

Optimized Crew Scheduling at Air New Zealand

E. ROD BUTCHERS

*Network Logistics, Air New Zealand Ltd.
Private Bag 92007, Auckland, New Zealand*

PAUL R. DAY

*Optimal Decision Technologies Ltd.
P.O. Box 128-233, Remuera, Auckland, New Zealand*

ANDREW P. GOLDIE

Optimal Decision Technologies Ltd.

STEPHEN MILLER

Network Logistics, Air New Zealand Ltd.

JEFF A. MEYER

Optimal Decision Technologies Ltd.

DAVID M. RYAN

*Department of Engineering Science,
University of Auckland, Private Bag 92019, Auckland,
New Zealand*

AMANDA C. SCOTT

Network Logistics, Air New Zealand Ltd.

CHRIS A. WALLACE

Network Logistics, Air New Zealand Ltd.

The aircrew-scheduling problem consists of two important subproblems: the tours-of-duty planning problem to generate minimum-cost tours of duty (sequences of duty periods and rest periods) to cover all scheduled flights, and the rostering problem to assign tours of duty to individual crew members. Between 1986 and 1999, Air New Zealand staff and consultants in collaboration with the University of Auckland have developed eight application-specific optimization-based computer systems to solve all aspects of the tours-of-duty planning and rostering processes for Air New Zealand's national and international operations. These systems have saved NZ\$15,655,000 per year while providing crew rosters that better respect crew members' preferences.

Commercial airlines must solve many resource-scheduling problems to ensure that aircraft and aircrews are available for all scheduled flights. Since aircraft and aircrews are among the most expensive of airline resources, their efficient utilization is important. Because of the large savings possible from using aircrews more

efficiently, many airlines have tried to develop optimization methods to solve their crew-scheduling problems. Most early optimization attempts failed because of inadequate solution methods and lack of computer power. Even now many airlines still use heuristic or manual methods to solve crew-scheduling problems. Since the

Airline scheduling process

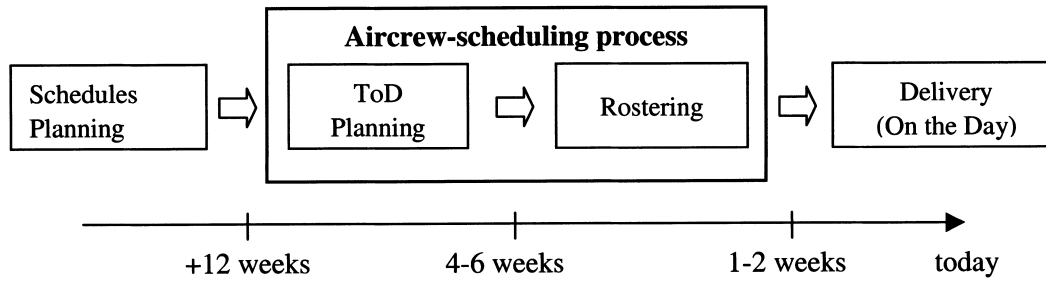


Figure 1: The aircrew-scheduling process consists of two distinct subproblems, the tours-of-duty problem and the rostering problem. The tours-of-duty (ToD) planning problem involves constructing the minimum-cost sequences of flights to crew the flight schedule. The ToDs from the ToD planning problem are passed to rostering four to six weeks before the start of the roster period. The rostering process involves allocating the planned ToDs to individual crew members to form a line of work over a roster period. Aircrew rosters are typically up to one month long and are normally published to crews one to two weeks before the start of the roster period.

mid-80s, improvements in optimization techniques and increasing computer power have resulted in the development of optimization-based solution methods for aircrew-scheduling problems.

The aircrew-scheduling problem is usually partitioned into two distinct subproblems (Figure 1).

The first, the tours-of-duty (ToD) planning problem, calls for constructing sequences of flights to crew the flight schedule, which are variously described as tours of duty, trips, pairings, or rotations. In some situations, these ToDs can be one-day periods of work, but they can also consist of sequences of flights and rest periods spanning many days. The ToD planning problem has prompted a great deal of research (Anbil, Forrest, and Pulleybank [1998] provide a useful summary), and over the past 20 years, airlines have regularly discussed the problem and various solution methods at meetings of the Airline Group of the International Federation

of Operations Research Societies (AGI-FORS). Aspects of the ToD planning problem for American Airlines were discussed by a previous Franz Edelman finalist [Anbil et al. 1991].

The second subproblem associated with aircrew scheduling is rostering. Rostering involves allocating the planned ToDs from the first subproblem to individual crew members to form a line of work (LoW) over the rostering period. In contrast to the extensive work on the ToD planning problem, little work has been reported on solution methods for rostering.

The two aircrew-scheduling subproblems can be considered from both management and crew points of view. Management seeks minimum-cost or maximum-productivity solutions that are legal in the sense that they do not break any rules or agreements with the crew and are feasible in the sense that all the work is performed. Management also seeks to quantify the costs associated with satisfy-

ing the rules and agreements. A further important feature of crew-scheduling solutions from the management point of view is their operational robustness or sensitivity to disruption of the planned flight schedule. For crews, the key issue is solution quality. Among all legal and feasible solutions, crews will prefer solutions that maximize some quality measure defined by the particular crew group. Typical quality measures are the fair distribution of work or satisfying the work preferences of individual crew members in seniority order or the avoidance of unduly arduous work patterns. While managers appreciate the importance of solution quality in promoting crew satisfaction, the quality objective remains the primary concern of crew members.

The Tours-of-Duty Planning Problem

A *tour of duty* (ToD) is an alternating sequence of duty periods and rest periods (or layovers). Each *duty period* comprises one or more flights and may also include passengering flights. A *passengering flight* is one on which a crew member travels as a passenger to get to a particular airport for a subsequent operating flight. The first

duty period of a ToD must start at a crew base, and the last duty period must end at the same crew base. An airline might have several bases (cities) at which crew members are domiciled. Each ToD will have an associated crew complement made up of a number of crew members of various ranks required to staff the ToD. For example, a crew complement might consist of one B767 captain and one first officer, or one inflight service director and three flight attendants (Figure 2).

In ToD planning, one constructs a minimum-cost set of generic ToDs that cover all flights in an airline schedule with the crew required (for example, captain, first officer, inflight service director, and flight attendants). Constructing ToDs does not include considering which individual crew members will perform the ToDs. Airlines that have stable and regular flight schedules can solve a standard daily or weekly representation of the flight schedule and then use that solution for the duration of the flight schedule. However, airlines with variable flight schedules must solve a fully dated problem, and the ToD solutions can differ from day to day or

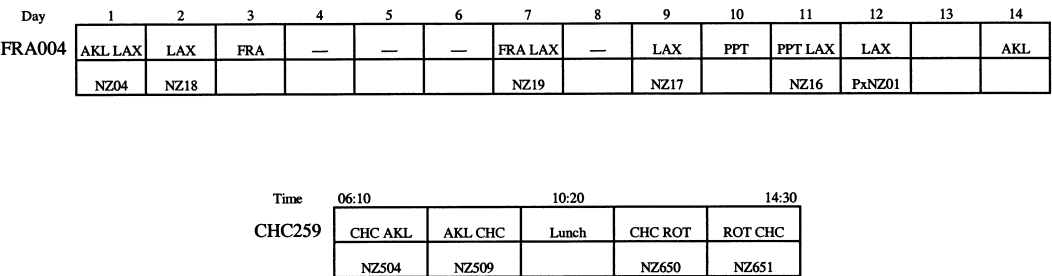


Figure 2: In these example tours of duty, FRA004 is a 14-day international tour of duty (ToD) for pilots and CHC259 is a single-day national ToD for flight attendants. The top row of the ToD indicates the airports visited (Auckland—AKL, Los Angeles—LAX, Frankfurt—FRA, Paapeete—PPT, Christchurch—CHC, Rotorua—ROT), and the bottom row shows the flight numbers. Px indicates a passengering flight and a dash indicates a day off overseas.

week to week.

The creation of legal duty periods and ToDs is governed by a complex set of rules. Some of these rules are specified by aviation regulatory authorities, such as the New Zealand Civil Aviation Authority or the Federal Aviation Administration. Other rules are specified by employment contracts and operational-robustness considerations. Typical duty period rules concern the length of a duty period, the latest time at which a flight may start within a duty period, the maximum allowable flight time within a duty period, and the number of crew members required to operate a flight. Typical ToD rules concern

Many airlines still use heuristic or manual methods.

the minimum rest between duty periods, cumulative duty-time and flight-time limits over a rolling time period (for example, seven days), and in international operations, the effects of time-zone changes or acclimatization at foreign ports. Operational robustness rules are sensible rules of thumb that minimize the impacts of disruptions on the day of operation. For example, providing rest periods slightly longer than the minimum legally required so that crews will still have their legal rest periods if their flights arrive late. Within Air New Zealand and generally in most airlines, the rules are different for pilots and for flight attendants, and for national and for international operations.

