




INFORMS Journal on Applied Analytics

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Prescriptive Analytics for Swapping Aircraft Assignments at All Nippon Airways

Saravanan Venkatachalam, Suresh Acharya, Kenji Oba, Yoshinari Nakayama

To cite this article:

Saravanan Venkatachalam, Suresh Acharya, Kenji Oba, Yoshinari Nakayama (2020) Prescriptive Analytics for Swapping Aircraft Assignments at All Nippon Airways. INFORMS Journal on Applied Analytics 50(2):99-111. <https://doi.org/10.1287/inie.2019.1016>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2020, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Prescriptive Analytics for Swapping Aircraft Assignments at All Nippon Airways

Saravanan Venkatachalam,^a Suresh Acharya,^b Kenji Oba,^c Yoshinari Nakayama^d

^a Industrial and Systems Engineering Department, Wayne State University, Detroit, Michigan 48202; ^b Decision, Operations and Information Technologies, Robert H. Smith School of Business, University of Maryland, College Park, Maryland 20850; ^c Juro International Systems, Kansas City, Missouri 64112; ^d All Nippon Airways, Tokyo 105-7140, Japan

Contact: saravanan.v@wayne.edu,  <http://orcid.org/0000-0002-7319-1456> (SV); sacharya@rhsmith.umd.edu (SA); jurois@aol.com (KO); yos.nakayama@ana.co.jp (YN)

Received: July 11, 2018

Revised: April 29, 2019; June 25, 2019

Accepted: July 9, 2019

Published Online in Articles in Advance:
March 9, 2020

<https://doi.org/10.1287/inte.2019.1016>

Copyright: © 2020 INFORMS

Abstract. In airlines operations, fleet assignment models are used to assign flight legs to different equipment types (such as Boeing 760 or 767 or Airbus 320). The fleet assignments have to satisfy certain operational constraints, such as coverage, maximum overnight stays, and airport compatibility. Fleet assignments are tactical decisions, and changes in demand and maintenance requirements require an intermediate decision-making process to capture these changes before a flight's day of departure. In this paper, we describe an implementation of a swapper optimization suite (SOS) for one of the largest airlines in Japan. The scale of the flight legs, the equipment types, complex operational constraints, maintenance requirements, and other complex criteria specified by the route planners necessitated the development of a sophisticated optimization suite to generate swaps of flight legs among the different equipment types for the allotted fleet assignments. The SOS uses optimization models to generate the optimal swaps. The company has seamlessly integrated the SOS in its information systems to incorporate optimization into the decision-making process.

Keywords: prescriptive analytics • decision support • fleet assignment model • equipment swapping application • airlines tactical decision • large-scale optimization

The International Air Transportation Association (IATA) year-end report for 2017 estimated that consumers would spend 1% of the world's gross domestic product (GDP) on air transport in 2018 (IATA 2018). Over 4 billion passenger departures occurred in 2017, and the number of jobs in the industry exceeded 2.7 million. Air transport boosts economic development worldwide by increasing connections between cities, enabling the flow of goods, people, capital, technology, and ideas. The total number of unique flights between two cities in 2017 exceeded 20,000, which is more than double the connectivity by air transport 20 years earlier. In addition to connectivity, an important contributor to the increase in air transport volume is a decrease in the cost of air travel. According to the IATA, the total revenue for 2017 was \$750 billion (and it was estimated to be \$834 billion for 2018). Also, although the year-to-year passenger population increased by 7.1% in 2017, the earnings before interest and taxes (EBIT) margin (ratio of earnings before interest and taxes to net revenue earned) decreased by 0.9%. Thus, it has become imperative for airline companies to make their operations lean for fiscal sustainability and growth.

The optimal alignment of demand and supply is paramount to the success of any business. In the airline industry, this alignment is achieved using decision

support systems in the areas of demand management and capacity planning. Demand management encompasses the core areas of forecasting, revenue management, and pricing. Determination of the optimal policy for accepting or rejecting bookings and being able to shape demand using price levers are central to an airline's ability to improve revenue. Capacity planning, however, addresses the supply side of the equation and includes, for example, long-term fleet planning; fleet rationalization; and the more tactical areas of fleet assignment, crew scheduling, gate assignments, and fleet swapping. A common practice in airlines' scheduling processes is to determine the equipment assignment for a flight schedule based on a forecasted, unconstrained demand. This exercise is typically done about a year in advance for budget planning purposes, and it is adjusted two to three months before finalizing the schedule. However, factors, such as the near-term trend, local events, and competitive pricing actions, can have a profound impact on short-term forecasts, possibly leading to significant mismatches between demand and capacity. In this paper, we introduce the swapper optimization suite (SOS), which was designed to optimize the near-term challenge of potential mismatches between the types of aircrafts assigned to flights and the projected demand based on on-hand bookings and the

booking pace. The SOS is an optimization solution for maximizing profitability while satisfying all operational requirements.

All Nippon Airways (ANA), the largest airline in Japan, has 85 international routes and 119 domestic routes in Japan. It has a unique dual-hub model that enables passengers to travel to Tokyo and connect through the two airports in metropolitan Tokyo, Narita and Haneda, to various destinations throughout Japan. The airline also offers same-day connections between various North American, Asian, and European cities. The ANA Group carried 53.8 million passengers in fiscal year (FY) 2017 and has a fleet of approximately 260 aircraft. ANA's operating revenues for FY 2017 were ¥1,971.7 billion. The airline had 122.8 billion available seat km (ASKs) and 89.4 billion revenue passenger km (RPKs). ASKs are the sum, across all segments of the airline's routes, of the number of available seats on the segment multiplied by the distance of the segment expressed in kilometers. RPKs are the sum, across all segments of the airline's routes, of the number of passengers on the segment multiplied by the distance of the segment also expressed in kilometers (Figure 1).


