Modifications to the Pure Preferential Bidding System for Rostering Pilots at Air New Zealand

Natalie M. Rouse
Department of Engineering Science
University of Auckland
New Zealand
natterlie@hotmail.com

Abstract

At Air New Zealand, flight crew are assigned to particular flights by complex computer-based rostering systems. One of these crew rostering systems is called the Preferential Bidding by Seniority (PBS) system. This system constructs rosters for international pilots by accepting bids for certain work structures from the crew and performing a series of individual optimizations (called the "squeeze" process) which ensures that bids are satisfied in strict order of seniority. While the PBS system produces excellent quality rosters in strict seniority order, the amount of bid satisfaction typically achieved by junior crew members is very low. In addition, the series of optimizations are slow and expensive to solve.

We have demonstrated that it is possible to avoid the squeeze process and improve roster quality for junior crew while maintaining an acceptable degree of the strict seniority enforced by PBS. Improvements have also been made to the branching strategy used by the Branch and Bound algorithm to produce sensible integer solutions. A representative group of Air New Zealand pilots was consulted during the course of the project to ensure that a solution was developed that could be approved by the pilots and implemented as an improvement to the current PBS system.

1.0 Report Background

1.1 The Crew Scheduling Problem at Air New Zealand

Air New Zealand is New Zealand's national carrier, with around 5,000 employees in New Zealand, and a large fleet of Boeing 737's, 767's, and 747-400's. Each aircraft type requires a different complement of crew, and operates under different sets of rules and regulations. The two types of crew required on every flight are the cabin crew (or flight attendants) and the technical crew (or pilots). There are different ranks associated with each type of crew. The technical crew ranks are Captain, First Officer and Second Officer. Each aircraft or fleet type has crew grouped in these ranks who can operate only on this specific fleet type. The notation for these specific rank and fleet groups are as follows – "c" denotes a Captain, "f" a First Officer, and "s" a Second Officer. A B747-400 fleet type is indicated by 44, and a B767 fleet type by 6. For example, the second officers on the B747-400 are denoted by s44. When considering the crew scheduling problem, each type of crew scheduling problem is divided into National (domestic) and International. This paper has focused solely on an International Technical Crew scheduling problem.

Airline crew scheduling problems comprise of two processes – planning and rostering. The planning process generates a set of generic tours of duty (TODs) which cover all relevant flights, and the rostering process allocates the TODs generated by the planning process to an individual crew member. The TODs are generated so that each TOD begins and ends at a crew base (Air New Zealand has crew bases in Auckland, Wellington and Christchurch), and is comprised of an alternating sequence of duty periods and rests. The allocation of these TODs must ensure that all flights are crewed with the correct complement of crew, and that each crew member has a legal, feasible sequence of flights, termed a line of work (or LoW) over the given roster period. An international roster period at Air New Zealand covers 28 days. [1] The notation for the fourth roster period in the year 2001 is p0401.

Air New Zealand has been at the forefront of the airline industry quest for utilizing optimization processes to solve their crew scheduling problems for the past 16 years, and today the company is totally dependent on optimization-based computer systems to solve all of their crew scheduling problems. There are optimization-based computer systems to solve the problems of National planning, National rostering, International Flight attendant planning and rostering, and International Technical crew planning and rostering. [1]

The International Technical crew rostering system contains a solution method for PBS, or Preferential Bidding by Seniority, which is the method of rostering pilots at most airlines worldwide. PBS involves the pilots submitting bids for particular activities to be included in their LoW for the coming roster period (the method is discussed in more detail in Section 1.2). The solution method developed by Megan Thornley during 1992 and 1993 is unique to Air New Zealand, as it has the unique property that the maximum number of bids are guaranteed to be satisfied in strict seniority order, where heuristic methods used by other airlines cannot guarantee this. [1]

This project was primarily focused on improving the overall quality of the rosters produced by the PBS solution method incorporated in the International Technical crew rostering system.