Airlines measure the quality of ToD solutions in terms of their dollar costs and measures of operational robustness and good practice. For example, good practice

dictates that airlines use the layover locations crew members prefer if this does not increase costs. The costs incurred in operating a flight schedule include the explicit costs of crew members being at foreign ports (hotel accommodation, meals, ground transport, and various other expenses and allowance entitlements), and the implicit crew costs of salaries, overtime, and other pay-related airline-specific metrics, such as flight time, trip credits, or incentive pay. Robustness quality measures might favor crew members continuing to operate on the same aircraft (to minimize delays during disruptions) or certain flight combinations over others. Airlines prefer solutions that are close to the minimum dollar-cost solution but have better robustness or good practice features.

A desirable feature of ToD solutions is that all crew members on a flight perform the same sequence of flights for as much of their duty period as possible and, even better, that they also stay with the same aircraft. This greatly reduces the propagation of disruptions from one flight to other flights on the day of operation. This is called *unit crewing*. ToD solutions with a high degree of unit crewing generally cost more than the minimum-cost ToD solution but are much more robust. Airlines must find some balance between cheaper solutions with less unit crewing and more expensive solutions with more unit crewing. The unit-crewing idea is not normally applied across crew types because pilots and flight attendants work under different rules and regulations that preclude their operating exactly the same duty periods.

| Day | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 | 24 | 25 | 26 | 27 | 28 | |
|--------------|-----|---|---|---|--------|---|---|---|---|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|--------|----|----|--|
| Crew PILOT01 | SIM | O | O | O | FRA004 | | | | | | | | | | | | | | O | O | O | O | O | O | O | SYD008 | | | |

| Day | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 |
|-----------|-----|-----|---|---|-----|------|------|-----|-----|----|-----|-----|----|----|
| Crew FA01 | 259 | 279 | O | O | 324 | CALL | CALL | 370 | 385 | O | 418 | 436 | O | O |

Figure 3: In these example lines of work, PILOT01 is a 28-day international pilot line of work (LoW) and FA01 is a 14-day national flight-attendant LoW. The LoW activities include a simulator training session (SIM), a Frankfurt ToD (FRA004), days off (O), a Sydney ToD (SYD008), on-call days (CALL), and various national ToDs (three-digit codes).

The Rostering Problem

The rostering problem consists of constructing a line of work (LoW) for each crew member. A LoW is a sequence of roster activities and rest periods over a specified roster period, such as 14 days or 28 days duration. Roster activities include ToDs, training, leave, days off, and call duties (Figure 3).

The rostering process consists of assigning all roster activities to crew members over the roster period, with the following requirements:

- The roster assignment must be legal; that is, it must satisfy all rostering rules, including legislative rules, employment-contract conditions, operational or in-house rostering agreements and best practice.
- The roster must be feasible; that is, all roster activities to be assigned must be assigned to crew members of the correct rank. In particular, no flying ToDs should be left unassigned in a feasible solution.
- The roster quality should be maximized; that is, the best possible roster for the crew should be obtained for any agreed measure of roster quality.

The creation of legal LoWs is governed by a complex set of rules derived from legislation, crew-employment contracts,

and operational-robustness considerations. Typical rules concern the number of days off per roster period, cumulative duty time and flight time limits over a rolling time period (for example, 24 hours, seven days, or 28 days), the rest period between ToDs, and the qualifications required to operate a particular ToD. Pilots generally have more complex operating experience requirements than flight attendants.

The earliest rostering approaches the airline industry used were based on bid-lines in which airlines constructed legal LoWs to produce feasible rosters without taking any account of the individual members of each crew rank. Crew members then bid for the LoWs they preferred. The airline then processed crew members in seniority order and assigned each person his or her most preferred LoW of those remaining unassigned. Many rostering inefficiencies occurred because the airlines constructed LoWs without taking into account crew availability, determined by annual leave, training, or carry-in activities from the previous period. This resulted in ToDs remaining unassigned at the end of the rostering process. It is now recognized that assignment-based rostering systems, in which roster activities are assigned to specific crew members, respecting their

carry-in activities and other existing assignments, are likely to produce much more efficient rosters. However, many airlines still use bidline systems because of long-standing industrial agreements.

Two major variants of assignment-based rostering systems have evolved. Preferential bidding by seniority (PBS) is based on satisfying crew bids in strict seniority order. A crew bid is an expression of preference for work content or days off in a LoW. The seniority order results in long-employed crew members usually achieving all or most of their bids, while more-recently-hired crew members are often unable to influence the content of their LoWs. Equitability assignment, on the other hand, is designed to fairly and evenly spread satisfaction across all members of a crew rank, where the satisfaction measures the achievement of crew member's individual bids and various agreed collective quality measures.

Air New Zealand

Air New Zealand is the largest national and international airline based in New Zealand. It employs over 2,000 crew members and operates flights to Australia, Asia, North America, and Europe, and between the major centers within New Zealand.

Air New Zealand has two closely integrated business units, National and International. Air New Zealand National operates eight B737-200 and nine B737-300 aircraft on national flights and longer flights to Australia and the Pacific Islands. Air New Zealand International operates 13 B767 and eight B747 aircraft on predominantly international operations. The Air New Zealand flight schedule is based on

an interconnected network, rather than a hub-and-spoke network like those operated by many American airlines.

Air New Zealand has four major crew types: international pilots, international flight attendants, national pilots, and national flight attendants. These different crew types have different employment contracts, different scheduling rules, and different pay schemes. Each crew type can also be partitioned into a number of ranks. For example, international flight attendants have the ranks of inflight service director, inflight service coordinator, flight attendant premium service, and flight attendant Pacific class. An unusual industrial situation within Air New Zealand is that pilots within the same fleet and rank

An optimization model can reliably detect infeasibility.

can belong to different contract groups and therefore have different scheduling rules. This causes additional complexity in formulating and solving crew-scheduling problems.

The first contact with Air New Zealand occurred in 1983, when David Ryan met with senior Air New Zealand managers to discuss their crew-scheduling problems. Despite a somewhat sceptical initial response from management, the first student project in 1984, investigating a small part of the national-tours-of-duty-planning problem, produced impressive results, and so began the collaboration on which this paper is based, and which continues today. Professor Ryan and postgraduate students at the University of Auckland have developed the prototypes of the optimiza-

tion systems implemented at Air New Zealand. In most cases, students were subsequently employed by Air New Zealand to extend their research work. By 1987, Air New Zealand had hired two OR graduates, and today it employs 11 graduates or postgraduates, with a further five having made significant contributions during the 1990s. In 1997, five of the early collaborators formed a company called Optimal Decision Technologies, which now provides much of the support and development of the systems. We have developed three ToD optimizers:

- (1) for national pilots and flight attendants,
- (2) for international pilots, and
- (3) for international flight attendants.

These ToD optimizers each have special capabilities reflecting the route network and business rules for that problem. We have also developed four rostering optimizers:

- (1) for national pilots,
- (2) for national flight attendants,
- (3) for international pilots, and
- (4) for international flight attendants.

These rostering optimizers implement different rostering methods and use different roster quality measures, which we developed and agreed upon with each crew group. The Air New Zealand rostering problems include both PBS for international pilots and equitability assignment for all other crew types. We developed separate solvers because of fundamental differences in the business problems or in the problem behaviour.

Models and Solution Methods for Aircrew Scheduling

The set-partitioning model provides an

underlying mathematical model for both the ToD planning and the rostering subproblems of the aircrew-scheduling problem. The set-partitioning problem (SPP) is a specially structured zero-one integer linear program with the form

$$\begin{array}{ll} \text{SPP:} & \text{minimize } z = \mathbf{c}^T \mathbf{x} \\ & \text{subject to } A\mathbf{x} = \mathbf{e} \\ & \text{and } \mathbf{x} \in \{0,1\}^n \end{array}$$

where $\mathbf{e} = (1,1,1,\dots,1)^T$ and A is a matrix of zeros and ones. Because of the computational difficulties in solving very large and practical instances of the set-partitioning problem, many early attempts to use optimization solution methods to solve aircrew-scheduling problems, and in particular the ToD planning problem, were unsuccessful, and researchers resorted to a variety of heuristic solution methods.

While most heuristic methods are fairly easy to implement and may have reasonably inexpensive computer resource requirements, they suffer from two major disadvantages: (1) they can provide no bound on the quality of any feasible solution that they produce, and (2) they are unable to guarantee a feasible solution will be found even if one exists. The heuristic method may fail to find a feasible solution either because the heuristic method is inadequate or because the problem is truly infeasible. It is important to distinguish between these two possibilities, particularly in the rostering problem, which can sometimes be infeasible because of insufficient crew. In contrast, an optimization model can reliably detect infeasibility. During the past two decades, the development of optimization methods and techniques for the solution of set-partitioning

problems and the increase in computer power has meant that realistic-sized models that arise in aircrew-scheduling problems can be solved.

