Probabilistic Airline Reserve Crew Scheduling Model*

Christopher Bayliss[1], Geert De Maere[1], Jason Atkin[1], and Marc Paelinck[2]

1 **ASAP Group, School of Computer Science, University of Nottingham, Jubilee Campus, U.K**
{cwb,gdm,jaa}@cs.nott.ac.uk

2 **KLM Decision Support, Information services department**
KLM Royal Dutch Airlines, KLM Headquarters, The Netherlands
Marc.paelinck@klm.com

Abstract

This paper introduces a probabilistic model for airline reserve crew scheduling. The model can be applied to any schedules which consist of a stream of departures from a single airport. We assume that reserve crew demand can be captured by an independent probability of crew absence for each departure. The aim of our model is to assign some fixed number of available reserve crew in such a way that the overall probability of crew unavailability in an uncertain operating environment is minimised. A comparison of different probabilistic objective functions, in terms of the most desirable simulation results, is carried out, complete with an interpretation of the results. A sample of heuristic solution methods are then tested and compared to the optimal solutions on a set of problem instances, based on the best objective function found. The current model can be applied in the early planning phase of reserve crew scheduling, when very little information is known about crew absence related disruptions. The main conclusions include the finding that the probabilistic objective function approach gives solutions whose objective values correlate strongly with the results that these solutions will get on average in repeated simulations. Minimisation of the sum of the probabilities of crew unavailability was observed to be the best surrogate objective function for reserve crew schedules that perform well in simulation. A list of extensions that could be made to the model is then provided, followed by conclusions that summarise the findings and important results obtained.

1998 ACM Subject Classification G.1.6 Optimization, G.3 Probability and Statistics

Keywords and phrases airline reserve, crew scheduling, probabilistic model

Digital Object Identifier 10.4230/OASIcs.ATMOS.2012.132

1 Introduction

This paper is related to the airline industry and in particular reserve crew scheduling. Much time and effort is put into improving the quality of airline schedules, mainly to minimise the cost of operations and improve the schedules' ability to recover from delays and cancellations. Outlines of the airline scheduling process are given in [3, 4, 6, 5]. The process starts with schedule design, in which airlines determine sets of origins and destinations and the corresponding departure and arrival times for the flights they will fly. These are called the flight legs. Following schedule design, fleet assignment determines which types

* This work was partially supported by the LANCS initiative.

12th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS'12).
Editors: Daniel Delling, Leo Liberti; pp. 132–143
OpenAccess Series in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

of aircraft will be assigned to each flight leg. Maintenance routing then assigns particular aircraft (tail numbers) to each flight leg in such a way that each aircraft spends sufficient time at maintenance stations. Crew scheduling can be performed once aircraft schedules are determined. Crew scheduling is typically solved sequentially as two sub-problems, firstly crew pairing in which anonymous feasible crew schedules are generated, and then the assignment of crew pairings to individual crew as the second sub problem. Crew schedules rarely go according to plan due to unforeseen disruptions, including crew absence, delays (ground and airside), unexpected maintenance requirements and delays due to the knock on effects of these. Crew constitute a large cost for airlines, second only to fuel costs [2], this coupled with the fact that crew schedules rarely go according to plan means that in order to keep airline operations running as smoothly as possible, reserve crew have to be scheduled so as to be ready to absorb crew related disruptions.

An outline of this paper is as follows: section 2 presents the probabilistic reserve crew scheduling model which forms the basis for the whole paper; section 3 discusses how input data for the model can be derived from the real world; section 4 gives results from experiments as to the best surrogate objective function for yielding reserve crew schedules that absorb large amounts of crew disruption; section 5 discusses solution methods that can be used to solve the model concentrating on the balance between solution quality and speed of solution; section 6 summarises the findings presented in the paper and discusses some extensions to the basic model to make it more realistic.

2 Probabilistic Reserve Crew Scheduling Model

2.1 Problem description

The problem is to assign duty start times to a fixed number of reserve crew. Reserve crew requirement is subject to the uncertainty of crew absence. This model schedules reserve crew start times such that the overall probability of crew unavailability for a sequence of departures from a single airport is minimised. A feature of this model is that possible reserve crew start times correspond to scheduled departure times, this step is based on the idea that standard duty start times are an inefficient approach, because many reserve crews beginning duty periods simultaneously leads to reserve capacity that cannot possibly be used.