1.2 PBS Rostering

The current method used by Air New Zealand to solve the international pilot rostering problem is based on the preferential bidding by seniority (PBS) system used by most airlines around the world. Under PBS, all of the crew members bid for certain "structure" of work they would like to have included or excluded from their LoW in the coming roster period. This could be tours of duty (certain trips), days off, or targeted flying hours. These bids are entered into the system, which then proceeds to build rosters by satisfying as many bids as possible given that crew members are considered in strict seniority order. The part of the PBS system which is unique to Air New Zealand is termed the "squeeze procedure", which is an optimisation process. This procedure maximizes the satisfaction of bids for an individual crew member without reducing the bid satisfaction of those crew members with higher rank, and without consideration of bids for those crew members with lower rank. The squeeze process guarantees that the maximum number of bids can be satisfied in order of seniority. This can be thought of as every crew member being considered to be infinitely more important in the problem than all crew members below him or her in the rank. The algorithm for the PBS system is detailed in the flowchart given in Figure 1.

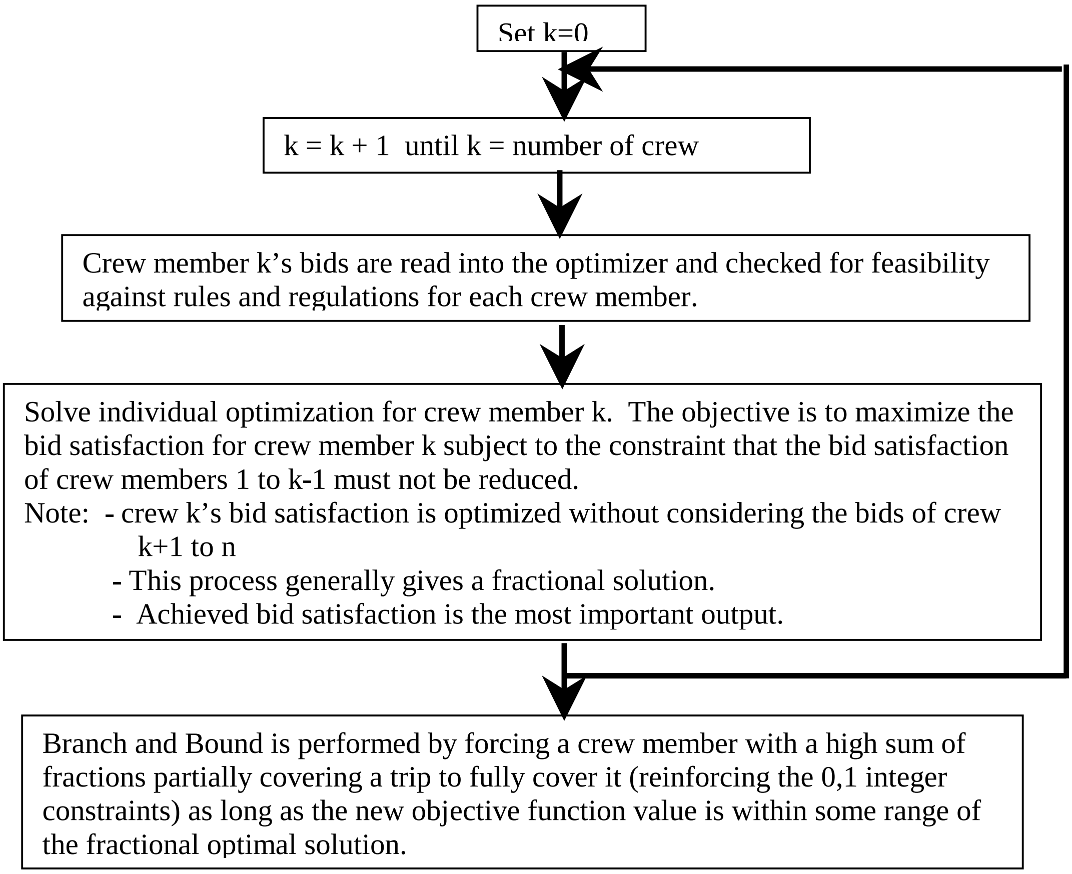

Figure 1: PBS Algorithm

The PBS system was developed during 1992 and 1993 by Megan Thornley [2] under the supervision and direction of Professor David Ryan. Over the last nine years, the current PBS system has been developed and maintained by the Operations Research group at Air New Zealand in conjunction with Engineering Science students and staff at the University of Auckland.

