training was only for a block of three days, it would be possible to force the solution generated to give that employee a day block where there is already another, but such a method would not be sufficient for employees going on three weeks holiday.

Currently as well as the problem of equalization of shifts, the issue of weekend shifts is also ignored. This should also be investigated and put into the model.

As well as the model limitations there are also the limitations of computing power and memory. During the first programming stages for the model, it was found that the problem could become too large to solve, due to the presence of numerous variables. The introduction of history allows the generation of smaller time steps, but requires human oversight to produce a complete solution. The production of solutions can also take a considerable amount of time. The time taken to produce an initial solution tends to be quite high, but also variable. Solutions have been known to take from as little as one minute to over twelve hours.

6 Conclusion

The application of the Rostering Problem to the operation of MOC can provide improved skill coverage and more shift coverage that the original roster. However, it is currently unable to produce satisfactory solutions due to a poor method of shift restriction to equalize shift allocations. Investigations into an appropriate algorithm for determine the restriction limits needs to be done, but it is expected a good equalizing algorithm will still produce improved solutions.

In addition to the issue of equalization, the model needs to allow for employees who wish to go on holiday. There is no current method in the model to take in such information. A method must be found that can respond to employees going on vacation before this can be replace the current method that is used by MOC.

The model of the problem has a lot of potential to provide an improved roster to MOC. If sufficient data is given, the model can be modified to take this information and use it to generate a solution that allows for the new data set. This could be in the form of known arrival/departure times, the aircraft that are used for the arrivals and flights, or otherwise.

Acknowledgements

I would like to thank my supervisors Professor David Ryan and Doctor Andrew Mason for the guidance they have provided me for this project. David was a huge help and inspiration, his knowledge and experience of the Rostering Problem was invaluable. Andrew also provided a lot of guidance during the length of the project and his own knowledge of rosters was also a great help. I would also like to thank Paul Keating, the

Air New Zealand liaison for helping me in this project, his enthusiasm for Operations Research was inspiring and was very helpful in obtaining useful data. I also thank Doctor Michael O'Sullivan Jr for the advice he gave me on this project with regard to CPLEX memory issues. I would also thank Carl Ho, whose previous project on rosters gave me some starting steps.

References

Ryan, D. 2010. 762 Set Partitioning Problem Notes. University of Auckland.