2.2 Assumptions

The probabilistic reserve crew scheduling model is based on the following simplifying assumptions;

1. The reserve crew scheduling problem consists of a set of departures from a crew base where each departure has probabilities of crew absence and therefore the need for reserve crew.
2. The maximum demand for crew per departure is 1. This could be interpreted as a single team of reserve crew rather than an individual.
3. The chance of crew absence is captured accurately by independent probabilities for each departure.
4. Reserve crew cover for the first crew absence that occurs within their fixed length duty period, if more than one reserve crew member is available then the crew member who started their duty period first is used.

5. Reserve crew can undertake a maximum of a single duty within any one duty period. This can be justified by the fact that when reserve crew are used they typically adopt the remainder of the pairing (string of flight duties) of the absent crew member they are covering for.
6. Reserve crew duties begin at times corresponding to scheduled departures.

Assumptions 1 and 3 mean that the base problem can be represented as a vector of probabilities (P) where p_i represents the probability that the crew scheduled for departure i will be absent. Assumption 2 means that the problem can be represented as a vector as opposed to a matrix for specifying the probabilities of different numbers of crew for each departure (see section 6). Assumptions 4 and 5 state the way that reserve crew are to be used: reserve crew duties are fixed in length equal to some integer number of departures (L) (note that the constant L can be replaced by L_i for the case where departures are not at equal intervals as assumed in the following), reserve crew cover for the first disruption that occurs within their duty period, they can only cover one flight and therefore once they undertake a cover duty they cannot cover for any of the remaining departures that occur within their duty period (see section 6 for the extension). They also state the precedence ordering for the use of reserve crew in the event that more than one reserve crew is available for a departure, which is to use the reserve crew who has been on duty the longest. Another assumption that could have been added is that no more than one reserve crew can be assigned to begin a reserve crew duty at the same time, but this is just a way of restricting the solution space to rule out solutions that will definitely be sub-optimal for this model. The reason being that if two or more reserve crews begin a duty at the same time, at least one of them will not be utilised in that period unless you have two flights at the same time and starting one later would cover at least one extra departure. Assumption 6 tries to minimise wasting reserve crew duty time by not scheduling them before the first time at which they may be required to cover crew absence.

2.3 Parameters and Variables

The parameters and variables are as follows;

$\boldsymbol{P}$: vector containing the probabilities of crew absence for each departure
$\boldsymbol{X}$: reserve crew schedule (solution)
$\boldsymbol{L}$: length of a reserve crew duty period
$\boldsymbol{R}$: number of reserve crew available for scheduling
$\boldsymbol{N}$: number of departures

2.4 Probability calculations

The problem can be represented as a vector of probabilities P, where each element denotes the probability of crew absence for a particular departure. A reserve crew schedule X can be represented as a list of reserve crew duty start times, or equivalently departure numbers.

$$P = \{p_1\ p_2\ p_3\ p_4\ ...\} \tag{1}$$
$$X = \{x_1\ x_2\ x_3\ x_4\ ...\} \tag{2}$$

For a given set of departures with associated probabilities of crew absence and a reserve crew schedule, we can determine the effect the given reserve crew schedule has on the probabilities of crew unavailability. The vector of probabilities of crew unavailability is denoted by P' and is a function of the reserve crew schedule and the probabilities of crew absence of the originally scheduled crew, $P' = f(P, X)$. The procedure for finding P' for a given reserve crew schedule is as follows and reflects the assumptions given in section 2.

Algorithm for calculating the effects of reserve crews on the probabilities of crew unavailability for a sequence of departures

In words the following pseudo code states that, for each reserve crew that begins a duty (i loop), initialise the probability of that reserve crew's availability (pr) to 1, set the probability of crew unavailability (pd) to 0 for that departure number. Then for each departure that occurs (j sub loop) in a reserve crew's duty update the probabilities that they are still available having not been used in previous departures and update the probability that no crew are available for that departure. The first line of the pseudo-code sets P' equal to P because P represents the initial problem and P' represents the probabilities after the reserve crew schedule has been taken into account.