The squeeze process violates bids of more junior crew in order to satisfy the bids of more senior crew members. The problem becomes very constrained or fixed quite quickly, meaning that the more junior crew members have very little chance of achieving satisfaction of any of their bids. This leads to highly inequitable rosters where the senior crew members almost always achieve all of the structure they bid for, and the junior crew members are assigned to whatever is left over when they are considered in the squeeze process. Figure 2 shows a graph of average bid satisfaction of a sample of crew under pure PBS grouped into small groups in decreasing seniority order.

This plot highlights another problem with the PBS system – the PBS bidding strategy. Under PBS, the most senior crew have all of their feasible bids satisfied before the bids of any other crew member are considered. This means that they are entitled to submit as many bids as they wish, and they will all be satisfied (where possible), before any more junior crew member is considered. The junior crew, therefore, have no chance for achieving satisfaction of a bid for a popular trip or day off. To this end, the junior crew members have taken to bidding not for duties that they would like to get, as things like popular trips are virtually unattainable to them, but for activities such as "days off" which they are more likely to get.

Figure 3 shows an average bid satisfaction graph based upon data obtained from executing pure PBS on a suite of bids created by the Pilot Scheduling Committee (PSC

– see Section 1.3) based upon the actual bids placed for the same rank and period as the graph in Figure 2. These bids were altered to reflect all crew members bidding for what they would like to do, which was necessary to predict the bidding strategy crew may adopt if they feel they have better chances of achieving any bids placed. The graph shows that under pure PBS, if every crew member bids for what they would like to achieve, the most junior crew do not have a chance of ever achieving any of their bids.

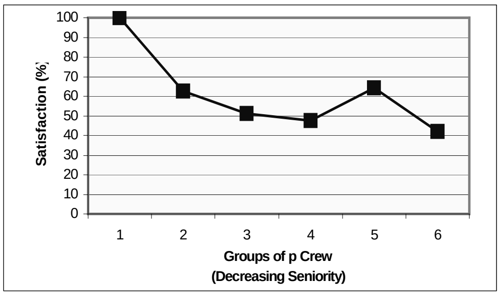

Figure 2: Average Bid Satisfaction achieved by pure PBS (PBS Bids)

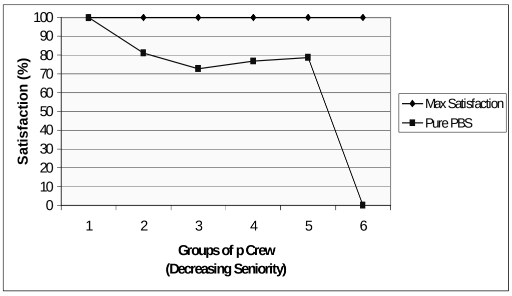

Figure 3: Average Bid Satisfaction achieved by pure PBS (PSC Bids)

Another problem with the PBS system and the squeeze process is the execution or roster build time. The squeeze procedure performs a separate optimization for each crew member in seniority order to ensure maximum bid satisfaction in strict seniority order. This sequence of separate optimizations is very expensive to perform both computationally and time-wise. Avoiding the squeeze process by performing a single optimization for all crew would significantly reduce the build time of each roster, enabling rosters to be rebuilt if necessary without serious disruption to operations.

1.3 Project Outline

The main objective of my project was to identify, develop and test different methods for avoiding the squeeze process and improving the roster quality for the more junior crew without serious reduction of quality for the senior crew. It was anticipated that the improved method would consist of a reduced level of squeeze, which would involve stopping the sequence of optimizations after a certain number of crew, and performing one optimization for the remaining crew. Seniority would be maintained by applying some weighting function to the crew remaining in the final optimization to ensure that bids entered by senior crew carried more weight than bids placed by more junior crew.

A major objective for my project was to reduce the time taken to generate good PBS solutions. This was achievable through avoiding squeeze, as the time taken for a single optimization for p crew is far less than the time for p small optimizations individually.

It was important to collaborate closely with the Pilot Scheduling Committee to produce solutions desirable to the pilots. The PSC is a group responsible for representing all pilots in matters concerning rostering decisions and quality. The direction of my project was continually influenced by recommendations from the PSC.