The Tours-of-Duty-Planning Model

In the basic ToD-planning model, each column or variable in the SPP corresponds to one possible ToD that could be flown by some crew member. Each constraint in the SPP corresponds to a particular flight and ensures that the flight is included in exactly one ToD. The elements of the A matrix can then be defined as

$$a_{ij} = \begin{cases} 1 & \text{if the } j\text{th ToD (variable)} \\ & \text{includes the } i\text{th flight (constraint),} \\ 0 & \text{otherwise.} \end{cases}$$

The value of c_j , the cost of variable j , reflects the dollar cost of operating the j th ToD. The calculation of c_j values is specified by the particular problem being considered but usually includes the cost of paid hours (both productive and unproductive), ground transport, meals, and accommodation, and the cost of passengering crew within the ToD. Many authors [Rubin 1973; Wedelin 1995] model the ToD-planning problem using the set-covering formulation in which the equality constraints are replaced by greater-than or-equals constraints. The overcover of a flight permitted by the set-covering constraint can be interpreted as passengering of the excess crew cover. This formulation results in fewer variables, but it has the major disadvantage of making it difficult to model accurately the rules and costs associated with passengering crew. For example, duty-time limits for passengering duties are generally longer than those for operating duties. This cannot be correctly

modeled using set-covering constraints. In the Air New Zealand applications, we used set-partitioning constraints, which allow us to accurately model passengering. This results in additional columns that explicitly include passengering flights with $a_{ij} = 0$. Also each column or ToD must correspond to a feasible and legal sequence of flights that satisfies the rules specified in civil aviation regulations or employment contracts or agreements. These rules or constraints can be thought of as being implicitly rather than explicitly satisfied in the ToD-planning model. For example, rules imposing limits on total work time and rest requirements are embedded in the variable generation process.

The ToD-planning model is usually augmented with additional constraints that permit restrictions to be imposed on the number of ToDs included from each crew base. Because these constraints typically have nonunit right-hand-side values, we describe the ToD-planning model as a generalized set-partitioning model.

Air New Zealand's international-pilot and international-flight-attendant ToD-planning problems have additional complexities that cannot be met using the basic ToD-planning model. We have extended the basic ToD-planning model to handle the additional complexities for these different crew types, resulting in superior solutions to those from the basic ToD-planning model. Air New Zealand forms and solves the ToD-planning model independently for each crew type and the flights they can operate.

The Rostering Model

The rostering problem is to construct a LoW for each crew member in a rank so

that each ToD is covered by the correct number of crew members from that rank. For each crew member, we can generate a set of many LoWs from which exactly one must be chosen.

The rostering problem can also be modeled mathematically using a generalized version of the set-partitioning model. Assuming there are p crew members and t ToDs, the model is naturally partitioned into a set of p crew constraints, one for each crew member in the rank, and a set of t ToD constraints corresponding to each ToD that must be covered. The variables of the problem can also be partitioned to correspond to the feasible LoWs for each individual crew member. The A matrix of the rostering set partitioning model is a 0–1 matrix partitioned as

$$A = \begin{bmatrix} C_1 & C_2 & C_3 & \dots & C_p \\ L_1 & L_2 & L_3 & \dots & L_p \end{bmatrix}$$

and $C_i = \mathbf{e}_i \mathbf{e}^T$ is a $(p \times n_i)$ matrix with \mathbf{e}_i the i th unit vector and $\mathbf{e}^T = (1, 1, \dots, 1)$. The n_i LoWs for crew member i form the columns of the $(t \times n_i)$ matrix L_i with elements l_{jk} defined as $l_{jk} = 1$ if the k th LoW for crew member i covers the j th ToD and $l_{jk} = 0$ otherwise. The A matrix has total dimensions of $m \times \sum_{i=1}^p n_i$ where $m = p + t$. The right-hand-side vector \mathbf{b} is given by $b_i = 1, i = 1, \dots, p$ and $b_{p+i} = r_i, i = 1, \dots, t$ where r_i is the number of crew members required to cover the i th ToD. We refer to the first p constraints as the crew constraints, and the next t constraints as the ToD constraints.

The cost vector \mathbf{c} reflects the cost of each LoW relative to all others. Since most airlines do not use optimization systems for rostering, there is no obvious or tradi-

tional measure that can be used to discriminate among feasible solutions in an optimization. Typically the rostering objective reflects either the preferential bidding by seniority (PBS) or the equitable rostering philosophy. We have defined different rostering objectives for each rostering system we have developed through consultation with representatives of the crew groups.

The rostering model has a special structure that deviates from pure set partitioning in that the right-hand-side-vector is not unit-valued and some constraints need not be equalities. The crew constraints of the A matrix also exhibit a generalized upper-bounded structure which is not commonly found in set partitioning.

Each column or LoW for a crew member must correspond to a feasible and legal sequence of ToDs that satisfies the rules specified in civil aviation regulations, employment contracts, and agreements. These rules or constraints can be thought of as being implicitly satisfied in the variable-generation process rather than explicitly satisfied in the rostering model.

We have further enhanced the basic rostering model by adding additional constraints to handle complicated qualification requirements within crew ranks. For example, at least a certain number of highly qualified people may be required on certain ToDs, or no more than two newly trained crew members are allowed on the same ToD. Air New Zealand forms and solves rostering models independently for each crew base and rank.

Solution of Generalized Set-Partitioning Problems

The generalized set-partitioning models

arising in both ToD planning and rostering are solved using linear programming (LP) relaxation and branch and bound. We solved the Air New Zealand ToD planning and rostering models using a special purpose SPP solver, ZIP, developed by Ryan [1980] to take advantage of the characteristics of the SPP model. Over the course of the project, we have enhanced the revised simplex method (RSM) implementation within ZIP to include both partial-pricing and steepest-edge-pricing techniques [Forrest and Goldfarb 1992; Goldfarb and Reid 1977]. The pricing step also includes a dynamic column-generator capability based on resource-constrained shortest-path methods [Minoux 1984; Lavoie, Minoux, and Odier 1988; Desrosiers et al 1991; Gamache et al 1999]. The nature of the underlying shortest-path network used depends on the particular application.

We have also implemented constraint branching [Ryan and Foster 1981] within the branch and bound. In contrast to the conventional variable branching within branch and bound, which forces a variable to value zero or one, constraint branching partitions sets of variables covering a pair of constraints (that is, the set of columns containing a one-element in either of the rows corresponding to the constraint pair) into a set covering both constraints and the complementary set covering just one of the pair of constraints. The pair of constraints will either be covered together by a single variable from the first set or be covered separately with two variables from the second set. In a fractional solution, it is always possible to find a pair of constraints that are covered fractionally by variables from each set. In each crew-

scheduling application, there is a natural choice of constraint pairs to define the branch.

Integer Properties of Set-Partitioning Problems

Some set-partitioning problems are very difficult to solve, whereas others are easier to solve. The LP relaxation of the ToD-planning and rostering models possess interesting structure and properties suggesting that they should be easier rather than harder to solve. It is known that three particular classes of zero-one matrices (unimodular, balanced, and perfect) define set-partitioning polytopes that have no fractional extreme points. Of these three classes, balanced [Berge 1972] and perfect [Padberg 1974] are particularly important for understanding why the ToD-planning and rostering set-partitioning problems are easier rather than harder to solve.

The ToD-planning SPP model generated with severely limited subsequence matrix-generation techniques [Ryan and Falkner 1988] has the properties of a near-balanced problem. Limited subsequence is a heuristic technique that restricts the choice of subsequent activities following any given activity. We have observed that models generated with limited subsequence have fewer variables than the full model but, when solved to optimality, exhibit strong integer properties in that many of the basic feasible solutions near the optimal solution have many variables at integer value and few at fractional values. The objective function of the ToD planning model penalizes the idle time between successive activities in a ToD and thus is consistent with the concept of limited-subsequence generation techniques. The

definition of limited subsequence may be changed dynamically during the optimization convergence.

In the rostering SPP model, the crew constraints are in the form of generalized upper-bound constraints that have the effect of creating perfect blocks of variables (LoWs) for each crew member. Padberg [1974] has shown that a perfect block of variables has no fractional extreme points. This implies that fractional solutions in the rostering model can never occur within the crew member's own variables but

A conservative estimate of the savings is NZ\$15,655,000 per year.

must always occur through two or more crew members competing for a ToD. Our choice of crew versus ToD constraint branching is based on this fractional-solution structure in the rostering model.