$P' = P$
for $i = 1$ to R **do**
  $pd = 0$, $pr = 1$
  if $x_i < N$ **then**
    for $j = 2$ to $\min(1 + N - x_i, L)$ **do**
      $pr = pr(1 - p'_{x_i+j-2})$
      $p'_{x_i+j-2} = pd$
      $pd = p'_{x_i+j-1}(1 - pr)$
    end for
  end if
end for

2.5 Non Linearity of P'

Section 2 so far gives a method for measuring the quality of a given reserve crew schedule, ideally we would like to minimise the vector P', to minimise the probability of crew unavailability. As a result P' can be used as an objective function which has to be minimised in some way and X is the variable. The first reaction would be to try to solve for the reserve crew schedule (X) as a MIP (mixed integer programming model), however the probabilistic model is nonlinear to a polynomial degree of R, the number of reserve crew to be scheduled. The reason is due to the possibility of overlapping reserve crew duties, this interaction of reserve crew duties means that the probabilities of reserve crew being required depends on whether a reserve crew was assigned previously.

2.6 Solution space

The size of the solution space is given by all combinations of R reserve crew in N positions.

$$Number\ of\ feasible\ solutions = \frac{N!}{R!(N-R)!} \tag{3}$$

The structure of the solution space is such that the natural definition of neighbouring solutions is that of moving a reserve crew's start time to a start time with no reserve crew assigned.

3 Input probabilities

The method for calculating the effect that a reserve crew schedule has on the probabilities of crew unavailability outlined in section 2 is underpinned by estimates of probabilities for crew absence for each departing aircraft. The accuracy of the model being put forward in this paper is therefore influenced by the accuracy of the input probabilities. Accurate estimates of crew absence probabilities can be derived from large sets of historic data and regression models can be used to estimate the systematic influence of factors such as origin-destination pairs, the time of day, month or year, and the potential interactions between them. In addition to probabilities for crew absence derived from historic data, characteristics from the schedule that increase the probability that a delay will propagate through the schedule could also be taken into account. Examples thereof include crew duty times close to the legal maximum (little slack) and delay propagation trees that model resource connections [1]. Finally, crew availability is influenced by the flexibility within a schedule to recover from delays, and thereby prevent delay propagation. Examples of such recovery patterns include crew swaps or move-up crews [9] and aircraft swaps.

4 Experimental results

4.1 Input probabilities for experiments

For the current investigation of the probabilistic mathematical model, probabilities of crew absence (P) will be generated randomly from a uniform distribution for each sequence of departures from a crew base. Due to the use of uniform random numbers the expected number of crew absences will usually be approximately half the number of departures (N) considered and this can be used as a reference point when interpreting solutions derived from the probabilistic model. For example if N is 25, 12.5 crew absences can be expected, now if 9 reserve crew are available and a given reserve crew schedule gives simulation results with an average crew shortage of 4.5, this can be interpreted as reserve crew utilisation of 0.88 ($(12.5 - 4.5)/9$), where higher reserve crew utilisation indicates superior reserve crew schedules. Equation 4 shows how to calculate expected reserve crew utilisation for any given instance of P.

4.2 Comparison of objective functions

The best reserve crew schedule will be the one that covers most of the potential for crew absence. This can be achieved by seeking to minimise the vector P' in some way as this is equivalent to minimising the average probability of crew unavailability. The following experiment is designed to find the best surrogate objective function involving P' that leads to the reserve crew schedule with the most desirable properties in terms of covering possible crew absences under simulation.

Nine alternative objective functions are considered (table 1), all minimisation based i.e. minimise the, A) sum of P', B) max of P', C) standard deviation of P', D) coefficient of variation of P', E) product of the mean and standard deviation of P', F) weighted sum of mean and max of P' and finally G), H) and I) with mean absolute deviation replacing standard deviation in C), D) and E) respectively. The reasoning is as follows, A) minimising the sum of P' is equivalent to minimising the average probability of crew absence. B) minimising the maximum which is equivalent to reducing the variation among P' as well as suppressing max P'. C) minimising the standard deviation of P' will act to minimise the difference between

Table 1 Objective functions.