My project was also to encompass making improvements to the Branch and Bound algorithm. The performance of this algorithm under pure PBS conditions with full squeeze has not been very important, as the squeeze process produces solutions very close to integer. Thus the Branch and Bound algorithm has very limited capacity for altering the solution, and integerizing the solution is a fast process. However, when the squeeze process is removed, the solution space for Branch and Bound to explore is much larger. This means that the way in which Branch and Bound operates becomes much more important. The current branching process is based on a constraint branch in which the crew member closest to covering a particular trip is forced to fully cover it. It was believed that a smarter branching strategy may be necessary when the squeeze process was reduced and the solution space consequentially expanded.

The means by which these objectives were to be achieved were through creating scenarios and executing these on actual bid sets generated for selected crew ranks and roster periods. These different scenarios were to include different levels of squeeze, and different weighting functions for the unsqueezed group. The results were to be shown on a spreadsheet table and graph to enable results to be easily interpreted and compared with others.

2.0 The Weighting Functions

It was decided in conjunction with the Pilot Scheduling Committee that the weighting scheme should have a fixed top value of 100, and the bottom value should be variable with the rest of the crew being weighted linearly in between. The PSC suggested a factor of 4 between the top and bottom crew members, that is a weighting scheme from 25 to 100. The slope for the linear weighting function for this was:

$$\text{Slope} = \frac{(1-n/100)}{(\text{number of crew} - 1)} \qquad (1)$$

where n is the desired minimum weighting value. The linear weighting function then used this slope to determine linearly spaced weights:

$$\text{Weight}(i) = 1 + \text{slope}*(1 - i) \qquad (2)$$

It was decided that as well as linear weighting functions, other weighting functions should be investigated too, such as exponential weighting functions. The approximation agreed upon is represented by the equation below:

$$\text{Weight}(i) = \frac{\text{Linear Weight}(i)}{(\text{cube root}(i))} \qquad (3)$$

This weighting scheme yielded weighted values much harsher than the linear weights, but also much gentler than the true exponential case. The weighting curve looked like a curve at the top, and then a sloped linear function at the bottom, which was the desired outcome. This ensured that there was enough variation in the weighting values so that even the junior crew members were treated differently in the problem.

3.0 The Report Generator

3.1 The Percentage Bid Satisfaction Reports

An Excel tool was created with the purpose of present meaningful reporting of results. This included a table of average bid satisfaction results and a plot of these values. For each scenario (Pure PBS, Equal Weights, Linear Weights and Exponential Weights), the run time, average percent satisfaction for each group of p crew, (p a user-defined size), and the overall average percentage satisfaction of the rank were given in the table. The crew were separated into a squeezed group and equitable groups of size p in the table, to make it easier to see where the squeeze was stopped.

For most ranks, pure PBS achieves higher bid satisfaction (marginally) for the first group of crew after the squeeze, but in all other groups there is a considerable gain in satisfaction achievement for both the linear and equal weights scenarios. Pure PBS achieves much higher (and much closer to the other two schemes) percent satisfaction for the most junior due to the bidding strategy of the junior pilots under the pure PBS scheme. As these pilots know that under PBS, the likelihood of them achieving their bids is low, they tend to bid for more easily achievable structure.

3.2 The Linear and Exponential Weighted Satisfaction Reports

These reports contained the weighting function shown as a line representing the maximum achievable satisfaction for the crew, with the satisfaction for each crew member as a function of the weight applied to that person in the optimization. The results table for these reports contained the averaged weights, the averaged satisfaction values of the linear or exponential scenario (as a function of the weighting scheme), and the percentage achieved of maximum achievable satisfaction (or percent difference).

The most equitable weighting function is one with each crew member achieving a similar percentage of their achievable bid satisfaction. This would be shown by a horizontal percentage of achievable bid satisfaction line (shown in Figure 4), for example if all crew members were to achieve 90% of achievable satisfaction. In contrast to the case given, another possible weighting scheme could cause all crew members to achieve a similar amount of their bid satisfaction (shown in Figure 5), for example if all crew members achieve 5 units of satisfaction less than their maximum achievable. This case represents too much weight on the senior crew and not enough on the junior members, and the line of percent of achievable satisfaction would have a negative gradient. The converse of this is represented by Figure 6, where there is not enough weight on the senior crew. This is apparent because the junior crew achieve a higher percentage of their achievable satisfaction.

