Call Repair for Long-Haul Flight Attendants

Amanda Day & Rochelle Meehan
Air New Zealand
New Zealand
amanda.day@airnz.co.nz
rochelle.meehan@airnz.co.nz

Abstract

An important part of any airline scheduling problem is "on the day" rescheduling. This occurs when the scheduled roster cannot be operated. This may be caused by mechanical problems on an aircraft, crew sickness, bad weather etc. To allow for disruptions to the scheduled roster, some crew members are designated as "on call" at the rostering stage. These crew members are not rostered any work, but are required to report to the airport at short notice in case of a disruption.

This paper looks at the selection of a crew member from among the "on call" crew for each piece of work for which a crew member is required. This selection is made in order to maximise the number of crew that remain "on call" in the future.

1 Introduction

Operations research techniques have been used to solve airline problems for many years. Almost every aspect of an airlines operation can be considered a scheduling problem. Aircraft, crew, passengers, freight and sundry items must be organised into an efficient operation on a daily basis. Therefore the potential finanicial benefits for using optimisation methods to produce optimal or near optimal schedules is significant.

1.1 Airline Scheduling Problems

The airline-scheduling problem is one of the areas that airlines can most benefit from using optimisation techniques. The solution to the airline-scheduling problem is a timetable of flights to be operated, with an aircraft, pilots and cabin crew assigned to each flight.

Once the airline-scheduling problem has been solved, the schedule must be maintained. This involves incorporating changes to aircraft, crew availability and flight requirements. This area is known as the "day of operations". One aspect of the operations management is call repair which it the subject of this paper.

1.1.1 Day of Operations

The airline-scheduling problem is solved well ahead of when the flights are operated. Any changes that occur after the rosters are published need to be resolved with minimal change to what was published. These changes are referred to as “day of operations” or “on the day” changes.

In a small or medium sized airline it is not practical to have spare aircraft waiting to fill in when a disruption occurs, so aircraft may need to be diverted to allow as much of the flight schedule as possible to be operated. However, it is practical to have spare crew available. When the rosters are built, some crew are assigned a call (or reserve) line. A call line is a period (up to 28 days long) during which the crew member must be available to be called out at short notice and used to fill in for the rostered crew in a disruption.

1.1.2 Air New Zealand Long Haul Cabin Crew

Most of Air New Zealands Long Haul Cabin Crew are currently based in Auckland. They are able to operate on all the Boeing 747-400, 767-200 and 767-300 that Air New Zealand have in their current fleet. They have a 28 day roster period and receive their roster ten days before the roster begins. The Long Haul Cabin Crew is currently made up of the following ranks.

- Inflight Service Director (ISD)
- Inflight Servide Co-ordiniator (ISC)
- Flight Attendant Premium Service (FAPS)
- Flight Attendant Pacific Clas (FAPC)

The number of crew required on each flight differs depending on the aircraft type.

2 Call Repair Process

2.1 Call Lines

Call lines must be created so that there are enough crew members of each rank on call at any particular time. As there are call lines starting every day, and each call line must contain a legal sequence of call duties, this is not a straightforward problem.

The first and most difficult part of call line generation is working out how many crew members are required on each day. Some earlier work has been done in this area (Teodorovic 1988).

The second part of the call line generation is the creation of the call lines themselves. This involves turning the call requirements calculated in the first step into a series of call lines which provide enough coverage everyday and meet all the legal and contractual rules. The call lines are made up of on call periods and off days. Call line contents differ depending on the crew type. Some of the on call days will be in the morning and others will be in the afternoon. For some crew groups there are 24-hour call activities which mean that the crew member must be available at short notice anytime of the day or night.

Some work was done on this problem at Air New Zealand (Deaker 1994), where call requirements specified by Air New Zealand were used to develop a model which produced a set of legal call lines with the required call coverage.

Call lines are assigned at the pre-assignment stage of the roster-build. Call lines are generally unpopular among crew members due to the unpredictability they represent. For this reason, they are rostered on a strictly rotational basis.