Degeneracy and Set-Partitioning Problems

The solution of the LP relaxation of SPPs is often difficult because near-integer basic feasible solutions are very degenerate. Our use of limited-subsequence techniques in the generation of the constraint matrix results in many near-integer basic feasible solutions being visited during the LP convergence, and basic feasible solutions with as many as 80 percent of the basic variables at zero value are common. At such degenerate bases, the RSM tends to stall, sometimes for many thousands of degenerate pivots, even when conventional techniques, such as maximizing the pivot element, are used to determine the leaving variable. This degeneracy phe-

nomenon in the SPP has been discussed by a number of authors (for example, Albers [1980], Falkner [1988], and Marsten and Shepardson [1981]). To avoid the difficulties, Marsten [1974] solved the relaxed LP using a dual algorithm in his SETPAR code. For problems with very large numbers of variables or when dynamic column generation is used, this is not an attractive option. Other authors have avoided the problems of degeneracy by developing alternative bounding strategies based on Lagrangian relaxation or dual-variable adjustments.

We overcame problems of degeneracy in our solution of the ToD planning and rostering models by using steepest-edge pricing [Boyd 1995] and Wolfe's method [Ryan and Osborne 1988; Wolfe 1963] and by using a carefully chosen right-hand-side perturbation scheme. Wolfe's method provides a guaranteed termination of the stall, but we use it only when we detect a sequence of degenerate pivots. The perturbation scheme of the rostering model consists of setting $b_i = 1 + \varepsilon$, $i = 1, \dots, p$, $\varepsilon > 0$, thus creating a tension between the crew constraints of the model and the ToD constraints that still have integer right-hand sides. For a small value of ε , such as 10^{-7} , the values of basic variables are perturbed sufficiently to avoid degeneracy, and in practice, few truly degenerate pivots are observed during the convergence of the RSM. Typically, as few as 100 to 200 degenerate pivots occur in RSM convergences of more than 10,000 iterations, and Wolfe's method is seldom required.

National ToD Planning

In the national ToD-planning system, we construct ToDs for both flight atten-

dants and pilots. All flights are operated by B737 aircraft with identical crew requirements. We build ToDs separately for flight attendants and pilots to cover between 300 and 900 flights per week depending on crew type. ToDs are between one and four days in duration. The crew bases are in Auckland, Wellington, and Christchurch. We produce fully dated ToD solutions because the flight schedule differs from week to week. The major complications of the national ToD planning problem include different pilot and flight-attendant rule sets, multiple home bases with different numbers and types of ToDs required at each home base, crew-base-rank imbalances causing nonunit crewing, a dynamic schedule, and the expected combinatorial complexity of the SPP model. ToDs must satisfy many contract rules and soft rules, which can be complex in terms of definition or implementation. For example, “A maximum duty time of 11 hours is allowed in any rolling 24 hour period” is a rule that is conceptually simple but difficult to implement. An example of a rule that is both difficult to define and difficult to implement concerns the ToD meal-break requirements. In principle, crew members must be provided with meal breaks every six hours. However, the definition of what constitutes a meal break differs for pilots and flight attendants and also depends on where the meal occurs within the duty period. Meal breaks can occur only at certain airports and times of the day. Pilots may also be provided with inflight meals on certain flights, and there is some inconsistency in the exact positioning of meal breaks within a ToD. The contract provides only superficial detail,

and so the rules in the system must reflect complex interpretations based on custom and precedence.

The system generates optimized ToDs for all national crew types, ranks, and bases. The main objective is to minimize the total dollar cost, which includes salary costs based on the number of crew days in the solution and expenses from operating the ToDs. We use a flight-based shortest-path network in the dynamic column generator. The constraint branch used within branch and bound is based on consecutive flight pairs, particularly same-aircraft flight pairs, as this increases operational robustness. A typical problem consists of between 600 and 2,000 constraints, with between 7,000 and 25,000 variables generated by the end of the optimization.

We developed the original ToD planning system in 1984 and 1985 and implemented it as a mainframe computer system in 1986, replacing a manual process. We believe the system was one of the first full optimization-based ToD solvers in production in an airline at that time. Air New Zealand used the system extensively after its introduction in response to the arrival of competition (Ansett New Zealand) in the domestic market. Using the ToD system, the airline could create new schedules and develop its associated crew plans in just two days in response to Ansett initiatives. The system remained in production essentially in its original form until 1997 when we redeveloped it to include improved optimization methods and implemented it on a Unix workstation. Significant benefits and improvements resulting from the optimization system include—higher quality ToD solutions produced

in shorter time,

- manpower savings (reduced ToD crew days and fewer ToD planning staff),
- the ability to analyze flight schedules,
- the ability to repair ToD solutions when the schedule changes or crew availability changes,
- the ability to produce reports for operational and budgeting purposes, and
- its use as a management tool in crew-basing studies, what-if scenarios, and award negotiations.

Features of the system include the ability to force or ban partial ToDs, hot starting from any previous solutions, and the automatic construction of fully dated solutions. The national ToD-planning system is fully integrated with Air New Zealand's schedule-data system and rostering system, and includes a graphical user interface that provides flexible user controls.

Air New Zealand uses the national ToD optimizer every day to evaluate proposed schedules and every two weeks immediately prior to the beginning of the national rostering phase to produce the ToDs covering that 14-day roster period.

National Rostering

The national rostering system constructs rosters for both flight attendants and pilots over a 14-day roster period, resulting in a LoW for each crew member. There are two ranks for flight attendants (pursers and flight attendants) and two ranks for pilots (captains and first officers). The airline builds separate fortnightly rosters for each crew type, rank, and base. The largest national-pilot roster group is Auckland-based captains with 50 pilots, and the largest national-flight-attendant roster group is Auckland-based flight attendants with

85 crew members. These problems cover few crew members compared to the international problems, but their combinatorial complexity makes them the most difficult to solve of all Air New Zealand's rostering problems. Most ToDs last only one or two days and are fundamentally similar in work content. This means that each crew member could be allocated any one of say 25 alternative activities (ToDs, day off, training and so forth) on each of the 14 days in the LoW. While many of the 25¹⁴ possible LoWs will be illegal because they violate some rostering rule, an extremely large number of alternative legal LoWs will remain valid for each crew member.

Air New Zealand uses a "fair and equitable" basis for national rostering. Once it achieves the management requirement of crewing all ToDs, the main objective is to maximize crew satisfaction and LoW robustness. During system development, we invited the flight-attendant-crew union to identify appropriate crew-satisfaction measures, such as the number of days between onerous (or difficult) ToDs, the maximum number of onerous ToDs in a roster period, and the granting of requests for days off, early finishes, and late starts. We incorporate robustness within LoWs by inserting additional rest between adjacent work activities. Some additional complications of the national-rostering problems include the complex rule sets, the existence of part-time and casual flight-attendant staff, and a variety of qualification requirements.

We build the national rosters in two stages, using a days-off optimization followed by a ToD assignment optimization. This approach is possible because, for the

national rostering problems, the daily crew-resource requirements can be accurately calculated prior to roster construction. The initial days-off optimization is necessary to reduce the combinatorial size of the rostering problem.

Initially, the days-off optimization considers only the rules that apply to days off to build a legal days-off LoW for each crew member. Each days-off LoW contains all of the crew member's pre-assignments plus the required total of five days off. We modeled the days-off optimization problem mathematically as a set-partitioning problem. We perform an a priori generation of all possible days-off LoWs for each crew member and then solve the resulting optimization to find a days-off roster ensuring that the correct number of crew members are available to work each day. The objective of this problem takes into account the days-off patterns crew members preferred, for example a days-off LoW containing two blocks of two days off and one single day off is preferred to a days-off LoW containing one block of two days off and three single days off.

We also modeled the ToD assignment optimization as a set-partitioning problem with additional constraints to handle some complex qualification requirements. For example, at least one of the crew members on a ToD must have a premium-service qualification. The solution procedure starts with an "optimal" days-off LoW for each crew member but can use an alternate days-off LoW if this is required to generate a legal LoW. Because of the combinatorial size of the problem, it is not possible to build the full fortnightly roster in one step, so subroster builds are used. A sub-

roster is a practical-sized subsection of the roster (usually five to seven days long) that can be formulated as one tractable set-partitioning problem. For each crew member, we generate a subset of partial LoWs over the subroster period using a limited-subsequence filtering technique. Limited-subsequence filtering restricts the choice of subsequent ToDs that will be considered each day. The crew-satisfaction and operational-robustness measures that contribute to the cost for each partial LoW were chosen in collaboration with crew members. The partial LoWs generated must be consistent with any previous subroster solutions, the days-off LoW, and pre-assignments, such as leave and training. The constraint branch used within branch and bound is based on crew-ToD pairs. A typical five-day subroster problem consists of 300 constraints and 100,000 variables. If the subroster is feasible (that is, all ToDs are covered and all crew are assigned a legal partial LoW), we confirm the subroster solution and consider the next subroster period. If the subroster is infeasible, we generate a new set of partial LoWs based on the dual information from the LP solution. If the subroster still remains infeasible, we unconfirm the last day of the previous subroster period and step back one day, removing all ToD allocations on that single day. We then consider a new subroster period starting on the unconfirmed day. The cause of the infeasibility may be related to ToD assignments on the day just unconfirmed, so the new subroster period may now be feasible. This subroster confirm-or-step-back procedure (Figure 4) continues until the roster is completed or the roster is infeasible.