This shows that the line of % bid satisfaction will indicate whether a weighting function is too severe, just equitable, or perhaps not strong enough. The desired result is a Class 0, which represents an equitable weighting scheme that still reflects seniority.

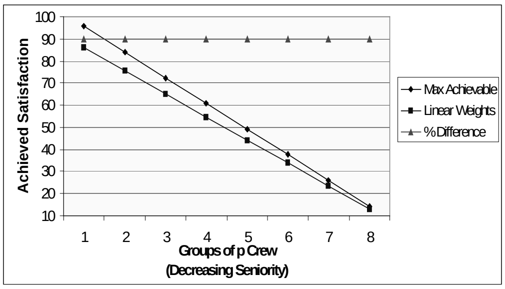

Figure 4: Class 0 Linear Weighted Satisfaction Graph (10-100)

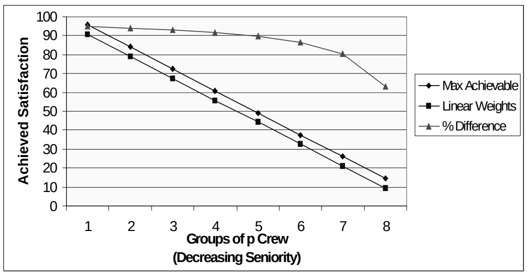

Figure 5: Class 1 Linear Weighted Satisfaction Graph (10-100)

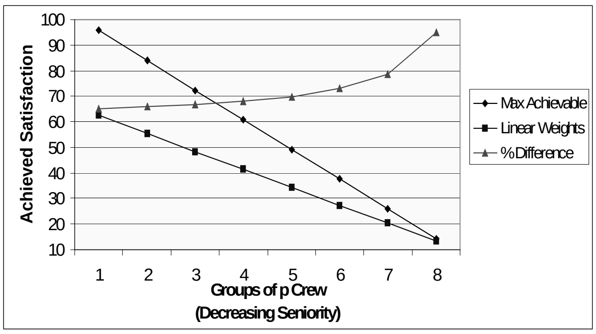

Figure 6: Class 2 Linear Weighted Satisfaction Graph (10-100)

However, in practice, because the results were both functions of the weighting scheme and the bids placed by the crew, a straight line was not easily achievable, and there were often oscillations in this line. For example, a group of crew bidding for pieces of work that were more easily achieved than another group would achieve higher satisfaction. Thus, the aim is to achieve a horizontal best fit of the percentage line.

4.0 Testing The Scenarios

The 0% squeeze case with equal weights would represent the maximum achievable average satisfaction and the minimum objective function value, as this scenario should find the best possible solution for every rank without seniority biases or separate optimizations to hinder it.

4.1 Variable Squeeze

The initial scenarios tested were with varying levels of squeeze, and equal weights for the unsqueezed group in each run. Runs of 0% squeeze (fully equitable), 25% squeeze, 50% squeeze, 75% squeeze and 100% squeeze (pure PBS) were performed. These scenarios with variable squeeze levels and equal weights for the unsqueezed crew showed decreases in time and objective function values from the pure PBS method, varying according to the level of squeeze set.

As the level of squeeze increased, and the scenario grew more and more like pure PBS, the problem became increasingly more constrained. In fact, a 75% and often even 50% squeeze level removed all of the flexibility out of the rank, and the solution was exactly the same as pure PBS (100% squeeze).

4.2 Variable Weighting Schemes

No (0%) squeeze scenarios were executed with exponential weighting schemes and linear weighting schemes. For some of the larger ranks, the exponentially weighted scenarios required a larger Branch and Bound window (solution space) in order to find their integer solutions. The linear weighting schemes all had a maximum weight of 100, and a variable minimum weight. All ranks were tested with a minimum weight of 25, 20, and 1, and the s6 p0901 original and created bid sets were also tested with a minimum weight of 10. The c44 p0401 rank also needed to be tested with minimum weights of 30 and even 35, until the results displayed enough weight on the junior crew. The reports generated from the results of these scenarios were evaluated against the three class scenario examples (See Figures 4, 5 and 6) to determine whether there was too much, just the right amount, or not enough weight on the junior crew.