When a call line is created, there are certain rules that must be followed. These rules also apply when a call line is altered once a crew member has been called out. There are also rules concerning which trips a crew member maybe called out for. These tripes are generally reffered to as Tours of Duty (ToD). These are a series of flight sectors, starting and ending at a base where are crew member is stationed. ToDs are built without a particular crew member in mind. Once the ToDs have been created they must be assigned to a crew member to form a roster. Every crew member must have a legal sequence of duties, or a Line of Work (LoW).

2.1.1 Call Line Rules

Every Call line (and the call duties inside it) must meet the following rules:

- There must be an activity on every day (even if it is an OFC day)
- The activities must no overlap
- All preassignments must be included in the LoW. These may include previous callout ToDs as well as activities before or after the call line
- There must be the required number of days off
- There must be a 36 hour period free of duty every seven days
- There must not be more than four consecutive standby (or call) duties (ONC,SB1,SB2,SBY) without an intervening OFC day

An OFC day is an off day in a call line.

2.1.2 Callout Rules

When a crew member is being considered for a callout ToD, they must meet the following requirements:

- The ToD must form a legal LoW with the crew member's other assignments
- Long haul cabin crew members must report to the airport within an hour of being called out
- The ToD must start either during the crew members call duty or within two and a half hours of the end of the call duty
- The ToD must end by midnight on the last day of the crew members call line

2.2 Call Repair

Each day, there are several ToDs which require extra crew members due to disruptions (they are called "open ToDs"). Call Repair is the selection of crew members from amongst those who are on call, to fill all of the open ToDs. It also includes the regeneration of each person's call line once they return from their callout ToD.

When a crew member is called out on a ToD, all of the duties which they were assigned for the duration of the ToD must be removed. This reduces the call coverage available on those days. The objective of call repair is to call out crew members for ToDs and regenerate crew members' call lines in such a way that effect of coverage reduction due to call out is minimised.

The call repair process is quite complicated, as there are up to 125 long haul cabin crew on a call line on any given day. Some of these will be on an OFC day, and others will not meet all of the rules for being called out on a ToD.

3 Generating Lines of Work

Currently the call repair process is manual. It is easy for the crew controllers who make callout decisions to select a crew member who is available to operate each open ToD, however it is difficult for them to see the effect this callout will have on future call coverage. The main reason for optimisation of the call repair process is to allow callout decisions to take the effect of call coverage into account. The objective of the optimisation is to maximise call coverage while assigning the required number of crew members to each ToD.

The call repair optimisation considers all of the long haul cabin crew members who have one or more call duties in the next 28 days, and all of the long haul cabin crew ToDs which have less than the required number of crew assigned for the next 28 days.

The optimisation is written as three separate Fortran 77 programs: the prebuild and generator together create the LoWs, and the optimiser finds the optimal set of LoWs.

3.1 ToD Eligibility

The first part of the call repair optimisation is called the prebuild. It determines which crew members are eligible for each open ToD. At this stage all of the call out rules are checked along with all the rules that apply to the assignment of ToDs (e.g. duty and flight time limitations).

The output from the prebuild is a list of open ToDs for each crew member which would form a legal LoW with all the other ToDs and non-call activities currently assigned to the crew member. Some crew members maybe eligible for multiple ToDs and many will not be eligible for any ToDs.

3.2 Generating Call Lines

Once the prebuild is finished, the call repair generator is run. This is the part of the optimisation which considers possible call line continuations, taking into account the open ToDs the crew member is eligible for. The generation process starts by creating a series of availables for each crew member under consideration. The availables are activities which the crew member is able to be assigned, including current assignments, open ToDs and call activities.

Some crew members are not considered by the generator because they do not have any call activities in the period of time which is being generated for (e.g. the next ten days). The generator then uses the availables created to form partial LoWs for each crew member. A partial LoW is a LoW which starts on or before the day the optimiser is being run and ends within the next ten days. The generation process creates a tree with partial LoWs for nodes and availables for arcs.

The final part of the call repair generation is to complete all of the partial LoWs. At this stage some LoWs will be suppressed because they are considered undesirable. If the optimisation is infeasible these LoWs will be included to allow a solution to be found.