Symbol	Objective function	Equation
A	Sum of P'	$\sum_{k=1}^{N} p_k^{'}$
B	Max P'	$\max_{k=1..N}(p_k^{'})$
C	Standard deviation of P'	$s = \sqrt{\frac{\sum_{k=1}^{N}(p_k^{'} - \bar{P'})}{N}}$
D	Coefficient of variation of P'	$\frac{s}{\bar{P'}}$
E	Product of mean and standard deviation of P'	$s\bar{P}'$
F	Weight sum of mean and max P'	$a\bar{P}' + b\max_{k=1..N}(p_k^{'})$
G	Mean absolute deviation of P'	m
H	D with mean absolute deviation	$\frac{m}{\bar{P'}}$
I	E with mean absolute deviation	$m\bar{P}'$

the lowest probabilities (guaranteed to be zero) and the higher probabilities, the easiest way to do this is to try to minimise the largest members of P', which is similar to minimising the maximum probability. The only difference is that calculating the standard deviation is computationally more demanding on computer resources. Also, its probability minimising properties are implied rather than direct as they rely on the problem/model structure. D to I) are based on similar reasoning.
The experimental method is to generate and store 20 example problem instances based on $N = 25, R = 9, L = 3$ and P generated using uniform random numbers (section 4.1), then enumerate the 2042975 feasible solutions (equation 3) to find and store the optimal reserve crew schedules for each objective function and problem instance. After this, the optimal reserve crew schedules for each objective function will be tested in simulations in order to find which objective function leads to reserve crew schedules which have the highest reserve crew utilization.
The simulations use random numbers in conjunction with the initial input probabilities to generate realisations of crew absences for each of the 25 departures in each problem instance. Reserve crew are used in the way specified in assumptions 4, 5 and 6 in section 2. 100 repeat simulations are performed for each problem instance (2000 simulations in total). Reserve crew schedules corresponding to each objective function for each problem instance will be tested by the same 100 repeat simulations (each objective being tested simultaneously). 100 repeated simulations are thought to be enough to sample a wide range of possible outcomes and also to get some feel for the average outcome for each problem instance. The results are summarised in tables 2 and 3 which give reserve crew utilisation and flight cancellation totals from 2000 simulations for each objective function.

Table 2 shows that the sum of P objective function has the fewest unused reserve crew over 2000 simulations and therefore is the probabilistic surrogate objective function with highest reserve crew utilisation. The fourth column gives the ranks of the objective functions in terms of reserve crew utilisation. Note the the reserve crew utilisation values are calculated by the number of reserve crew scheduled in 2000 simulations (18000) minus unused reserve crew divided by 18000.
Table 3 gives the total number of flights that were cancelled in the simulations as the originally assigned crew were absent and no reserve crew were available. The sum of P objective gave the highest number of legal flights followed closely by objective function I (product of mean and mean absolute deviation). The results of table 3 make it clear that in order to guarantee feasibility of all departures a number of reserve crew equal to the number

Table 2 Overall reserve crew utilisation results from 2000 simulations.

Objective function	Total unused reserve crew	Reserve crew utilisation	Rank
A	981	0.9455	1
B	1255	0.9303	7
C	1009	0.9439	4
D	2255	0.8747	9
E	1004	0.9442	3
F	1062	0.9410	6
G	1017	0.9435	5
H	1951	0.8916	8
I	1000	0.9444	2

Table 3 Total cancellations in 2000 simulations (50000 flights).

Objective function	Total cancellations	Cancellation rate	Rank
A	9242	0.1848	1
B	9515	0.1904	7
C	9319	0.1864	5
D	12374	0.2475	9
E	9296	0.1859	3
F	9354	0.1871	6
G	9301	0.1860	4
H	11891	0.2378	8
I	9278	0.1856	2

of departures would be required, which would be very costly. So ideally a balance between the cost of the reserve crews and the cost of cancelling flights is required. This gives rise to a possible multiple objective optimisation formulation. However, in reality airlines have the ability to call out freelance reserve crew as well as other possibilities, so the balance has to include these as well.
Tables 2 and 3 give two very important simulation derived measures of the quality of reserve crew schedules in comparison to other reserve crew schedules. High reserve crew utilisation means that airlines are paying for reserve crew that will be used. Secondly cancellation rates indicate what proportion of flights are infeasible in relation to a given reserve crew schedule. Another interesting point from tables 2 and 3 is that there is a close connection between the performance each objective demonstrates for each measure of solution quality, this can be seen in the similarity of the rank ordering of objective functions in both tables 2 and 3. In conclusion it has been found that minimising the sum of P is the best surrogate objective function for deriving quality reserve crew schedules from this probabilistic model.