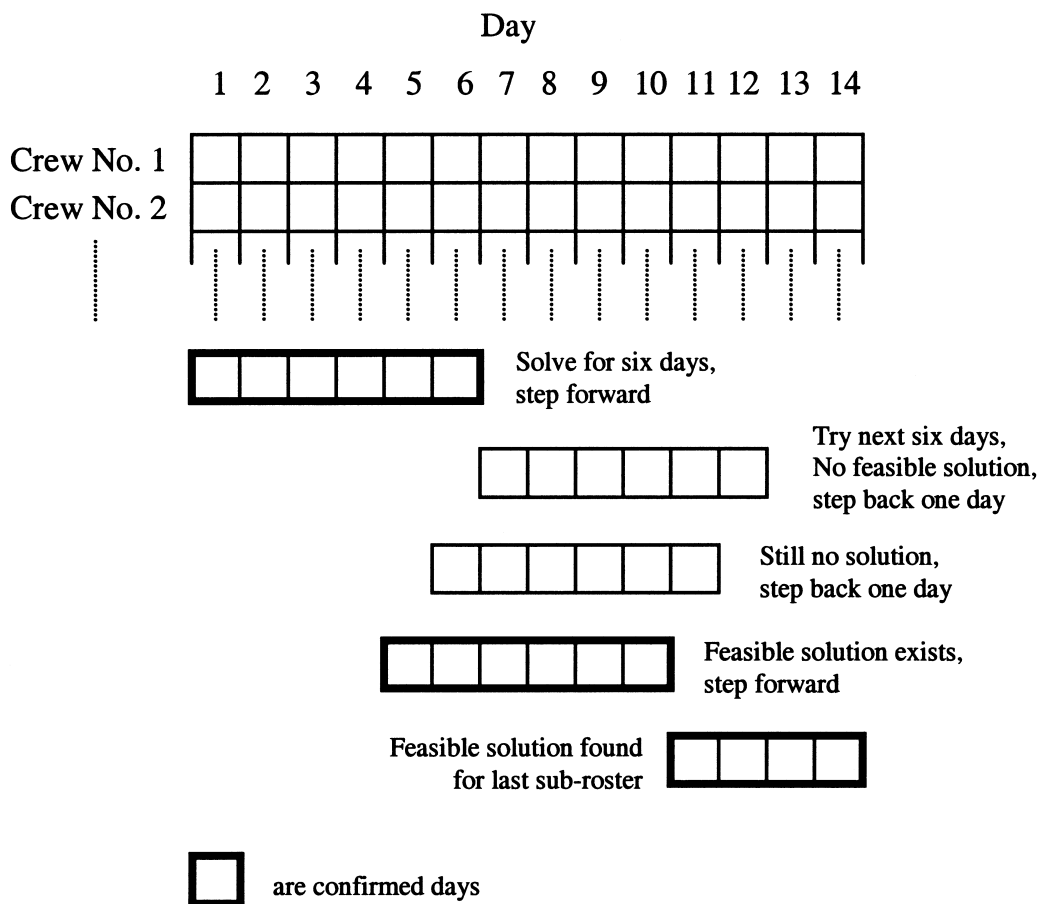


Figure 4: In the subroster solution procedure, the step-back process unconfirms the last day or days of the previous subroster to overcome infeasibility. The remaining subrosters are then built to complete the whole roster. Cells with dark borders represent confirmed days, or subrosters that have been built.

ble. We report detailed information about an infeasible roster to the roster builder so that he or she can make an appropriate decision to remove the infeasibility, for example, to deny a crew request. Since crew satisfaction and LoW robustness are mainly subjective measures, the system aims for high-quality solutions rather than the “optimal” solution. Features of the national-rostering systems include flexible user control of the solution procedure and detailed reports of infeasibilities and roster solutions.

Two previous attempts [Clarke 1989; Mueller 1985] to solve the national rostering problem were unsuccessful, but Day finally solved the flight-attendant crew-rostering problem in 1992 using the sub-roster procedure during PhD research sponsored by Air New Zealand [Day 1996; Day and Ryan 1997]. Air New Zealand implemented a production system based on this research in 1993. In 1998, Scott adapted the same solution methodology to produce pilot rosters under quite different operating rules. The two national rostering

systems are now fully integrated into the Air New Zealand rostering system. The national rostering systems produce rosters of excellent quality for all crew ranks and bases in less than one hour of computation. Two roster builders now construct rosters in one to two days immediately prior to roster publication, 10 days in advance of the first day of the roster period. They manually resolve infeasibilities caused by pre-assignments and then use multiple optimization runs to further improve roster quality. Using the previous manual rostering methods, eight roster builders took two weeks to build the national rosters.

International Flight Attendant ToD Planning

The international-flight-attendant ToD-planning system constructs ToDs for an international schedule containing approximately 300 flights per week. ToDs can last up to 15 days and are mostly for Auckland-based crew, although Air New Zealand has recently established a new base for flight attendants in London. Generally the airline constructs a ToD solution for a standard week because of the stability of the international schedule. The international-flight-attendant ToD-planning problem is particularly difficult in that flight attendants are qualified to operate on different aircraft types. The complexity arises because the different aircraft types require different numbers of crew members. For example, Air New Zealand requires a crew of 14 on the B747-400, eight on the B767-300, and seven on the B767-200. The airline can make considerable savings by having some members of a crew work a different set of flights

from the other crew members. This is known as crew-complement splitting. For example, a full B747-400 crew may split up so that some members of the crew fly a B767 flight in their next duty period. The remaining members of the B747-400 crew could combine with other crew members to make up the crew for some other flight or they could fly as passengers to another airport to join other crew members. Crew splitting is further complicated by the existence of four ranks for international flight attendants and the ability to substitute crew members of higher rank in place of crew members of a lower rank, called rank overcover. This crew-splitting complication does not exist for pilots, who are usually qualified to fly just one aircraft type.

We developed a unique optimization solver for international-flight-attendant ToD planning that automatically incorporates crew splitting within the optimization process [Wallace 2000]. We are unaware of any other optimization-based ToD-planning system that incorporates this feature. The flight constraints for flights that might be covered by multiple ToDs are expanded during the optimization into multiple constraints to model crew splitting. The extra constraints allow alternative complement splits and also ensure that only one alternative is actually chosen for each flight. We implemented this capability using row generation, in which flight-complement splits are evaluated (and appropriate constraints, or rows, added) during the optimization procedure. We have applied a pricing calculation similar to that used in column-generation pricing but involving multiple columns. Ideally the flights that are candidates for

crew splitting should be carefully chosen to minimize the number of splits and to encourage unit crewing. In practice, splitting is permitted everywhere and incurs a penalty cost.

The main objective is to minimize total dollar cost, which includes salary costs based on crew days in the solution and expenses from operating the ToDs. The optimization system uses a duty-period-based shortest-path network in the dynamic column generator. The constraint branch used within branch and bound is based on consecutive flight pairs, usually separated by a rest period. A typical problem consists of between 300 and 800 constraints with between 5,000 and 25,000 variables generated by the end of the optimization. The ToD planners run the system in an iterative manner, refining the solutions to achieve a balance between full- and split-complement patterns and cost savings. We incorporate operational robustness within ToDs by including extra buffer times when crew change aircraft during a ToD. We focus on reducing legal but undesirable (arduous) flight sequences within ToDs and on distributing work equitably between ToDs. This is possible because many different ToD patterns exist at approximately the same cost. The airline uses the international-flight-attendant ToD optimizer every day to evaluate proposed schedules and every four weeks immediately prior to the beginning of the international-flight-attendant-rostering phase to produce the ToDs covering that 28-day roster period.

International-Flight-Attendant Rostering

The international-flight-attendant-rostering problem, involving 1,500

Auckland-based flight attendants in four crew ranks, is the largest rostering problem solved at Air New Zealand. The problem has three processes.

The first international-flight-attendant-rostering process is to create and assign call lines. A *call line* is a sequence of call duties and days off over a 28-day period. A *call duty* is a period of 12 hours during which a crew member can be called out to fill a vacancy in a ToD. Call lines are allocated to crew on a strict rotational basis. Typically, 125 flight attendants will be on call in each roster period. The sequence of call duties and days off within each call line must meet employment contract rules, and the set of all call lines must provide the required number of call duties specified for each day. We have modeled this problem mathematically as a small generalized set-partitioning problem [Deaker 1994], which we solve to generate call lines taking into account varying levels of call requirement for each day. We then solve an assignment problem to allocate the call lines to crew members due for call, maximizing achievement of crew requests.