Each LoW is assigned a cost so that the optimisation can discriminate between them. The cost is calculated as follows:

- If the LoW is suppressed it is assigned maximum cost
- LoWs with open ToDs are costed to bias the solution towards crew finishing their call line soon
- A small cost is added to all LoWs reflecting the number of days off near the start of the LoW
- An additional cost is added to LoWs that have less than the standard number of days off

4 Call Repair Optimisation

The set of LoWs from the generator is passed to an optimiser which uses a generalised set partitioning model. The optimisation will consider all non-suppressed LoWs generated by the call repair generator, all of the open ToDs within the window of time being looked at, and all of the crew members who had a LoW generated for them. Only open ToDs and call duties from the generated LoWs are included in the optimisation - all other activities the crew member may have already been assigned are excluded.

4.1 Model Formulation

The basic format of the problem is:

$$\begin{array}{ll} \text{Minimise} & z = \underline{c}^T \underline{x} \\ \text{Subject to} & \mathbf{A}\underline{x} = \underline{b}, \qquad x_j \in \{0,1\}, \qquad j = 1,...p \end{array}$$

The columns of the problem are the LoWs generated for each crew member, $\mathbf{A}$ is a zero-one matrix consisting of open ToD, call duty and crew constraints, $\underline{c}$ is a vector containing the cost of each LoW, and $\underline{b}$ is a vector containing the right hand sides for the constraints. The solution is given by $\underline{x}$, where

$$x_j = \begin{cases} 1 & \text{if the LoW } j \text{ is included in the solution, or} \\ 0 & \text{if the LoW } j \text{ is not included in the final solution} \end{cases}$$

The LoWs are represented in the problem by columns of zero-one values (the jth LoW is the jth column of the $\mathbf{A}$ matrix, a_{ij}). Each row represents one open ToD or call duty, and the ones in the LoW column signify which ToDs and call duties appear in the LoW.

$$a_{ij} = \begin{cases} 1 & \text{if the ith open tod or call duty is in the jth low} \\ 0 & \text{otherwise} \end{cases}$$

The index of the crew member each LoW belongs to is stored as a one at the end of each column, with the position of the one representing the index of the crew member. If two crew members have identical LoWs, the LoW will appear in the optimisation twice, with different crew member tags.

The model has three sets of constraints: the open ToDs must be covered, the call duties must be covered and every crew member must have exactly one LoW.

4.1.1 ToD Constraints

There is one constraint for each open ToD with the following format:

$$\sum_{j=1}^{s} a_{ij}x_j + v_i - y_i = d_i \tag{1}$$

where i is the index of an open tod
s is the total number of LoWs for all crew members
d_i is the number of crew members required on the ith open ToD
v_i and y_i are slack and surplus variables

There is one open ToD constraint for each crew rank required on each open ToD. If a ToD requires two FAPC crew members and an ISC, there will be one constraint for the FAPC crew and another for the ISC.

4.1.2 Call Duty Constraints

The constraints for call duties are similar to those for open ToDs

$$\sum_{j=1}^{s} a_{ij}x_j + w_i - z_i = e_i \tag{2}$$

where i is the index of a call duty
s is the total number of LoWs for all crew members
e_i is the number of crew members required on the ith call duty
v_i and y_i are slack and surplus variables

There is a constraint for each call duty in the optimisation i.e. there is one constraint for each call duty type, for each crew rank which can be assigned this type, for each day that is being optimised.

4.1.3 Crew Constraints

The crew constraints work in quite a different way from the ToD and call duty constraints:

$$\sum_{j=1}^{s} a_{ij}x_j + u_i = 1 \tag{3}$$

where $a_{ij} = \begin{cases} 1 & \text{if the } j\text{th LoW belongs to the } i\text{th crew member} \\ 0 & \text{otherwise} \end{cases}$
i is the index of the crew member
j is the index of the LoW

If variable u_i has a value 0, then the crew constraint ensures that the ith crew member has exactly one LoW in the solution.