5 Solution methods

5.1 Comparison of solution methods

In this section it is shown that the objective function value for a given reserve crew schedule implies the average reserve crew utilisation level and cancellation rate that the reserve crew schedule will give in simulation. This conclusion is based upon applying a variety of heuristic solution techniques in order to minimise P'. The same 20 problem instances which were used in section 4 are used again here because the optimal solutions to these have already been enumerated. Table 4 gives the average objective values and solution times[1]. Additionally table 4 compares simulation results with expected results derived theoretically from the surrogate objective values for the criteria of reserve crew utilisation and cancellation rates. The objective function in each case is sum of P' as this was found to be the most effective surrogate objective function in section 4. The first row contains the results found from enumeration in section 4 and correspond to the optimal solution, this gives a benchmark for judging the effectiveness of the various search and optimisation techniques. The solution methods tested include local search, population based algorithms, greedy algorithms, pruned dynamic programming algorithm and problem specific heuristics. The pruned dynamic programming algorithm was the only method which identified all optimal solutions. Note that this is not guaranteed to occur because the pruned dynamic programming algorithm does allow the possibility that a partial solution corresponding to an optimal solution might be pruned at an early stage of the search due to the use of heuristics for estimating upper and lower bounds of partial solutions (see section 5.2).

5.2 Description of solution methods

Pruned dynamic programming algorithm

The method referred to as a pruned dynamic programming algorithm (DP) is based on dynamic programming in that it uses the idea of states (number of reserve crew assigned) and stages (departure number). It is also a branch and bound algorithm because entire branches can be eliminated early during the search. The algorithm constructs and searches a binary tree in a breadth first manner where each level of branching represents a departure time and each path from the root of the tree to a leaf represents a partial solution. Each iteration of the algorithm considers the next departure and adds a layer of depth to the tree. The algorithm branches on each leaf remaining from the end of the previous iteration until all departures have been considered. To cut down the amount of the complete binary tree of feasible solutions that needs to be searched in order to find a good solution upper and lower bounds corresponding to each partial solution are estimated. Lower bounds and upper bounds are heuristically estimated, the lower bound heuristic is such that it always gives a solution of better quality (quality here means lower objective value) than the upper bound heuristic both starting from the same partial solution. Partial solutions are than eliminated if their lower bound is greater than the minimum upper bound of partial solutions in an equal or higher state. The use of upper and lower bound heuristics for pruning partial solutions makes this method a heuristic, but a heuristic with a tunably high probability of obtaining optimal solutions. The choice of upper and lower bound solutions influences how ruthless the pruning strategy is. The algorithm can be made faster or more accurate by using (respectively) lower or higher quality upper bound heuristics, provided that the lower bound heuristic is always the best and most intelligent possible.

Population based heuristics

Genetic algorithm and ant-colony algorithms were investigated. The implementation of a genetic algorithm (GA) uses a binary vector representation of candidate reserve crew schedules, two-competitor tournament selection, single point crossover and a mutation rate

[1] Matlab, dual core 1.86ghz, 2gb, windows vista

of 0.001 applied to every chromosome. The constraint on using a fixed number of reserve crew means that crossover can lead to infeasible solutions with either more or less than the required number of reserve crew. This issue was dealt with by applying greedy heuristics (backwards and forwards heuristics, see the constructive heuristics subsection of section 5.2) to obtain candidate reserve crew schedules with the required number of reserve crew. Table 4 shows that the backwards heuristic alone performs better than the genetic algorithm which actually makes use of the same constructive heuristic. Due to the stochastic nature of genetic algorithms this outcome varies from run to run. The reason for this result is that the best and worst case results for the genetic algorithm span those of the constructive heuristics.
In the ant colony approach each of the 100 ants visits R positions of N, for each move made by an ant a cumulative distribution corresponding to the ant's next possible moves is created from the pheromone vector, then a random input is compared with the cumulative distribution to determine where the ant will move to. The sum of P' is computed for each ant's tour, an evaporation factor of 0.9 is applied to the pheromone vector, then the ant with the smallest objective value is used to lay pheromone in such a way that replenishes the amount evaporated.

Constructive heuristics

The backwards heuristics starts with the (optimal) infeasible solution of reserve crew assigned to each period, reserve crew are removed one at a time choosing the one that increases the objective value the least.
The forwards heuristic starts with no reserve crew assigned and adds one at a time choosing the one that decreases the objective value the most.
The basic greedy approach positions reserve crew corresponding to the R highest probabilities in the original probability of crew absence vector P.
The even distribution heuristic allocates reserve crews evenly across departures so that the number of reserve crew on duty for each departure remains constant.