The second international-flight-attendant-rostering process involves assigning flight attendants who speak various languages (such as German, French, Mandarin, Cantonese, Thai, Japanese, and Samoan) to ToDs requiring those language skills. We developed a languages-assignment optimization to perform this process [Waite 1995]. The language requirements are specified in terms of the number of speakers of each language on selected flights. About 250 crew members have language skills, and each roster period has approximately 700 language re-

quirements. In addition, there are several grades of language skills (ranging from grade 1 for beginners to grade 5 for fluent speakers), and some crew members may speak more than one of the languages. The language requirements can be complex, for example, “One person who speaks both Mandarin and Cantonese at grade 5 level, plus two people who can be either Mandarin or Cantonese speakers so long as at least one of these is a grade 4 or better.” The basic problem is to build partial LoWs for language-skilled crew to cover the language requirements. Language requirements are flight based and independent of crew rank so crews of different ranks and on different ToDs can be combined to cover the language requirement on a given flight. Since Air New Zealand does not have enough language-speaking flight attendants to cover all language requirements, the main objective is to minimize the undercoverage of language requirements. All language requirements are solved as a single optimization problem. Crew members without language skills must also be considered because denying their requested ToD pre-assignments may improve overall coverage of language requirements.

We modeled the problem mathematically as a generalized set-partitioning problem, containing crew constraints, ToD complement constraints, and language-requirement flight constraints. It incorporates a penalty for undercoverage, and it is formulated to spread any undercoverage across lower-priority flights. We included other penalties that take into account crew request satisfaction and ToD variety in a LoW. The language-assignment optimization

system uses a ToD-based shortest-path network in the dynamic column generator. The constraint branch used within branch and bound is based on crew-ToD pairs taking into account language requirements. A typical problem consists of 4,000 constraints with 25,000 variables generated by the end of the optimization. The solver can also be used to roster any other flight-based skill requirements (for example, service-level skills). Compared to previous manual methods for language assignment, the model can solve for more languages and meets more of the language requirements while automatically considering crew requests. The previous manual process took two people several days, whereas one person can complete the process in several hours using the optimization system. This aspect of flight-attendant rostering provides Air New Zealand with important commercial benefits because many of its passengers, particularly those from Asia and Europe, do not speak English. By providing language-qualified crew members on relevant flights, the rostering system supports Air New Zealand’s focus on quality customer service. The system is fully integrated with the rostering optimization. Language ToD assignments are treated as pre-assignments in the main rostering-optimization procedure.

The final international-flight-attendant-rostering process is the allocation of ToDs to form a LoW for each crew member. The ToDs can last one to 15 days, and typically each crew member will be allocated three or four ToDs over the 28-day roster period. The overall problem is therefore less combinatorial than the national-rostering problem. However, many more crew

members are on the international roster. International flight attendants can request a limited number of specific ToDs and days off. The airline considers all of the crew members in each rank to be equal and rosters them on a fair-and-equitable basis.

This means that over time, all crew members should have approximately the same levels of satisfaction of requested ToDs and days off and approximately the same frequency of ToDs visiting desirable and undesirable ports. We use historical statistics for each crew member to determine roster quality. A statistic is a record of the number of times a crew member has been assigned a particular activity and the date of the last assignment. The statistics recorded for crew members include satisfaction of requested days off, ToD destinations, and ToDs that last more than 10 days. The statistical rules are divided into two groups: hard statistical rules that must be enforced and soft statistical rules that can be violated with a penalty. The main objective of the rostering system is to maximize overall crew satisfaction based on statistics. We also apply other soft rules not related to statistics, such as forcing an onerous through-the-night ToD to provide a day off immediately before and after. The airline builds a 28-day roster for each rank independently. There are two crew bases and four crew ranks, with the largest crew rank consisting of approximately 550 crew members. Some additional complications of the international-flight-attendant-rostering problem include complex rule sets and matching LoWs for spouses within the crew rank. The international-flight-attendant-rostering op-

timization system uses a ToD-based shortest-path network in the dynamic column generator. The column generator does not generate LoWs that break hard statistical rules and assigns a cost to each LoW based on the violation of soft statistical rules. The constraint-branch used within branch-and-bound is based on crew-ToD pairs. A typical problem consists of 1,100 constraints with 35,000 variables generated by the end of the optimization. Roster builders intervene to resolve roster infeasibility by altering pre-assignments or relaxing hard statistical rules.

Air New Zealand implemented the original international-flight-attendant-rostering system based on a priori generation as a mainframe computer system in 1989 [Ryan 1992]. We believe that this was the first optimization-based rostering system used by any airline. At the time of its original implementation, the optimized solution demonstrated that it was possible to construct rosters with a five percent reduction in the number of flight attendants (approximately 30 fewer in a rank of 650) and, at the same time, significantly improve the quality of the rosters from a crew point of view. The reduction in crew numbers occurred over time through attrition, and during this time, Air New Zealand increased scheduled flying without hiring additional crew members. We revised the system to incorporate column-generation methods in 1996, which resulted in a further two-percent savings (approximately 20 fewer crew members) and further improved achievement of requested ToDs and days off. The airline achieved these overall savings by increas-

ing crew productivity, reflected in the reduction from an average 13.6 days off per roster in 1988 to an average of 10.4 days off in 2000. The legal minimum-days-off requirement is 10 days off per roster. Because extra days off are allocated to crew on call, we believe an average of 10.4 days off is very close to the achievable minimum. In developing and implementing this system, we consulted representatives of the flight-attendant union, who defined the roster quality measures that we incorporated in the optimization. The international-rostering optimizer is fully integrated into the Air New Zealand rostering system. It produces rosters in two to six hours, depending on crew rank. The airline runs the rostering optimization as close to roster publication as possible, and it publishes rosters 10 days in advance of the first day of the 28-day roster period.

International-Pilot ToD Planning

The international-pilot ToD-planning system constructs ToDs for Auckland-based B767 and B747 pilots. Each week Air New Zealand operates about 50 B747 flights and 250 B767 flights. ToDs can last up to 14 days. Because of the reasonably stable nature of the international flight schedule, the airline can usually build ToDs for a standard week. That is, the airline builds ToDs for one week and then replicates the one-week solution to create ToDs for an entire roster (four weeks). When schedules vary, the airline must build fully dated solutions. The standard crew complement is two pilots (one captain and one first officer). The major complication of the international-pilot ToD-planning problem is the potential use of augmented crew (or third pilot) ToDs,

which allow the airline to extend duty periods by including extra pilot(s) on some flights. For example, the operation of two flights separately may require two pilots per flight, whereas the operation of both flights within the same duty period may require three pilots because of the longer duty time. The airline must then evaluate whether it is better for the solution to contain two short duty periods, each requiring two pilots, or one long duty period requiring three pilots. Also, some pilot rules concern acclimatization in time zones, meaning that the number of pilots required depends on the content of the entire ToD, not just the duty period. Because of the nature of the flight network, the solution to the pilot-ToD-planning problem for Air New Zealand often requires three- and four-pilot crews. This problem is referred to as the crew-augmentation problem.

We have developed a ToD planning solver that automatically constructs augmented-crew ToDs within the optimization. We believe this feature is unique; we understand that the ToD planning systems used by other airlines construct augmented-crew ToDs in a subsequent step. We handled the augmented-crew ToDs by incorporating additional constraints in the optimization model. These constraints ensure that when a ToD requiring additional pilots is in the solution, other ToDs will also be included in the solution that provide extra crew for the required duty periods. The main objective is to minimize total dollar cost, which includes salary costs based on crew days in the solution and expenses from operating the ToDs. The optimization system uses a

combination of a priori and dynamic column generation. The dynamic column-generation subproblems construct legal ToDs using a duty-period-based shortest-path network. The constraint branch used within branch and bound is based on consecutive flight pairs. A typical problem consists of between 140 and 1,100 constraints with between 60,000 and 275,000 variables generated by the end of the optimization. Additional features of this system include the capability to rank-overcover a crew requirement by allowing captains to operate as first officers, the ability to force or ban duty periods and rest periods, the automatic creation of fully dated solutions, and the ability to lock or confirm any partial ToD solutions and reoptimize the remaining flights.

Air New Zealand implemented the international-pilot-ToD-planning system in 1996 [Goldie 1996]. The system provides significant benefits and improvements over the previous manual procedure, including an estimated three-percent cost savings per year. The international-pilot-ToD-planning system is fully integrated with Air New Zealand's schedule data system and rostering system and includes a graphical user interface. The airline uses the international-pilot-ToD optimizer every day to evaluate proposed schedules and every four weeks immediately prior to the beginning of the rostering phase to produce the ToDs covering that 28-day roster period.