4.2 Variable Costs

The objective of any optimisation is to find the feasible solution with the lowest cost. The cost assigned to each variable determines which solution is optimal. In this problem the costs are assigned as follows:

- LoWs have costs assigned by the generator
- Crew slacks have the highest cost possible
- ToD slacks also have a very high cost, as a solution which does not cover all of the open ToDs is of no practical use
- Call slacks have low to medium cost
- ToD surpluses have a high cost
- Call surpluses have a zero cost, because extra call coverage is desirable

4.3 Initial Models

The first model provided good solutions but had a major downfall - it was indiscriminate about which call duties it undercovers, other than a bias towards future calls. A typical solution using the above model may include call duties that are undercovered by four or five crew members. A much better solution would be to undercover four or five different call duties by one crew member each. This issue was addressed by modifying the call duty constraints and the addition of some new constraints. The new constraint contains three slack variables for each call duty with the first two having a maximum value of one. The cost of the first slack is the cost of undercovering the call duty by one unit. This cost is low. The cost of undercovering the call duty by a second unit is medium and is in the second slack. The cost of undercovering the call duty by a third or subsequent unit is high. The optimisation will always choose the correct variable for the first unit of undercoverage due to the associated costs.

After these changes were made it was noted that this model performs much better than the original model, however it still has one major operational flaw. This is associated with the selection of crew assigned to each call duty. The new model ensures that call coverage is as close as possible to the required number of crew each day, but there is no control over the characteristics of the crew that are assigned each call duty.

The main area where control is desirable is the number of days a crew member has left in their call line. This is important because the current model could result in the correct number of crew being assigned, but none of them being available for callout. For example, consider the number of FAPS crew on morning call on Thursday. The requirement for this call duty might have been specified as six and the optimisation may have been able to satisfy this. There maybe a twelve day ToD which departs early in the morning and requires a FAPS crew member on Thursday. Although there are six crew available, say one of them is on the last day of their call line, two only have three days to go, two more have eight days to go, and one has two days to go. None of the six crew on call are available for the twelve day ToD.

To overcome this problem more changes to the model were required. This time a certain number of the crew assigned to each call duty must be available for medium or long ToDs. The definition of the length of medium and long tods is arbitrary however the classifications the optimisation used is that a medium ToDs is lasting seven or more days, and a long ToD as lasting 12 or more days. For each call duties two new constraints are required: one for the coverage of medium-length ToDs and one for coverage of long ToDs. This was the final model that was used.

4.4 Optimisation Process

The call repair problem is actually four independent problems being solved simultaneously. Every constraint only applies to one crew rank. Similarly the variables all apply to only one rank. This means that the call repair problem could be solved as four separate problems (one for each rank). The combined problem is being solved instead of four separate problems for two reasons: it is easier for users to only have one optimisation to run, and the combined model makes it much easier to allocate a crew member of a different rank if the original problem is infeasible.

A common feature of set partitioning models, which also appears in the call repair problem is constraint matrix sparseness. The sparseness of the call repair constraint matrix allows a big saving on storage space - instead of storing the entire variable, each variable is stored as a list of the rows in which a one occurs.

The call repair process was solved using the revised simplex method, with an integer programming package called ZIP (Ryan 1980). ZIP carries out the primal revised simplex iterations using information passed to it by the main program. At each stage ZIP only has information about the variables which are currently in the basis. ZIP does not have any information about the real-world problem that is being solved.

5 Results

The call repair optimisation has been tested extensively in Air New Zealand. Each day the call repair optimisation results were compared to the decisions made by the crew controllers who manually assigned open ToDs. The call coverage in the future was where the majority of benefits were obtained. Because the optimised solution could not be implemented we are unable to see the flow on effect of the extra call duties that could be assigned. So the next day the optimisation was run, these same new call duties would be assigned again. If the first solution had been implemented then there would be less opportunity for improvements in subsequent runs. In other words the optimisation makes the same improvements day after day because the last time it made the improvement it was not implemented. For this reason it is difficult to quantify the improvements made by the call repair optimisation.

5.1 Run Times

An impportant feature of the call repair optimisation is speed. It must be able to provide solutions very quickly, so that solutions may be optimised and re-optimised several times as call and open ToD requirements change. The bottleneck of the call repair optimisation is the text file extraction from the database. This generally takes about $\frac{1}{2}$ and hour to 45 minutes. The files are currently created during the night, so they are available for use when Air New Zealand's first shift begins at 6am. The many test scenarios that have been completed have a run time of less than 25 seconds. This is the time taken to do a complete optimisation (prebuild, generate and optimise), although the user will often need to alter data files and re-run parts of the optimisation as open ToD requirements change.