4.3 The Branch and Bound Algorithm

A large jump in the objective function values from the optimal LP solution to an integer solution indicated that the bound gap used in Branch and Bound was too large for the more flexible problems yielded by scenarios with small amounts of squeeze. With the bound gap set at 15% for all scenarios, the algorithm accepted an integer solution that was much worse or higher than necessary. When the bound gap was reset to 5%, a much better integer solution was identified.

The Branch and Bound algorithm was operating on an equitability branching scheme (EB), where the next branch selected was the highest sum of fractions out of the entire fractional crew, with no regard to seniority. While this method worked for pure PBS and the strict seniority imposed by the squeeze, it was unacceptable with a more equitable optimization process. My task then was to impose a seniority branching strategy, where the best senior (fractional) branch was then selected out of the five most senior fractional crew (SB).

4.4 The Pilot Scheduling Committee

For the duration of my project, the approval of the PSC was a paramount aim. Results and progress of this project were discussed with the PSC at every step. Suggestions were made along with the results, shaping the direction of the project.

The PSC had the final say in the decision on weighting schemes, with the general consensus being that a weighting factor of 4 between the most senior and most junior crew members was a reasonable linear gradient to test (that is, a weighting scheme from 25-100). It was also suggested that this factor may be varied in order to identify the best gradient. The PSC then undertook to take a suite of PBS bids, and alter them to represent a bid suite that they would anticipate under a linear, unsqueezed system. This was done on the bids of the s6 p0901 rank and period.

It was also required that actual rosters be prepared for the results of the scenarios run on both suites of bids (including pure PBS) for the s6 p0901 rank and period, to allow more meaningful analysis of what the new schemes would mean for individual pilots.

5.0 Results

5.1 Partial Squeeze Scenarios

The results of testing the partial squeeze scenarios showed that squeeze levels of 50% and over imposed an increasing lack of flexibility on the problems, and tended to produce solutions very close to and sometimes equal to those generated by pure PBS. It was then decided that 0% and 25% partial squeeze scenarios were the most equitable alternatives.

In general, large increases in average bid satisfaction for junior crew members was observed under 0% squeeze, with very small and sometimes no reduction in bid satisfaction for the most senior crew members.

Bid Satisfaction Report (0% Squeeze)

Rank/Period: s6 p0301 Crew Size: 25 p (>2): 3

OPTION	Time (s)	Average Bid Satisfaction (groups of p crew) Equitable									Overall Average Satisfaction
		3	3	3	3	3	3	3	3	1	
Pure PBS	2233	83.503333	92.593	94.6867	73.5833	100	93.94	54.54667	45.46	54.55	78.7796
Equal Weights (SB)	469	88.89	85.577	83.94	94.84	99.66333	83.84	85.18667	88.47333	100	89.2492

Figure 8: Bid Satisfaction Report for s6 p0301 with 0% Squeeze

5.2 Seniority Branching Scheme

After comparing results from the EB and SB schemes, it was obvious that the SB scheme was selecting more sensible branches. Results indicated that both branching strategies generate solutions of a similar overall quality, but the SB scheme allows the spread of the average bid satisfaction to be biased more towards the more senior crew members, which is sensible in this environment.

5.3 Weighting Schemes

Weighting schemes were applied to all 0% squeeze scenarios for all crew ranks. These weighting schemes included exponential crew weights and linear crew weights between 1-35 and 100 according to weighting function and crew rank size.

The exponential weighting scheme generated solutions very similar to pure PBS in most cases, while the linear weighting scheme generated solutions generally between the exponential and equal weighting strategies. It was then decided that the exponential weighting scheme placed too much weight on the more senior crew in the problem, and generated solutions too close to pure PBS to be considered equitable. The linear weighting scheme was found to be the most equitable scheme while still maintaining a degree of seniority. This degree of seniority could now be determined by varying the gradient of the linear weighting function.

The linear scheme was tested on all crew ranks with minimum weights of 1, 20 and 25. Figures 9 and 10 respectively show the effect of changing the value of the minimum weight from 1 to 25 for the s6 p0401 rank. The "% Difference" line indicates that the 25-100 weighting scheme is more equitable because all crew members are achieving very similar percentages of their maximum achievable bid satisfaction. Figure 9 shows that this is not the case with the 1-100 scheme. The minimum weight should be selected for each rank individually, and range from between 20 and 35.