International-Pilot Rostering

Many airlines worldwide roster international pilot crews using preferential-bidding-by-seniority systems (PBS) that are based on greedy-sequential-heuristic

roster-construction methods. In PBS systems, crew members bid for work and days off. The airline then constructs rosters by satisfying as many bids as possible but considering crew members strictly in seniority order within each crew rank. The international-pilot-rostering system at Air New Zealand constructs LoWs for Auckland-based B747 and B767 pilot groups. Each crew group has three ranks: captains, first officers, and second officers. The system builds a 28-day roster for each crew rank and fleet. The airline has about 80 B747 captains, 80 B747 first officers, 40 B747 second officers, 110 B767 captains, 115 B767 first officers, and 25 B767 second officers. The ToDs can last up to 14 days, so the overall problem is less combinatorial than the national-rostering problem. However, its major complication is that it must satisfy the maximum number of bids in strict seniority order. That is, a solution in which the most senior pilot achieves all of his or her bids and all other pilots achieve no bids is preferred over one in which the most senior pilot achieves 90 percent of his or her bids and all other pilots achieve all of their bids. Crew-member bids include requests for specific or generic ToDs or days off, bids to avoid certain ToDs and bids for work-pattern characteristics, such as a lot of work, or days off grouped together. Crew members can make any number of bids and can specify the relative importance of each bid. The system satisfies these bids in crew-member-seniority order, with the proviso that all of the ToDs must be assigned and every crew member must have a legal LoW. Some additional complications of the international-pilot-rostering

problem include qualification and training requirements.

Between 1992 and 1994, Thornley [1995] developed a new optimization model and solution method for PBS. The solution method incorporates a unique squeeze procedure that violates the bids of more junior crew members to satisfy the bids of more senior crew members. This guarantees that the airline can satisfy the maximum number of bids in seniority order and assign all ToDs. Other airlines use heuristic methods that cannot provide such a guarantee. Crews negotiated the PBS objective as part of a crew contract. Unlike the national rostering system, the international rostering system cannot assign days off independently of the ToDs. The international-pilot-rostering optimization uses a ToD-based shortest-path network to generate legal LoWs for each crew member. It gives each LoW a cost based on bid satisfaction. The constraint branch used within branch and bound is based on crew-ToD pairs. A typical problem consists of up to 500 constraints with 25,000 variables generated by the end of the optimization. Users must intervene to resolve roster infeasibilities, that is, when no LoW exists for a crew member, or when no feasible roster exists with all ToDs assigned, ignoring all bids. The squeeze procedure guarantees satisfaction of bids in strict-seniority order by solving a series of LP problems that eliminate all LoWs that do not provide the highest bid satisfaction for the current crew member being considered. If this causes the roster to become infeasible, it makes the LoWs that provide the next-highest bid satisfaction for the current crew member available until feasi-

bility is restored. When the current crew member has a fractional LoW coverage at different bid-satisfaction levels in a solution, this implies integer infeasibility and the system applies a back-up procedure. It sequentially moves back to more-senior crew members and allows LoWs with lower bid satisfaction for them into the solution until the LoWs in the solution for the current crew member all have identical bid satisfaction. The squeeze procedure then recommences, starting from the more-senior crew member for whom we added the lower-bid-satisfaction LoWs.

The system implements a difficult PBS contract requirement and has provided high bid satisfaction for senior pilots. Recently we introduced soft rules to the system to improve the quality of days off allocated to all crew members prior to the bid-satisfaction squeeze procedure. We have also made a wider variety of bid choices available to crew members.

Air New Zealand introduced the system as part of a major cultural change it initiated in its employment contracts in 1991. Prior to the change, the airline built international-pilot rosters manually using fair-and-equitable assignment rostering and with salary-based pay. The rostering rules were very restrictive and were based on historic practice. The new contract specified that the airline would build rosters using an automated preferential-bidding-rostering system and would base pay almost exclusively on flight hours (incentive pay). The more hours a pilot worked, the more he or she would be paid. While PBS itself does not provide any direct savings, an acceptable PBS rostering system was a prerequisite for relax-

ing the restrictive rostering rules. Air New Zealand’s successful implementation of PBS between 1994 and 1997 enabled it to build rosters with 30 fewer international pilots. In addition, the more-restrictive historic-practice rostering rules required Air New Zealand to maintain a crew base of about 30 pilots in Los Angeles to operate flights to Europe. The airline no longer needed this crew base once the rules were removed, leading to additional savings. The PBS system is fully integrated into the Air New Zealand rostering system. It takes about two days to produce rosters. The airline runs the rostering optimization as close to roster publication as possible, and it publishes rosters 10 days in advance of the first day of the 28-day roster period.

Implementation

Between 1986 and 1999, the aircrew-scheduling project developed eight optimization-based systems to solve all aspects of the planning and the rostering processes for Air New Zealand’s national and international airlines (Table 1). Initially the airline incorporated these systems into its existing mainframe environments. It has since integrated them into the Sabre AirCrews database environment.

A major factor contributing to the suc-

cess of the crew-scheduling project at Air New Zealand has been the close collaboration between the OR team and the groups affected by the introduction of the systems: the scheduling staff, managers, and crews.

—Our collaboration with the scheduling staff provided us with a comprehensive understanding of all aspects of the ToD-planning and rostering problems. Also the scheduling staff gained confidence in the optimization technology as we developed the systems. This two-way communication helped us to develop customized systems that have a natural fit, both in supporting Air New Zealand’s business practices and delivering high-quality solutions.

—One Air New Zealand manager sponsored the first student project in 1984 and has continued to support development of the crewing optimizers. Since 1984, managers in many areas of the company have recognized the power and benefits of the optimization methods. They have also accepted that considerable research and development are necessary before successful implementation of a production system can be considered. Also management support has been ongoing because each system we have developed and implemented has delivered more than was expected.

| ToD Planning | Pilots | | Flight Attendants | |
|---------------|--------------|------|-------------------|----------------------|
| | Rostering | | ToD Planning | Rostering |
| National | 1986 | 1998 | 1986 | 1993 |
| | Revised 1997 | | Revised 1997 | |
| International | 1996 | 1994 | 1998 | 1989 revised 1996 |

Table 1: Air New Zealand’s aircrew-scheduling project developed eight optimization systems between 1986 and 1999 for tours-of-duty (ToD) planning and for rostering for both national and international pilots and flight attendants.

—To gain acceptance of the systems, we had to provide crew members, who are directly affected by the solutions produced, with some understanding of optimization and reasons to trust the solution technology. We maintained ongoing communication with representatives of the crew groups. “The accessibility of OR staff and their readiness to explain processes and discuss problems proved to be a key factor in the successful implementation of the system, and it is this relationship that will ensure its continued success”, said Captain David Allard, flight instructor on the B767 fleet and scheduling committee member [Allard 2000].

A characteristic of optimization-based methods is that they produce solutions that exploit the rules. Air New Zealand’s introduction of optimized ToD-planning and rostering systems has resulted in increased crew productivity, while still providing very-high-quality solutions. Such systems can lead to solutions that are legal but which contain undesirable features. It is therefore important that solutions not only be optimal from a financial point of view but also crew friendly and safe. We achieved this by collaborating with the crew groups to identify additional soft rules. This is of fundamental importance in maintaining the integrity of the crewing optimizers in a production environment. Management and OR staff continue to work with the crew’s contract-management groups and unions to institute appropriate soft rules to address crews’ quality issues. We can relax the soft rules when we cannot find a feasible solution, so they have no impact on crew productivity yet have a huge positive benefit

on solution quality from the crew’s perspective. Air New Zealand has also collaborated with NASA in pioneering work on crew-fatigue measurement. We have incorporated the results of these studies into the ToD planning and rostering optimizers as additional constraints and rules.

Impacts

The crew-scheduling optimizers provide financial benefits to Air New Zealand by directly reducing the number of hotel-bed nights, meals, and other expenses for crew away overseas and by reducing the total number of crew members required overall. Each optimization application has also reduced the costs of constructing and maintaining the crewing solution for the flight schedule. Over the past 10 years, Air New Zealand’s aircraft fleet and route structure have increased significantly in size yet the number of people needed to solve the crew-scheduling problem dropped from 27 in 1987 to 15 in 2000.

A conservative estimate of the savings from the crew-scheduling optimizers is NZ\$15,655,000 per year.

We have measured savings in relation to the original solutions produced prior to implementation of the optimizations. During implementation, we benchmarked each optimizer against the previous solution methods (in most cases, manual) and identified the cost benefits. We have observed that savings were made even before the systems were fully implemented, because the users incorporated beneficial features of the initial optimized solutions they were evaluating into their manual solutions. Air New Zealand has audited the projects internally to validate the savings claimed for each system.

While the estimated savings are NZ\$15.6 million per year, the estimated development costs over 15 years were approximately NZ\$2 million. In 1999, the savings represents 11 percent of the net operating profit of NZ\$133.2 million for the Air New Zealand group (excluding one-off adjustments).