Local search based methods

Hill climbing searches the local neighbourhood and takes the best move only if it is better than the current best. The hill climbing algorithm uses the neighbourhood structure described in section 2.6.
The simulated annealing implementation (SA) is based on: a temperature reduction being applied every 4 iterations (epoch=4); an initial temperature of 500; a final temperature of 0.001 and a geometric temperature reduction factor=0.9.
The tabu search implementation maintains a recency tabu list in which swapping of reserve crew between two positions in the schedule is prevented for a tenure of 50 iterations after the swap was last made. The method uses 100 iterations, always accepting the best non-tabu move.
The variable neighbourhood search method (VNS) uses 5 neighbourhoods (in order): single swap, cut and swap, single point crossover, sideways shift and random. If a neighbourhood contains a better solution the solution is accepted as the current solution and a new iteration begins, starting from neighbourhood one. If a better solution is not found the next neighbourhood is tested, if no improving solution is found after cycling through all neighbourhoods the procedure is terminated.

Table 4 Percentage of optimal, simulation coverage levels and solution times for a variety of solution methods.

Solution method	Objective value	Reserve crew utilisation	Expected reserve crew utilisation	Cancellation rate	Expected cancellation rate	Solution time (s)
Enumeration	4.5315	0.9455	0.9433	0.1848	0.1813	1296
Even distribution	5.1054	0.8783	0.8795	0.2061	0.2042	0.0160
Basic greedy	4.8258	0.9220	0.9106	0.1917	0.1930	0.0470
Forwards H	4.6125	0.9408	0.9343	0.1877	0.1845	0.1250
Backwards H	4.5655	0.9429	0.9395	0.1841	0.1826	0.1410
Hill climbing	**4.5340**	0.9464	0.9430	0.1823	0.1814	1.6540
SA	4.5517	0.9411	0.9410	0.1852	0.1821	0.6080
Tabu search	**4.5316**	0.9422	0.9433	0.1784	0.1813	16.2650
VNS	**4.5338**	0.9483	0.9430	0.1851	0.1814	18.7040
GA	4.5814	0.9433	0.9377	0.1844	0.1833	31.6070
Ant colony	4.5957	0.9369	0.9361	0.1853	0.1838	12.0900
DP	**4.5315**	0.9452	0.9433	0.1820	0.1813	38.0480

Expected value calculations

$$\textit{Expected reserve crew utilisation rate} = \frac{\left(\sum_{i=1}^{N} p_i) - (\sum_{i=1}^{N} p'_i\right)}{R} \tag{4}$$

Intuitively equation 4 means that the average probability of each reserve crew being used is the expected total reduction in crew absence due to reserve crew availability divided by the number of reserve crew available.

$$\textit{Expected cancellation rate} = \frac{\sum_{i=1}^{N} p'_i}{N} \tag{5}$$

The expected cancellation rate is simply the expected number of flights without crew divided by the number of flights.

Table 4 shows that three of the local search methods came very close to optimality, namely hill climbing, tabu search and variable neighbourhood search, of these tabu search came closest to optimality. Of the greedy heuristics the method called backwards heuristic performed best followed by the forwards heuristic, both of these give good solutions fast, as a result these heuristics could be used as part of a more complex algorithm such as the pruned dynamic programming method discussed in the pruned dynamic programming subsection of section 5.2. Of the population based methods, the genetic algorithm (GA) performed the best but still worse that the backwards heuristic. The methods highlighted in bold are solution methods that have desirable properties in terms of either objective value and speed or both. From these results more complex algorithms can be built using quality building blocks. For example, one possibility is to begin a hill climbing algorithm seeded with the backwards heuristic method. The fastest method of solution (solution times include repeat simulations for each of 20 problem instances) was the even distribution heuristic which allocated reserve crew evenly across departures, however, the objective value attained via this method was the worst.

Observations

Table 4 indicates that the topology of the solution space for the probabilistic crew absence model has many local optima that are close to the global optimum solution. Table 4 also shows that there is a clear negative correlation between expected reserve crew utilisation and cancellation rate and the utilisation and cancellation rates achieved in simulation, this indicates that the probabilistic objective function approach gives results that will perform as expected over a large number of trials. The strong correlation between expected and simulation results is a result of the law of large numbers (averages) and says nothing about the full range of possible behaviours, that is, this method does not guarantee a minimum worst case performance, it only maximises the average expected performance. The results also demonstrate that the step from good solutions to optimal solutions requires either long computation times or intelligent search techniques and that a short cut to optimal solutions is yet to be found for this particular model.