The reasons for the speed of the optimisation are:

- The number of crew in the optimisation is small (when compared to rosters)

- The number of ToDs to be assigned is very small (always less than 20)
- The problem is very flexible, so it is very easy to find a feasible solution
- Solutions are naturally integer

5.2 Solution Quality

The main measure of quality in call repair solutions is the number of call duties assigned. However there are some other quality considerations.

One is the selection of crew members to be called out for ToDs. This is based solely on maximising the call coverage. The current (manual) selection process involves using crew members that are close to the end of their call line first as they would not be available for any long ToDs. In the test runs of the optimisation it became apparent that this was not necessarily conducive to maximising call coverage. The optimisation routinely chooses crew members who are at the start of their call line for short ToDs, because this allows it to rearrange their call line and provided better call coverage.

It is not acceptable to the Day of Operations area to call out crew members for short ToDs near the start of their call line, as they perceive this as reducing their ability to cover long ToDs, (even though the optimisation explicitly covers long ToDs). There is now a penalty to discourage the optimisation from calling out crew members close to the start of their call line. However if the call coverage that can be changed is significant, this will outweigh the penalty.

As the call repair optimisation is implemented, an ongoing consultation process will be required to ensure that the optimisation meets the Crew Control requirements. The process that Crew Control currently use may also need to be reviewed, to ensure that they have sound reasons behind them, rather than just being based on tradition.

6 Conclusions

6.1 Summary

This paper has investigated the generation and optimisation of call lines when crew members are called out on open ToDs. The first stage of the process determines which open ToDs each crew member is eligible for and the second stage enumerates the possible call line continuations over a ten day period. Finally the third stage finds the optimal solution from the generated call lines.

Several models were developed before an adequate solution was found. This allowed for undercoverage for call duties to be spread evenly. Crew selected for call duties were constrained by a minimum requirement for people to be available for medium (7 or more days) and long (12 or more days) to prevent everyone available being close to the end of their call line.

Visual Basic was used to develop a user interface for the optimisation. This interface allows the user to select a scenario, run the optimisation, edit parameters and data files, and view the output from the optimisation. The solution can also be viewed in a graphical, colour-coded format.

The optimisation produced good results, with significant increase in assignment of call duties over the current manual process. The run times are a significant improvement over the manual process and is generally completed in under 25 CPU-seconds.

6.2 Advantages of Optimisation

Optimisation of the call repair process has the following advantages over Air New Zealand's current manual process:

- Speed: the optimisation runs very quickly, allowing Air New Zealand staff to focus on other work
- Efficiency: the manual process does not usually include regeneration of the call line following assignment of a call out ToD
- Continual re-evaluation: while a crew member is away on a call out ToD their call line is re-evaluated on a daily basis to maximise call coverage
- Flexibility: the optimisation considers every possible call line continuation
- Automated rules checking: the optimisation automatically checks all of the rules, while the manual process is prone to human error
- Viewpoint: the optimisation is able to look at the overall call coverage and find the days which require more crew members, while a manual users is only able to consider a few crew members and their call lines at a time

The manual process does not have any known advantages over the optimisation.

6.3 Implementation

The call repair optimisation and user interface were developed by Day (1999) and have been ready for some time for use in Air New Zealand's production database. Software has now been developed using C++ to download the required information from the database every morning prior to the crew controller starting work. This means the program is ready for use first thing. Software has also been developed to save the solution to the database at the completion of the optimisation. This is currently in its final testing stages ready for implementation.

References

Deaker, J. 1994. "Construction of standby crew patterns for international flight attendants at Air New Zealand." *Proceedings of the 30th Annual Conference of ORSNZ.*

Ryan, D. 1980. *ZIP - A Zero-One Integer Programming Package for scheduling.* Harwell, Oxfordshire: A.E.R.E.

Teodorovic, D. 1988. *Airline Operations Research.* Gordon and Breach Science Publishers.