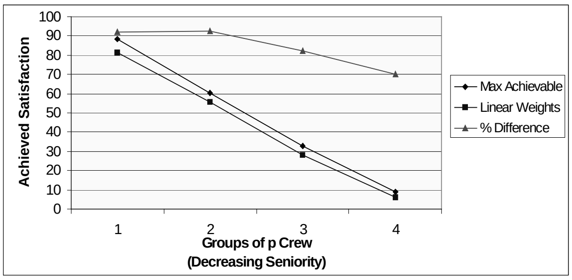

Figure 9: Linear (1-100) Weighted Bid Satisfaction Report s6 p0401

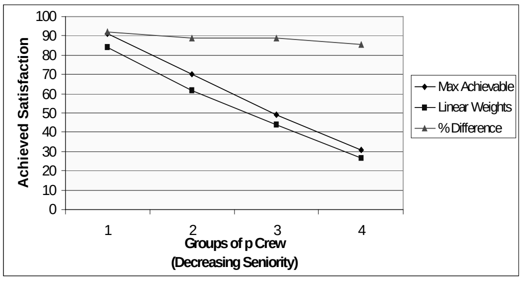

Figure 10: Linear (25-100) Weighted Bid Satisfaction Report s6 p0401

7.0 Conclusions

The main conclusion reached is that it is possible to develop a PBS system which produces good quality solutions while maintaining seniority without performing a separate optimization for every crew member.

7.1 Final Scheme

The best rostering scheme identified a 0% squeeze scheme, with all crew considered in a single optimization. Seniority is maintained in this scheme by weighting or scaling crew bids evenly or linearly in the problem. The best weighting function is unique to each individual crew rank. All weighting functions allow the bids of the most senior crew member a weight of 100, and the bids of the most junior crew member a weight between 20 and 35. This minimum weight is determined for each rank. For example, in period 4, the best minimum weight for the s6 rank was 25, but the best minimum weight for the c44 rank was 35.

After the minimum weight value is chosen for each rank, the result should be checked every few periods to ensure that the weighting function is still producing equitable rosters for the rank. The results can be checked by generating a Weighted Bid Satisfaction Report for each period, and examining the percentage of achievable bid satisfaction line. If the best fit of this line is not horizontal, the minimum value can be adjusted accordingly, ensuring the bids continue to be scaled effectively.

7.2 Implementation

The Pilot Scheduling Committee have approved the final scheme, and agreed that this scheme is a major improvement for both the pilots and the company. The pilots would see large improvements in overall roster quality, and individually big increases in bid satisfaction for the majority of crew members with only small decreases in satisfaction for the most senior crew. All pilots periodically change rank, either to a higher or lower rank, where they must begin as the most junior crew member. Because of this, the PSC believe that most pilots would support a shift to a more equitable rostering scheme. The company would see large reductions in execution and roster build time under the proposed scheme. These time reductions would not only save expensive computation time but also allow rosters to be built later, allowing time for last minute changes to be incorporated, reducing the likelihood of expensive roster rebuilds.

Currently, the PSC is awaiting employment contract finalizations on the old PBS system, which must occur before any changes can be made. When these finalizations have been made, a vote will be put to all pilots with details of the proposed system.

7.3 Further Work

The branching strategy used by the Branch and Bound algorithm has been altered from an equitability branching scheme to a seniority branching scheme. However, further work must be done to identify a “smarter” branching strategy to ensure that the large solution space generated by the more flexible 0% squeeze solutions is explored in an intelligent way. The ideal branching strategy would examine the effect of imposing a certain branch before branching, and only branch where the effect of imposing the branch is not detrimental to the solution.

Acknowledgements

I would like to acknowledge and thank my supervisor, Professor David Ryan, who has been revolutionizing the airline industry with the development of these systems for the past 16 years. I would also like to thank Amanda Day and the Operations Research group at Air New Zealand, for the amount of time and effort they put into helping me understand the current system and implement the required changes.

References

[1] Ryan, D. M. 2000 “Optimization Earns its Wings” *ORMS Today: Special International Issue* vol. 27(2) 26-30

[2] M.E. Thornley, “Crew Rostering Under A Seniority Preferential Bidding Environment Using Column Generation,” Master’s Thesis, University of Auckland (1993)