6 Conclusion

6.1 Main findings

We have introduced a probabilistic reserve crew scheduling model based on departures from a single airport and the mathematics of this model have been introduced. Through an investigation of possible objective functions for the model it was found that the sum of P objective function leads to reserve crew schedules with the most desirable properties (lowest cancellation rate and highest reserve crew utilisation rate). A multiple objective optimisation approach was considered where the number of reserve crew to assign is a variable along with the corresponding reserve crew schedule itself. Then, using the best objective function for the single objective model, an investigation of solution methods was carried out because enumeration typically takes a very long time. We found that with a pruned dynamic programming based algorithm it was possible to obtain optimal solutions, although this result is not guaranteed to happen because this method uses heuristics to estimate upper and lower bounds of partial solutions in order to prune the search tree. The main problem is that the heuristic bounds may not correctly reflect the potential quality of a partial solution, which introduces a risk of pruning partial solutions corresponding to optimal solutions. Local search techniques tended to perform very well compared to population based algorithms, with tabu search obtaining solutions very close to optimality. A constructive heuristic (backwards heuristic) was found to obtain good solutions very rapidly and as a result will be considered for use in the pruned dynamic programming approach as a lower bound estimation heuristic in order to improve the reliability of the pruned dynamic programming method.

6.2 Future directions

The current work can be made more detailed in several ways, including: a real time framework, multiple crew absences possible per departure, inclusion of the effects of disruptions other than crew absence, and other recovery actions. A real time framework would allow determination of whether reserve crew can feasibly begin and end flight duties within their duty period. A real time framework can easily be incorporated by making duty length (number of departures covered by a duty) a function of the departure number, this would require some preprocessing. Allowing the number of crew absent per departure to vary is a step towards a realistic model because airlines always require more than one crew member per flight, in fact they always require several crew of varying ranks [8]. This model extension would lead to the vector P

being replaced by a matrix describing the probabilities of different numbers of crew being absent for each departure (note that the model described in this paper could be used to schedule teams of reserve crew rather than single crew members).

The inclusion of delayed flights would make the problem more complicated because there is then the chance that reserve crew might become unavailable for covering a flight if it is delayed for too long.

Including other recovery actions such as aircraft swaps, crew swaps and delays means that costly reserve crew may not even be required if a cheaper recovery action is feasible.

Another area for improvement in the current model is to emphasise the feasibility of the crew schedule.

All of the areas for improvement suggested in this section are based on considering the feasibility of crew operations and airline operations in general. The current approach of using uniform random numbers as probabilities of crew absences leads to unrealistically high cancellation rates (tables 3 and 4), real data could be used to derive realistic probability values, which would have to be representative of the actual probabilities of absence.

One of the underlying assumptions of the model is that reserve crew can cover any of the departures occurring within their duty period (provided they have not already been used), in reality this will not always be the case because some departures might represent the beginning of long pairings that encroach upon the reserve crew's scheduled time off (or subsequent schedule [7]). To account for this, some preprocessing could be performed to find which departures can be feasibly considered for each individual reserve crew member.

References

1 Shervin AhmadBeygi, Amy Cohn, and Yihan Guan. Analysis of the potential for delay propagation in passenger airline networks. *Journal of air transport management*, 2008.

2 Micheal Ball, Cynthia Barnhart, George Nemhauser, and Amedeo Odoni. Air transportation: Irregular operations and control, 2007.

3 Cynthia Barnhart, Peter Belobaba, and Amadeo R Odoni. Applications of operations research in the airline industry. *Transportation Science*, 2003.

4 Cynthia Barnhart, Amy M. Cohn, Diego Klabjan Ellis L. Johnson, George L. Nemhauser, and Pamela H. Vance. Airline crew scheduling. *Handbook and Transportation Science*, 1999.

5 Ram Gopalan and Kalyan T. Talluri. Mathematical models in airline schedule planning. *Annals Of Operations Research*, 1998.

6 Diego Klabjan. Large-scale models in the airline industry. *Business and Economics journal*, 2005.

7 KLM. Private communication.

8 Marc Paelinck. KLM cabin crew reserve duty optimisation. In *Agifors proceedings*, 2001.

9 Sergey Shebalov and Diego Klabjan. Robust airline crew pairing: Move–up crews. *Transportation science*, 2006.