Before implementation of the SOS, ANA already used a fleet assignment model (FAM), which is an optimization tool that assigns aircraft to flight legs in a way that maximizes the overall profitability of the schedule. Whether determining the feasibility of

a future schedule or gauging the profitability of a current schedule, the FAM provides valuable insights into the flight scheduling process. The major specifications for the FAM include

- producing fleet assignments that are consistent with a base schedule and meeting maintenance requirements;
- adhering to specified ground turn times and block times (block times are the total amount of time that a flight takes from pushing back from the departure gate to arriving at the destination gate);
- maintaining "prefleeted" flights—a normal business practice for assigning the equipment to certain flights;
- respecting "through" and "forced turn" flights by maintaining the same equipment type for some flights that have one or two intermediate stops before arriving at their final destination and maintaining the same equipment type throughout the route for some flights; and
- enforcing crew-related restrictions.

In addition to these requirements, the FAM is also restricted by carrier-, airport-, and routing-specific constraints. The FAM is used as a tactical tool in which the initial equipment assignments are driven by representative forecast values. A "weekday" fleet assignment is made based on average daily forecasts and fares, and the equipment assignment is identical for all weekdays of the schedule. Similarly, a "weekend" fleet assignment captures different demand patterns for weekends.

Figure 1. (Color online) ANA's Domestic Hub-and-Spoke Network, Which Includes ANA, IBEX Airlines, AIR DO, Solaseed Air, and STARFLYER, Serves 132 Routes and 1,054 Flights per Day



The initial assignments made by the FAM are used for budgeting and long- and short-term planning. For long-term planning, the rosters for the aircraft provide overviews for strategic objectives and overall profitability. In the shorter term, the rosters are useful for publishing a schedule and for operations planning. However, there are challenges associated with using average demand forecasts for FAM assignments, because these forecasts fail to reflect day-to-day variations, special events, competitive actions, and other marketing decisions. Also, as the day of departure gets closer, the demand is better represented by actual customer bookings-based projections. Another important phenomenon that is not captured by the FAM is “group reservations” or “bulk bookings,” which generally tend to happen closer to the departure date. Group reservations are higher-volume bookings for special events, such as conferences, sporting events, corporate training, and large family functions. These reservations can substantially increase the bookings for certain legs, and these reservations typically occur after assignments are made by the FAM; therefore, there is a need to revise the FAM’s assignments to accommodate the additional revenues and passenger volume without violating other operational constraints, such as maintenance and pilot-crew requirements. Additionally, there can be changes in the availability of overnight maintenance routines, crew, or the schedule as well as other minor changes after the FAM assignments, and these developments can make some of the ground times unacceptable and thus, require a change in the equipment type for a flight. Because of these challenges, ANA was motivated to add an optimization process on top of the FAM to make better, more robust assignments for the equipment. The concept of matching more recent demand with capacity may sound like it should be simple; however, these changes need to be conservative for the airline’s operations and can give rise to additional challenges in terms of pilot and crew assignments, ground handling, passenger acceptance, and spares and maintenance for the aircraft. Hence, any deviations from the rosters generated


by the FAM create an additional burden for the operations team (OT), and therefore, the changes must be minimal. Given the large number of aircraft and flights and complex requirements for the fleet assignments, an efficient tool, such as the SOS, is required for producing meaningful recommendations.

Literature Review

Strategic and tactical planning for airline operations is broadly classified into four categories: network planning, fleet assignment, aircraft routing, and crew rostering (as shown in Figure 2). The SOS is a tactical tool, and it is used multiple times before a flight’s day of departure. The functionalities of the tool are restricted based on the number of days until departure.

Network planning is a very strategic phase in which flights are determined for a market based on demand and potential profits, and fleet assignment is another strategic step in which individual aircrafts are assigned to each leg based on cost and the supply of different seat segments. An FAM is then used for aircraft routing based on maintenance and operations restrictions, and finally, crew rostering is performed based on the crew costs and workload regulations. The SOS is used intermittently after the FAM has provided the initial routing up to the day of departure. The literature on FAMs is extensive, but the literature on swapper applications is very limited. The main objective of an FAM is to create a line of flight (LOF) for each aircraft in a way that maximizes profit. Aircrafts differ in their configurations, capacities, operating costs, and maintenance requirements, and the demand for each flight will also differ based on the day of the week and the time of day. Hence, the supply of aircraft has to be well matched with travelers’ demand, like any other business requirement. Allocating larger aircraft for a smaller demand will result in a low profit margin because of the higher operating costs and lower occupancy, and similarly, allocating smaller aircraft for high-demand flights will result in demand spill (unsatisfied demand), poor customer service, and multiple reroutings for customers. The literature has

Figure 2. (Color online) Operations Planning in the Airlines Industry



largely treated FAMs as multicommodity network flow problems with cover, flow balance, and aircraft availability constraints. We refer readers to Abara (1989), Daskin and Panayotopoulos (1989), Berge and Hopperstad (1993), and Subramanian et al. (1994) for earlier work on FAMs and to Rushmeier and Kontogiorgis (1997), Gopalan and Talluri (1998), and Sherali et al. (2006) for surveys of FAM models and methodologies. Dillon et al. (1993), Jarrah (1993), and Subramanian et al. (1994) present the FAMs used by USAir, American, and Delta, respectively.

Regarding the complexity of FAMs, Gu et al. (1994) have shown FAMs to be NP hard for three aircraft types, even without availability constraints. Hane et al. (1995) and Rushmeier and Kontogiorgis (1997) have presented different modeling strategies for large-scale FAMs. Other methodological advances include solving the linear programming relaxation of the mixed integer model, applying rounding heuristics, using branch-and-bound search techniques for the remaining variables (Rushmeier and Kontogiorgis 1997), Lagrangean relaxation (Rushmeier and Kontogiorgis 1997), branch-and-price solution schemes (Hane et al. 1995, Bélanger et al. 2006), and large-scale neighborhood searches (Ahuja et al. 2007).

Talluri (1996) presented heuristics for a swapping application with a focus on the required overnight maintenance of aircraft. Additional applications include balancing because of schedule disruptions and undergoing designated maintenance. The swapping also satisfies the most important requirements of a fleet assignment, such as flow balance, aircraft count, and coverage. Swapping aircraft is also a useful tactic in disruption management during operations.