In addition to the direct dollar savings, the optimization systems provide many intangible benefits:

- (1) Using the crewing optimizers, schedulers produce high quality solutions in minutes; manual solutions took them two or more days. For example, a B767 pilot ToD problem can be optimally solved in approximately 60 minutes, while the B747-400 pilot ToD problem can be solved in less than five minutes on a Unix workstation. Roster builders now solve the international-flight-attendant-rostering problem for 550 crew members in one rank in less than six hours.
- (2) Crew schedulers now function as analysts, preparing and validating data, and interpreting and evaluating solutions.
- (3) The airline no longer depends on a small number of highly skilled schedulers with intimate knowledge of employment contracts and scheduling rules. These contracts and rules are now embedded in the crewing optimizers.
- (4) Managers now use the crewing optimizers to investigate strategic decisions for crewing, for example, evaluating the costs of proposed rule changes or determining ideal crew numbers at crew bases.
- (5) With short build times, schedulers can easily produce solutions close to the day of operation, efficiently accommodating late schedule changes.

(6) The schedule planning group can obtain accurate, reliable feedback on crew costs for proposed schedules, leading to the development of more profitable and more robust flight schedules. This information was formerly difficult to obtain because manual ToD solutions were so time consuming.

(7) Prior to 1986, Air New Zealand operated fixed six-month winter and summer flight schedules. Now it operates flexible flight schedules that can vary from week to week, allowing quick response to market opportunities.

(8) Air New Zealand can now repair solutions quickly to accommodate small changes in the schedule.

(9) The rostering optimizers reflect crew-defined quality measures and soft rules.

(10) The rostering optimizers satisfy over 80 percent of all legal requests from international flight attendants for ToDs and days off and an even greater percentage of requests from national pilots and flight attendants.

(11) The rostering optimizers can accurately identify and minimize roster infeasibility and can also identify days on which there are insufficient numbers of crew.

(12) The international-flight-attendant-languages optimization has improved passenger service for which Air New Zealand is renowned. Air New Zealand has received the "Globe Award for the Best Airline to the Pacific" from the top British Travel Weekly in 2000, 1999, 1997 and 1996. The AB Road Airline Survey also ranked it first for inflight service in August 1999.

(13) The crewing optimizers guarantee and demonstrate that the airline satisfies

legislative and contractual rules. Manual systems could not guarantee compliance without complicated and time-consuming checking.

At the start of this project, Air New Zealand employed no OR analysts. Now Air New Zealand employs six permanent OR professionals in crew scheduling and contracts with up to five OR analysts for various projects and to provide support and maintenance. The OR staff now focuses on exploiting the what-if potential of the OR tools, concentrating initially on crewing-related issues. However, management believes that the OR team, in collaboration with the University of Auckland and OR consultants, will make its most valuable contribution by further developing the underlying tools and techniques and applying them in other areas of the company.

References

- Albers, S. 1980, "Implicit enumeration algorithms for the set-partitioning problem," *OR Spektrum*, Vol. 2, No. 1, pp. 23–32.
- Allard, D. 2000, supporting letter dated March 9, 2000 provided as part of Franz Edelman 2000 finalist submission.
- Anbil, R.; Forrest, J.; and Pulleyblank, W. 1998, "Column generation and the airline crew pairing problem," *Documenta Mathematica*, Extra volume ICM, III, pp. 677–686.
- Anbil, R.; Gelman, E.; Patty, B.; and Tanga, R. 1991, "Recent advances in crew-pairing optimization at American Airlines," *Interfaces* Vol. 21, No. 1 (January–February), pp. 62–74.
- Berge, C. 1972, "Balanced matrices," *Mathematical Programming*, Vol. 2, No. 1, pp. 19–31.
- Boyd, E. A. 1995, "Resolving degeneracy in combinatorial linear programming: Steepest edge, steepest ascent and parametric ascent," *Mathematical Programming*, Vol. 68, No. 2, pp. 155–168.
- Clarke, R. J. 1989, "Crew rostering using assignment algorithms," Masters thesis, Department of Engineering Science, University of Auckland, New Zealand.
- Day, P. R. 1996, "Flight attendant rostering for short-haul airline operations," PhD thesis, Department of Engineering Science, University of Auckland, New Zealand.
- Day, P. R. and Ryan, D. M. 1997, "Flight attendant rostering for short-haul airline operations," *Operations Research*, Vol. 45, No 5, pp. 649–661.
- Deaker, J. 1994, "Construction of standby crew patterns for international flight attendants at Air New Zealand," *Proceedings of the 30th Annual Conference of ORSNZ*, Operational Research Society of New Zealand, pp. 74–79.
- Desrosiers, J.; Dumas, Y.; Desrosiers, M.; Soumis, F.; Sanso, B.; and Trudeau, P. 1991, "A breakthrough in airline crew scheduling," *Cahiers du GERAD*, Report G-91-11, Montreal, Canada.
- Falkner, J. C. 1988, "Bus crew-scheduling and the set partitioning model," PhD thesis, Department of Engineering Science, University of Auckland, New Zealand.
- Forrest, J. and Goldfarb, D. 1992, "Steepest-edge simplex algorithms for linear programming," *Mathematical Programming*, Vol. 57, No. 3, pp. 341–374.
- Gamache, M.; Soumis, F.; Marquis, G.; and Desrosiers, J. 1999, "A column generation approach for large scale aircrew rostering problems," *Operations Research*, Vol. 47, No. 2, pp. 247–263.
- Goldfarb, D. and Reid, J. 1977, "A practicable steepest-edge simplex algorithm," *Mathematical Programming*, Vol. 12, No. 3, pp. 361–371.
- Goldie, A. P. 1996, "Optimal airline crew-scheduling using dynamic column generation," Masters thesis, Department of Engineering Science, University of Auckland, New Zealand.
- Lavoie, S.; Minoux, M.; and Odier, E. 1988, "A new approach for crew pairing problems by column generation with an application to air transportation," *European Journal of Operational Research*, Vol. 35, No. 1, pp. 45–58.
- Marsten, R. E. 1974, "An algorithm for large set partitioning problems," *Management Science*, Vol. 20, No. 5, pp. 774–787.
- Marsten, R. E. and Shepardson, F. 1981, "Exact solution of crew-scheduling problems using the set partitioning model: Recent successful applications," *Networks*, Vol. 11, pp. 165–177.
- Minoux, M. 1984, "Column generation tech-

- niques in combinatorial optimization—A new application to crew pairing," *Proceedings of XXIV AGIFORS Symposium*, Strasbourg, France, pp. 15–30.
- Mueller, M. 1985, "Air crew rostering for Air New Zealand," Third professional project report, Department of Engineering Science, University of Auckland, New Zealand.
- Padberg, M. 1974, "Perfect zero-one matrices," *Mathematical Programming*, Vol. 6, No. 2, pp. 180–196.
- Rubin, J. 1973, "A technique for the solution of massive set covering problems, with application to airline crew-scheduling," *Transportation Science*, Vol. 7, No. 1, pp. 34–48.
- Ryan, D. M. 1980, "ZIP—A zero-one integer programming package for scheduling," Computer Science and Systems Division Report 85, Atomic Energy Research Establishment, Harwell, Oxfordshire.
- Ryan, D. M. 1992, "The solution of massive generalised set partitioning problems in air-crew rostering," *Journal of the Operational Research Society*, Vol. 43, No. 5, pp. 459–467.
- Ryan, D. M. and Falkner, J. C. 1988, "On the integer properties of scheduling set partitioning models," *European Journal of Operational Research*, Vol. 35, No. 3, pp. 442–456.
- Ryan, D. M. and Foster, B. A. 1981, "An integer programming approach to scheduling," in *Computer Scheduling of Public Transport*, ed. A. Wren, North-Holland, Amsterdam, pp. 269–280.
- Ryan, D. M. and Osborne, M. R. 1988, "On the solution of highly degenerate linear programmes," *Mathematical Programming*, Vol. 41, No. 3, pp. 385–392.
- Thornley, M. O. 1995, "Crew rostering under a seniority preferential bidding environment using column generation," Masters thesis, Department of Engineering Science, University of Auckland, New Zealand.
- Waite, J. R. F. 1995, "Language assignment for flight attendant rostering," Project report, Department of Engineering Science, University of Auckland, New Zealand.
- Wallace, C. A. 2000, "Construction of optimal tours-of-duty for long-haul flight attendants," PhD thesis (in preparation), Department of Engineering Science, University of Auckland, New Zealand.
- Wedelin, D. 1995, "An algorithm for large scale 0–1 integer programming with application to airline crew-scheduling," *Annals of Operations Research*, Vol. 57, pp. 283–301.
- Wolfe, P. 1963, "A technique for resolving degeneracy in linear programming," *Journal of the Society for Industrial and Applied Mathematics*, Vol. 11, No. 2, pp. 205–211.