As Kohl et al. (2007) note, a number of different techniques are available to airline operations control centers to mitigate the effects of disruptions. These generally take the form of flight delays, cancellations, and “aircraft swapping.” Ageeva (2000) presented the idea of using aircraft swaps to avoid delays, and Eggenberg (2009) showed that aircraft swaps and increased idle and passenger connection times are useful for improving the recoverability of aircraft routings. More recently, Froyland et al. (2013) introduced a two-stage stochastic program in which FAM decisions are made in the first stage and swapping is considered as a recovery option in the second stage. Burke et al. (2010) considered a multiobjective optimization problem where reliability and flexibility were represented by the probability of being on time and swapping opportunities, respectively.

This paper focuses on what happens after the FAM-based assignment of flights to equipment. The swapping application that this paper describes is useful for meeting variations in demand and increasing revenue as well as accommodating changes in maintenance routines that have occurred after the FAM assignments.

Background

A “flight” is a scheduled service between a departure airport and an arrival airport with a scheduled departure time and a scheduled arrival time. There are different airplane fleet types (for example, Boeing 787 and Airbus A330), and within each fleet type, there are also configuration types referred to as equipment configurations (ECs), which specify the number of seats in the economy and business classes. Each aircraft corresponding to an equipment and configuration type will have multiple flights assigned to it; this assignment is the aircraft’s LOF, which is a sequence of connected flight legs that begins and ends at (possibly different) maintenance airports. A “maintenance airport” is one where the maintenance for the aircraft is performed during an overnight stay; not all airports are maintenance airports. A maintenance airport is also capacitated by aircraft type and configuration. An FAM process produces the LOF for an aircraft. The assignments of crew, minimum and maximum ground times, and maintenance routines will depend on the aircraft type, and the demand served will depend on the aircraft and configuration type. A “minimum” or “maximum” turn time is the minimum or maximum time required between a flight’s arrival and the same aircraft’s subsequent departure, respectively; these turn times will ensure that there is sufficient time to conduct critical activities, such as disembarking, removal of bags, refueling, loading of bags, and boarding for the next flight. The departure time for a flight can be changed slightly if it can facilitate additional swapping opportunities for the aircraft to increase revenue or decrease demand spill.

Swapping must adhere to a number of inclusion and exclusion rules, some of which we list below.

- Violates forced turn: certain flights are required to connect to other flights
 - Runway length: certain equipment types are prohibited at certain airports because of runway length requirements
 - Segment prohibition: some flights have explicit equipment prohibitions
 - Preferred equipment and configurations: planners can specify a preference for certain flights with certain equipment types and configurations
 - First and last flights: certain flights are designated as the first flight of the day, and some others are designated as the last flight of the day
- Any viable swap should adhere to all of the rules and prohibitions that are established by the OT.

ANA’s Previous Practice and the Need for Analytics

Before the implementation of the SOS, the route planners (RPs) at ANA selected the flights that warranted

higher-capacity assignments based on observed booking patterns and sent a request to the OT. The OT evaluated the list of flights based on firm on-hand bookings, crew conditions, and other operational requirements and provided the RPs with swapping recommendations. The RPs then had to decide whether to adopt each OT recommendation. The major limitation of this approach was missed opportunities. Because the swaps were identified manually, the process was cumbersome, requiring multiple interactions across multiple departments, and there was no consideration of cost savings from downgauging flights (that is, moving the flights from using larger to smaller equipment, because the revenue management planners focused primarily on revenue maximization). However, the OT wanted minimal changes to the operations to reduce its burden. The SOS was introduced to improve the swapping process at ANA. The major challenges for using this process before adoption of the SOS are listed in Table 1. Other key features, such as multiple iterations and downgauging, would be ideal for the swapping process; however, they were not part of this process because of their complexities in the absence of an optimization process.

To illustrate how important swapping can be, consider two aircraft, 318 and 319, with capacities of 124 and 107, respectively, and FAM-assigned LOFs as we show in the top panel of Figure 3. The current projected bookings (denoted as “pb”) for the flight from Haneda Airport (HND) to New Chitose Airport (CTS) has increased beyond the capacity of aircraft 318. Hence, the two flights, HND to CTS and CTS to HND, are exchanged with those of aircraft 319, which has enough additional capacity to accommodate the increase in the current bookings. Fortunately, aircraft 318 has the required capacity for the flights swapped from aircraft 319 between HND and Osaka International Airport (ITM) and between ITM and HND, and the swap is in accordance with all inclusion and exclusion rules. Additionally, we note that this swap does not change the aircrafts’ overnight maintenance requirements, and no other operational constraints, such as


minimum/maximum ground time and maximum aircraft or configuration changes, are violated.

Although it is trivial to spot the profitable exchanges for aircraft 318 and 319 in Figure 3, on an average day there are approximately 600 flights with 100 LOFs and seven different aircraft configurations. Also, the example in Figure 3 demonstrates only a two-level swap, whereas in practice, a multilevel swap can also happen. For example, a multilevel swap involving three aircraft A1, A2, and A3 might swap flights between A1 and A2 and between A2 and A3, with A2 and A3 making some swaps to accommodate the swaps between A1 and A2. In addition, the OT is interested in making very few changes to the committed schedule, whereas the RPs are interested in revenue maximization. Although the departure times are known and fixed, the SOS supports flexible departure times and allows planners to conduct what if analyses. All changes will constitute a network effect and produce other maintenance requirements for the flights toward the end of the day. Because of the complexities involved in swaps and the possible number of combinations for swaps, an optimization system, like SOS, is required to identify the LOF exchanges that will provide the maximum revenue benefit.

The swapping process at ANA has different types of restrictions based on the number of days until the departure date. As the departure date approaches, the number of restrictions will increase, because the maintenance rosters and flight arrival and departure times have already been published. In contrast, when the departure date is farther away, the swapping process is expected to recommend exchanges that will produce the highest monetary benefits. Table 2 presents the restrictions for the swapping process, which are based on the number of days until the departure date. For example, only round trip swaps are allowed when the swapping process is very close to the departure date, because this approach would not disrupt other published LOFs. Similarly, departure time changes are not allowed within three days of the departure date.

Table 1. Swapping Practice at ANA Before the SOS Implementation and the Drawbacks of the Process

Before SOS	Drawbacks
Manual process	Laborious process of identifying swaps by looking at bookings and the routing files
Labor intensive	1.5 FTEs required to identify, review, and confirm equipment swaps
Downgauging not considered	No cost savings by downgauging fleet assignments
Low-quality solution	OT acceptance was moderated, and fewer than six equipment changes per day were identified; error prone
Few iterations	Short-sighted planning process
Low importance given to swapping	Critical tasks were given priority over fleet-swapping efforts

Figure 3. (Color online) Movement of Flights from One Aircraft to Another Aircraft Owing to an Increase in Current Bookings

Note. KIX, Kansai International Airport; ISG, New Ishigaki Airport; NGO, Chūbu Centrair International Airport; FUK, Fukuoka Airport.

The SOS recommends an “iterative-optimal” solution: that is, in each iteration of the algorithm, the system optimally discovers the most promising LOFs for which flights can be swapped to generate maximal gains in the incremental profit. On discovering a swap, the system commits to the swap, records it, and continues to the next iteration to find the next most promising swap. Again, the system does so optimally for each iteration. The algorithm terminates when it cannot find any meaningful swaps, or it reaches the aircraft change limit established by the user. In addition, each swap that the SOS finds is fully feasible, because the SOS considers all operational constraints while identifying swaps. The number of LOFs that are involved in each swap is parametrized and can be changed to three or more, which means that the swapper system can discover a multiple-swap LOF solution. Although increasing the number of swaps provides more flexibility for the optimizer and can generate theoretically better swaps, executing such a solution can be quite challenging for an airlines OT.

SOS is designed as an iterative tool, because ANA must internally review and approve each swap. The review considers strategic and operational plans,

which are outside the scope of the SOS. The goal of the SOS is to provide a list of easy to understand and profitable swaps that the users can review and accept. Solving a single large mixed integer linear program is certainly possible; however, the results in multiple-line, multiple-flight swaps are hard for the planners to understand. In contrast, SOS ranks each swap based on the objective function value and provides a perspective to the user on the type of swap and its corresponding impact on overall revenue. Some swaps could be difficult from a practical perspective and may result in lower revenue and vice versa. Hence, the iterative process is helpful in distinguishing the swaps based on revenue.


Developing the SOS: Optimization Models

The SOS is an optimization tool that swaps flights for aircraft based on the assignments made by the FAM. The swaps are performed in a way that maximizes the overall profitability of the schedule. With advances in algorithmic techniques and computational power, the SOS has matured into a dependable tool in the schedule development process. Whether used to determine the feasibility of a future schedule or gauge the profitability of a current schedule, the SOS provides valuable insights into the flight scheduling process. As the airline industry becomes more and more competitive, tools, such as the SOS, are indispensable for a carrier’s survival and growth. The overall architecture of the SOS is presented in Figure 4. The optimization engine has three input streams. One provides the recent “unconstrained” demand bookings and current fares, another provides the existing assignments recommended by the FAM,

Table 2. Order of the Restrictions from Most Restrictive to Least Restrictive as Defined Within the SOS

Feature	Scenario
Run period	Daily/weekly/monthly
Swap	Round trip only/no restrictions
Overnight stay	Fixed aircraft and counts/no restrictions
Maximum aircraft changes	100/no restrictions

Figure 4. (Color online) The Overall SOS Architecture Shows that Three Streams of Data and a User Interface Guide the Optimization Engine



and the third provides operational constraints. “Unconstrained revenue” is the expected revenue from a flight if capacity constraints are ignored. In addition to these input streams, other user parameters guide the optimization process. The solution process comprises a preprocessing stage and solving a mixed integer linear programming (MILP) model. The preprocessing ensures that any swap will adhere to the inclusion and exclusion rules set by the OT, and it also helps reduce the size of the mathematical model, resulting in a better run time. We present the overall architecture and details of the optimization model in the appendix.

Preprocessing: Inclusions/Exclusions

Before presenting the viable options for swapping flights between two aircraft, preprocessing is performed to provide incentives or exclude potential swaps from the model. The general rules include a capacity and a compatibility check; we list some considerations below:

- an aircraft’s capacity cannot be less than the current bookings for a “swappable” flight;
- ground times for an aircraft are within the parameter-specified tolerance;
- a swap will not violate any forced turns;
- spare aircrafts are grounded within the parameter-specified tolerance;
- equipment is assigned to airports with compatible runway lengths and gate capabilities;
- equipment and configuration prohibitions for certain flight segments are not violated;
- parameter-specified special block times for certain combinations of airports and equipment are met; and
- parameter-specified allowable time changes for the legs and equipment types are met.

Swapper Model

As we mention above, the SOS is a tool that makes swapping decisions based on the overall profitability of a schedule. The profitability of a flight is the difference between the revenue generated and the cost incurred for that flight. Intuitively, revenue can be computed by multiplying the unconstrained expected demand in each fare class by the associated fare. Although this may seem to be a straightforward task, a sophisticated model or the assistance of an experienced forecaster is necessary to estimate the unconstrained expected demand.

The number of decision variables in the MILP model depends on the flights and aircraft configurations. The flight sequences are constructed as a network, with connections representing “legitimate” combinations of arrival and departure flights at an airport and arcs representing flights. Any pair of connections in the network is defined as decision variables (Figure 5). Operational feasibility requires ensuring that the combination of flights obeys all operational constraints, such as minimum, maximum, and special ground times between flights; forced turns; flight-airport compatibility; and so on. The preprocessing step is performed to obtain all legitimate combinations, and then, these combinations become the decision variables of the MILP model. The flag in each variable indicates whether the FAM’s assignment of the flight to a particular aircraft configuration was retained.

The objective of the MILP model is to maximize profit, which is revenue minus costs. The revenue and costs are calculated for each connection variable. The costs are calculated based on fleeting decisions, which are often referred to as “flight variable costs.” For example, fuel costs must be included, because the total cost of a flight will depend on the fuel efficiency of the aircraft. However, overhead and ownership costs can

be excluded, because they tend to be fixed in the short term and have less of a bearing on the equipment type.

The types of constraints can be broadly categorized into three major groups: (1) balance and coverage constraints, which ensure that each flight is assigned to an aircraft, the total number of LOF assignments is not greater than the total number of available aircraft, and mass balance constraints for the assignment of flights are satisfied; (2) operations constraints, which address all of the operational restrictions, including constraints pertaining to overnight maintenance requirements for each airport or aircraft, crew assignments to aircraft, and limitations for the slot counts at the airports; and (3) LOF constraints, which limit the number of changes in LOFs or the number of assignments for each aircraft.

Implementation

The graphical user interface (GUI) interface for the SOS was developed using the .NET framework. Based on a user's preferences, all of the input data required for the optimization engine are prepared, and the optimization model is then triggered. The optimization model was developed using the Java programming language and is solved using the commercial solver CPLEX 12.6.3. For our hardware, we used an Intel Core i7-4790K 4.0-GHz CPU with 32 GB RAM. Constraints and functionalities for the optimization model are invoked based on a run type as specified in Table 2. The optimization model results are generated as output and then, uploaded back into the SOS IT system.

Table 3 provides data characteristics for some of the daily runs at ANA without any of the restrictions indicated in Table 2. ANA operations involve 61 airports. The preprocessing helps by removing all of

the exclusions for possible connections. For example, the total number of connections for ECs and flights could potentially be $17,101$ ($2,443$ flights $\times$ 7 ECs); however, given the exclusions, that number is only $3,517$. Similarly, the number of connections, for example 1, could potentially be $3,517 \times 3,517$; however, the number is only $171,533$ with exclusions. Some other elements of the model are block groups, special ground times, and special block times. Block groups refer to the grouping of aircraft types for the purpose of calculating the total block times. Special ground times are adjustments to the standard ground times to reflect time of day and day of week congestion patterns at certain airports. These times can also be influenced by the markets served, because leisure markets typically have longer turnaround time because of passenger and baggage handling needs. Special block times refer to adjustments in the block times that reflect air traffic congestion and taxi times owing to airport maintenance and renovation. Also, "res" in Table 3 is the number of restrictions for equipment configurations and airports. Some airports have restrictions or preferences for certain types of aircraft owing to runway lengths or gate capabilities. The specified restrictions and preferences are all added to the model as either a penalty or an incentive for the equipment configuration and airport combination.

Table 4 provides an indication of the size of the mathematical models run on SOS. The characteristics are provided for the data instances reported in Table 3. Columns (2)–(5) in Table 4 denote the size of the mathematical model for the first iteration, and the last column "Time" denotes the total run time for all iterations. Based on the number of nonzeros as shown in Table 4, the model is very sparse. To provide clarity for the OT regarding implementation of the SOS

Table 3. Data Characteristics for Some SOS Runs at ANA

Instance	No. flights	No. eqp	No. ECs	No. combs	No. conns	No. BGs	No. SGs	No. SBs	No. res
1	2,443	7	7	3,517	171,533	2	249	259	274
2	2,545	7	8	4,105	127,188	2	249	259	274
3	2,595	9	12	5,817	162,045	2	353	256	295
4	2,443	7	7	3,548	174,177	2	249	259	274
5	2,530	8	11	5,151	138,489	3	306	350	319
6	2,606	9	12	5,836	159,700	2	353	256	295
7	2,545	7	8	4,108	127,324	2	249	259	274
8	2,590	9	12	5,829	160,670	2	353	256	295
9	2,626	9	12	5,923	168,628	3	343	350	360
10	2,545	7	8	4,193	130,258	2	249	259	274

Notes. There are 61 airports, and the model size increases as the number of connections increases. No. flights is the number of flights for a particular day, and No. eqp and No. ECs represent the numbers of aircraft (for example, Boeing 765 and Airbus 320) and the different configurations (for example, Boeing 765-300 and Airbus 320-500) available for the day, respectively. No. combs and No. conns represent the numbers of combinations of flight and equipment configurations and connection variables as represented in Figure 5, respectively. No. BGs represents block groups, No. SGs represents additional ground times, No. SBs specifies the special block times, and No. res is the number of restrictions for equipment configurations and airports.

recommendations, the model recommends an LOF change in each iteration. The recommended LOF swaps will provide the highest revenue. During each iteration, the swaps from previous iterations are fixed, and a new pair of swaps is obtained. In this way, the OT can quantify the impact of the swaps, which provides the team with guidance for choosing the easiest swaps that can provide higher revenues. The revenue benefits between consecutive swaps decrease monotonically, and the model terminates when an iteration is unable to provide revenue benefits beyond the threshold set by the user. In the next section, we present representative revenues per iteration.

Benefits Summary

ANA's initial interest in the development of the SOS was driven by its need for a process that would enable it to easily identify the swaps that could satisfy the RPs and OT at ANA. The SOS benefits ANA in the following ways.

- It provides a reliable and faster response; it takes less than five minutes to run, and the solutions are not only automated but optimal. Given the number of possibilities for swapping, the SOS discovers a considerable number of swaps that would have gone undetected in a manual environment.
- It reduces the required effort. Only 0.1 full-time equivalent (FTE) staff member is required for running the SOS, whereas 1.5 FTE staff members were required for identifying, reviewing, and confirming swaps before the implementation. Moreover, the additional effort needed for multiple runs is minimal.
- The SOS requires minimal human intervention. ANA has a policy of rotating staff across various operations departments, which always posed a challenge when new staff had to manually perform swaps. The adoption of SOS greatly improved the rotation process, with significantly higher productivity for new teams.
- Utilization and acceptance of swaps are higher. Up to 20 equipment changes are possible in each run, and the average number of swaps is between 10 and 14. Currently, the OT's acceptance rate for swaps is 90%.

- The acceptance from operations is better. Recommendations from the OT that are similar to the SOS outputs have better acceptance for implementation, because the SOS process is exhaustive in terms of considering all of the inclusion and exclusion criteria for any swap.

- "Spares" are better utilized. Spares are a special class of aircrafts that have high capacities and are utilized for charter services or to cover legs when a significant surge in demand occurs. Before adoption of the SOS, utilization of spares was always a challenge for ANA; however, since the SOS implementation, spare aircrafts have played a significant role in the swapping process.

- Likewise, crew constraints are designed to serve the needs of a nonunionized workforce in Japan. This differs from those models that serve the U.S. and European markets. SOS optimizes the swaps while adhering to the rules and regulations for crew.

- The SOS provides monetary benefits. The incremental benefit of using the SOS for the results shown in Table 5 for the month of July 2017 was ¥339 million.

Implementation Challenges

After the integration of SOS into ANA's system, there were challenges in adopting the recommendations for operations. The challenges included the following.

- ANA's flight schedule is highly optimized for safety and utilization. Although the SOS generated feasible swap recommendations, these recommendations were occasionally viewed as potential causes for crew and/or other operational disruptions. Eventually, this issue was resolved, because undesirable swaps were weeded out by the optimization model.
- ANA is an iconic brand in Japan and enjoys very strong customer loyalty. As such, any initiative that is perceived to have an impact on the passenger experience is heavily scrutinized. The SOS went through a thorough validation during its phased rollout.
- The SOS results impact many departments within ANA. Although revenue management and fleet planning departments were the drivers of this

Table 4. Model Characteristics for the Data Shown in Table 3

Instance	No. binary variable	No. nonzeros	No. constraints	No. iterations	Time (seconds)
1	172,564	1,013,416	4,425	16	5,877
2	127,434	750,114	5,053	24	10,209
3	163,233	958,477	6,837	24	4,735
4	174,279	1,024,765	2,545	20	5,225
5	138,835	808,258	6,035	18	1,129
6	163,233	958,477	6,837	25	1,191
7	127,434	750,114	2,649	27	1,745
8	164,218	960,963	6,948	21	11,097
9	168,740	986,372	2,628	24	7,693
10	130,368	768,120	2,649	18	1,129

Table 5. Benefits Accrued from Using the SOS for Swapping Recommendations in a Single Day

Itr	No. LOF changes	Inc. cost	Inc. pax	Inc. rev	Dec. pax	Dec. rev	Inc. profit
1	3	436	170	4,514	0	0	4,078
2	3	−452	139	3,539	−53	−1,222	2,769
3	3	655	118	2,161	0	0	1,506
4	2	−159	97	1,928	0	0	2,087
5	2	30	37	798	0	0	768
6	2	264	47	906	0	0	641
7	2	−161	39	689	0	0	850
8	3	−85	139	3,713	−118	−3,262	537
9	2	24	18	255	0	0	232
10	2	−53	83	1,549	−94	−1,451	151
11	3	−1,085	436	12,304	−116	−2,251	11,137
12	3	5	61	1,426	0	0	1,421


Notes. The first two columns, Itr and No. LOF changes, give the SOS iteration number and the number of line of flight changes for the iteration, respectively. Inc. cost is the incremental cost. Inc. pax and Dec. pax represent the increase or decrease in passenger volumes owing to the swaps, respectively, and Inc. rev and Dec. rev represent the increase or decrease in revenue, respectively. The final column, Inc. profit, represents the increment in profits. Cost, revenue, and profit are expressed in 1,000 ¥.

initiative, the results of the solution also impacted the crew as well as maintenance and sales departments. In a consensus-driven culture, it was important that all stakeholders accept the results.


- As we mention above, ANA's rotational management philosophy can lead to the need for retraining of new staff any time that a rotation occurs. Although this policy can help make the recommendations even more actionable by incorporating constraints and business rules that had not been considered previously, it requires a level of training to help the new planners get up to speed on the SOS.

Table 5 depicts some results of the SOS implementation for a single day in July 2017. ANA analyzes the projected losses and gains from each swap before approving the swap, and hence, this breakdown specifically helps ANA to choose the swaps for operations. We note that the results are monotonic with the objective function of the optimization model. However, given the built-in bonus and penalty structures for certain preferred assignments, it only follows a general but not strict monotonicity with regard to profitability. It is important for ANA to review the detailed impact of revenue and cost for each recommended swap. In iteration 2, for example, the impact of the swap led to a decrease in incremental cost of ¥452,000, an increase in projected revenues of ¥3,539,000 for

flights that were upgauged, and a decline in projected revenues of ¥1,222,000 for flights that were downgauged, resulting in a net profit of ¥2,769,000. We also note that iteration 11 results in significantly high profitability but appears only toward the end of the swap list. This ordering is because ANA has configured the system to find easy to implement swaps (for example, round trip swaps) before enabling it to find more complicated ones involving multiple flight legs. These are exactly the cases ANA likes to review while deciding to accept or reject a swap. Figure 6 indicates the revenue from each SOS iteration. As the revenue monotonically decreases over the iterations, the expected revenue is weighted against the operational challenges. The topmost iterations that make a few minor changes to the operation schedules are the most preferred.

Figure 6. (Color online) Revenue per Swapper Iteration

Notes. Each of the series represents the revenue by using SOS during the run made in a day. Revenue (in million ¥) is the additional revenue for each SOS iteration.

Figure 5. The MILP Decision Variables Are Defined as Connections in the Network, and Each Connection Represents a Legitimate Combination of an Arriving Flight and a Departing Flight

Conclusions and Future Directions

Flight swapping is a necessary but complex tactical problem, and a solution needs to minimize customer dissatisfaction and maximize revenue for airlines. Our SOS, implemented for ANA, allows for a seamless integration within the ANA scheduling process to reduce demand spillage and maximize revenue while enhancing customer satisfaction. Additionally, the SOS provides a visualization for the OT of the sequence of swaps and their estimated change in demand spillage and revenue increase. From the operational perspective, the SOS is a critical component for daily business operations because of its essential capabilities and monetary benefits. A natural application of the SOS is as a what if tool throughout the fleet-planning process. The SOS has the potential to make additional contributions in the future. One possibility would be to use the SOS as an operational tool to identify swaps that can reduce the disruptions that occur in flight schedules owing to adverse weather or equipment failures, and it can also potentially be used to estimate the swappable opportunities for each flight, which could indicate the criticality of a flight within the network. Similarly, the solution can be made more robust by incorporating the stochastic nature of projected demand and revenue. Because SOS is able to find the least disruptive swaps, it can also be modified for schedule repair purposes owing to operational disruptions, such as equipment or crew unavailability.

Acknowledgments

The authors acknowledge the following staff members at All Nippon Airways for their contributions to the project: Seiichi Takahashi, Toshiyuki Tabe, Nobuya Ikeda, Osamu Kubo, Hiroshi Kobetto, Takaaki Endo, Masanori Iwamura, Haruka Suzuki, Ryohei Ishigami, Mio Watanabe, Moemi Ando, Kotaro Noda, Akira Sekine, Tomoki Watanabe, Kosuke Abe, Hajime Fudeshima, Kenji Suzuki, and Yutaro Kurashige.


Appendix.

A.1. Decision Variables

$$X_{nf} = \begin{cases} 1 & \text{if the pair of flights } n \text{ is assigned} \\ & \text{to fleet type } f, \\ 0 & \text{otherwise.} \end{cases}$$

$$Y_{\tilde{\ell}} = \begin{cases} 1 & \text{if there is a line of flight violation} \\ & \text{for } \tilde{\ell} \in LoF_f, \\ 0 & \text{otherwise.} \end{cases}$$

Figure A.1. (Color online) The Flowchart Shows the Flow of the SOS Algorithm



Notes. The functionalities are defined in three phases, and they are realized sequentially. The different phases are the inclusion of spare flights, multiple line of flight changes, and time changes. SOS also limits the optimization run in terms of the total time and the number of iterations.

Table A.1. Sets

Set	Description
A	Set of airports, $a \in A$
F	Set of fleet types, $f \in F$
L	Set of flights, $\ell \in L$; arrival and departure times for a flight ℓ are represented as $a(\ell)$ and $e(\ell)$, respectively; $s(\ell)$ and $d(\ell) \in A$ denote the source and destination airports for a flight ℓ , respectively
T	Set of block times, $t \in T$
C	Set of pilot-crew pairs, $c \in C$
N	Set of pairs of consecutive flights in accordance with operational restrictions, $n \in N$; $n(\ell_1)$ and $n(\ell_2)$ denote the first and second legs of n , respectively
$\tilde{N} \subseteq N$	Set of pairs of overnight (consecutive) flights, where $\ell_1, \ell_2 \in \tilde{N}$, such that $e(\ell_2) < a(\ell_1)$
LoF_f	Collection of sets of LOFs for the fleets, $f \in F$

A.2. Model Formulation

$$\text{Max} \sum_{n \in N, f \in F} p_{nf} X_{nf}. \quad (\text{A.1})$$

Subject to

$$\sum_{n \in N: n(\ell_1) = \ell, f \in F} X_{nf} = 1, \quad \forall \ell \in L, \quad (\text{A.2})$$

$$\sum_{n \in N: n(\ell_1) = \ell} X_{nf} = \sum_{n \in N: n(\ell_2) = \ell} X_{nf} \quad \forall \ell \in L, f \in F, \quad (\text{A.3})$$

$$\sum_{n \in N} X_{nf} = s_f \quad \forall f \in F, \quad (\text{A.4})$$

$$\sum_{n \in \tilde{N}: d(n(\ell_2)) = a} X_{nf} \leq o_{af} \quad \forall a \in A, f \in F, \quad (\text{A.5})$$

Table A.2. Parameters

s_f	Total number of available aircraft of type f
p_{nf}	Profit for the pair of flights n assigned to fleet f
o_{af}	Number of allowed overnight maintenance fleets of type f at destination airport a
$u_{caf}(t)$	Required number of pilot-crew pairs c from airport a for fleet f during time period t within the planning horizon
M	A large integer value that denotes the maximum number of flights for any LoF_f
$mLoF$	Maximum number of allowed LOF changes

$$\sum_{n \in N: c(n(\ell_2))=a, f \in F, a(\ell_1) \leq t \leq c(\ell_2)} X_{nf} \geq u_{caf}(t) \quad \forall c \in C, a \in A, t \in T, f \in F, \quad (A.6)$$

$$\sum_{n \in N: n(\ell_1) \in \bar{\ell}, n(\ell_2) \notin \bar{\ell}} X_{nf} \leq MY_{\bar{\ell}} \quad \forall f \in F, \bar{\ell} \in LoF_f, \quad (A.7)$$

$$\sum_{f \in F, \bar{\ell} \in LoF_f} Y_{\bar{\ell}} \leq mLoF, \quad (A.8)$$

$$X_{nf} \in \{0, 1\}, \quad \forall n \in N, f \in F, \quad (A.9)$$

$$Y_{\bar{\ell}} \in \{0, 1\}, \quad \forall f \in F, \bar{\ell} \in LoF_f. \quad (A.10)$$

Based on the specified inclusions and exclusions, all possible first and second legs are combined for the set N . Along with the inclusions and exclusions, the possible departure times are also used to define the variables in the set N . Because of the growth in the number of variables owing to time changes, only a discrete limited number of minutes is allowed for time changes. The set F defines all possible fleets, where a fleet is a combination of an equipment type and configuration. For example, Airbus 320 may have three configurations (A, B, and C), where each configuration defines changes in the seat configurations. The objective function defined in (A.1) maximizes the total profit from the flight assignments, and the profit is calculated based on revenue, cost, and preference/penalty for each assignment. Constraint (A.2) ensures that every flight is covered by a fleet. Constraint (A.3) defines a balance for the assignments for each of the fleets, making sure that the assignments are in a cycle or a loop so that fleet f will follow the loop every day. Based on the capacity of each fleet f and the set $\bar{N}$ defining the last flight of the day, constraint (A.4) defines the number of LOFs to be declared for the fleet f . Similarly, the number of the overnight equipment allowed for maintenance is limited at each airport, and this limitation is captured in constraint (A.5). Certain aircrew within a time window must be covered by certain fleet types in particular airports, and this requirement is captured by constraint (A.6). Any change from the original assignment $\bar{\ell}$ of LOFs is captured using the binary variable $Y_{\bar{\ell}}$ in constraint (A.7). The total number of LOF changes is limited by the parameter $mLoF$ using constraint (A.8), and the definitions of the variables are given in constraints (A.9) and (A.10).

The model is utilized iteratively. In a first pass, the model is allowed only two LOF changes ($mLoF$ is set to two), and departure time changes are allowed. After the optimization process has found all of the two-LOF changes, the parameter

$mLoF$ is set to four, and time changes are not allowed. As a result, the optimizer will generate the LOFs with the highest revenue gains within the first few iterations and will provide the user with insights about how the revenue gains deteriorate as the iterations continue. The optimization process terminates if either the revenue gain between two consecutive iterations is less than the user-defined threshold or the runtime exceeds the stipulated maximum.

References

- Abara J (1989) Applying integer linear programming to the fleet assignment problem. *Interfaces* 19(4):20–28.
- Ageeva Y (2000) Approaches to incorporating robustness into airline scheduling. PhD thesis, Massachusetts Institute of Technology, Cambridge.
- Ahuja RK, Goodstein J, Mukherjee A, Orlin JB, Sharma D (2007) A very large-scale neighborhood search algorithm for the combined through-fleet-assignment model. *INFORMS J. Comput.* 19(3):416–428.
- Bélanger N, Desaulniers G, Soumis F, Desrosiers J, Lavigne J (2006) Weekly airline fleet assignment with homogeneity. *Transportation Res. Part B: Methodological* 40(4):306–318.
- Berge ME, Hopperstad CA (1993) Demand driven dispatch: A method for dynamic aircraft capacity assignment, models and algorithms. *Oper. Res.* 41(1):153–168.
- Burke EK, De Causmaecker P, De Maere G, Mulder J, Paelinck M, Vanden Berghe G (2010) A multi-objective approach for robust airline scheduling. *Comput. Oper. Res.* 37(5):822–832.
- Daskin MS, Panayotopoulos ND (1989) A Lagrangian relaxation approach to assigning aircraft to routes in hub and spoke networks. *Transportation Sci.* 23(2):91–99.
- Dillon J, Gopalan R, Ramachandran S, Rushmeir R, Talluri KT (1993) Fleet assignment at USAir. *Proc. TIMS/ORSA Joint National Meeting*, Phoenix, AZ.
- Eggenber N (2009) Combining robustness and recovery for airline schedules. PhD thesis, École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland.
- Froyland G, Maher SJ, Wu C-L (2013) The recoverable robust tail assignment problem. *Transportation Sci.* 48(3):351–372.
- Gopalan R, Talluri KT (1998) Mathematical models in airline schedule planning: A survey. *Ann. Oper. Res.* 76(January):155–185.
- Gu Z, Johnson EL, Nemhauser GL, Yinhu W (1994) Some properties of the fleet assignment problem. *Oper. Res. Lett.* 15(2):59–71.
- Hane CA, Barnhart C, Johnson EL, Marsten RE, Nemhauser GL, Sigismondi G (1995) The fleet assignment problem: Solving a large-scale integer program. *Math. Programming* 70(1):211–232.
- IATA (2018) IATA economics. Accessed April 14, 2018, <https://www.iata.org/publications/economics/Reports/Industry-Economic-Performance/IATA-Economic-Performance-of-the-Industry-end-year-2017-report.pdf>.
- Jarrah AI (1993) The fleet assignment model. *Proc. TIMS/ORSA Joint National Meeting Presentation*, Chicago.
- Kohl N, Larsen A, Larsen J, Ross A, Tiourine S (2007) Airline disruption management-perspectives, experiences and outlook. *J. Air Transport Management* 13(3):149–162.
- Rushmeir RA, Kontogiorgis SA (1997) Advances in the optimization of airline fleet assignment. *Transportation Sci.* 31(2):159–169.
- Sherali HD, Bish EK, Zhu X (2006) Airline fleet assignment concepts, models, and algorithms. *Eur. J. Oper. Res.* 172(1):1–30.
- Subramanian R, Scheff RP, Quillinan JD, Wiper DS, Marsten RE (1994) Coldstart: Fleet assignment at Delta Air Lines. *Interfaces* 24(1):104–120.
- Talluri KT (1996) Swapping applications in a daily airline fleet assignment. *Transportation Sci.* 30(3):237–248.

Verification Letter

Seiichi Takahashi, Senior Vice President, Marketing, All Nippon Airways, Ltd, Shiodome City Center, 1-5-2, Higashi-Shimbashi, Minato-ku, Tokyo 105-7133, Japan, writes:

“The purpose of this letter is to verify the paper titled ‘Prescriptive Analytics for Swapping Aircraft Assignments at All Nippon Airways.’ The swapper optimization suite (SOS) described in this paper is currently being used within the operation processes on a daily basis by many of our businesses. We have achieved the quantitative and qualitative benefits described in the Benefits section of this paper from using the SOS in our revenue management and operations business unit.”

Saravanan Venkatachalam received his MS and PhD in industrial and systems engineering from Texas A&M University, College Station. He is currently an assistant professor of industrial and systems engineering at Wayne State University. His research interests are in stochastic programming, large-scale optimization, and discrete event modeling and simulation, and applications of interest include supply chain management, healthcare, pricing and revenue management, and energy management.

Suresh Acharya is a professor of practice in the Robert H. Smith School of Business at the University of Maryland. He also serves as Vice President of Data Science at JDA Software, Inc. His interests are in industrial applications of optimization in the travel, transportation, and supply chain verticals. Suresh received his MS in operations research from the University of North Carolina and a MS in mathematical sciences from Clemson University.

Kenji Oba is the president of Juro International Systems, Inc., an airline solution and consulting company based in Kansas City, Missouri. Prior to this, he served in various management roles at TWA, United Airlines, Northwest Airlines, and US Airways. Kenji has significant experience developing solutions in passenger forecasting, fleet optimization, and airline marketing analytics. He received an MBA in operations research and finance from the University of Missouri, Kansas City.

Yoshinari Nakayama is an assistant manager of All Nippon Airways Co., Ltd (ANA). In his tenure at ANA, Yoshinari has served in various capacities encompassing revenue optimization, forecasting, fleet planning, and sales and marketing strategy. He received his BA in economics from Hitotsubashi University.