Optimization vs. Persistence in Scheduling: Heuristics to Manage Uncertainty in On-Demand Air Travel

Itir Z. Karaesmen
University of Maryland, College Park, MD, USA, ikaraes@rhsmith.umd.edu

Wei Yang
Long Island University, Long Island, NY, USA wei.yang@liu.edu

Pinar Keskinocak
Georgia Institute of Technology, Atlanta, GA, USA, pinar@isye.gatech.edu

We study the scheduling problem for fractional management companies (FMCs) that provide on-demand air travel services. FMCs operate in a highly uncertain and dynamic environment where frequent changes in supply (e.g., due to aircraft break-downs) and demand (e.g., trip cancellations, new trip requests) occur throughout a scheduling horizon. Certain logistical factors require FMCs to prepare their schedules in advance without complete knowledge of the aircraft availability and the trips. These advance schedules are later updated as more information becomes available. However, modified schedules are desired to remain both cost-effective and persistent (i.e., close to the original schedule). We propose a reserve-fleet heuristic that pro-actively solves the persistence problem by enforcing idleness of the aircraft in creating the original schedule. We discuss how the reserve-fleet size can be determined, what the composition of the reserve-fleet should be, and how the original schedule is modified using this reserve-fleet idea. The reserve-fleet heuristic takes advantage of the set-partitioning formulations and branch-and-price methods developed earlier for FMCs. We also present results of computational experiments that quantify the persistence vs. optimization trade-off for FMCs.

Keywords: transport scheduling, applications, heuristics, aviation, rescheduling, robustness
History: submitted January 31, 2007; revised June 8, 2007.

1 Introduction

On-demand air travel has seen steady increase over the past decade and is now common among business executives, celebrities, and others seeking the convenience and flexibility of being able to travel any time, anywhere. Charter companies and fractional management companies (FMCs) are two major providers of on-demand air travel. These companies take full responsibility for managing aircraft, crew and other aspects of air travel while their customers enjoy the benefits of private aviation. While charter companies work similar to limousine companies with no commitments or availability guarantees, FMCs provide their services through fractional ownership or jet-card programs where the customers are entitled to a certain amount of flight time on the company jets, and the service is guaranteed with as short as 4 hours of notice. On-demand air travel services by FMCs started in the late 1980s with only a handful of fractional-owners and aircraft. The industry forecasts the number of fractional-owners to reach 7000 by the end of 2007 [14], and the largest FMC in the U.S. and Europe, NetJets, currently has a fleet of 665 private jets, making the company one of the largest airlines in the world [15].

Despite the growth in the industry, it is a big challenge for FMCs to remain profitable, given their high fixed and variable costs, such as flight costs, maintenance and repair costs of the aircraft, and costs associated with crew overtime and routing the crew (e.g., from/to their home-base to/from the location of the aircraft at the start/end of their duty). While the FMC charges a fee to its customers for the actual flight time, the company bears the cost of re-positioning (routing) an aircraft from its location to a customer's point of origin. In fact, re-positioning may constitute as much as 25-35% of the total flight time of the fleet for a FMC [8, 21]. Another cost component for FMCs is subcontracting. A FMC subcontracts a trip to a charter company or to another FMC when there are no aircraft in its fleet to serve a customer's request on time. Netjets reports having spent $200 million on charters in 2005 [8].

A FMC must have efficient operations on a daily basis to manage all these costs. Scheduling of the aircraft and the crew to satisfy customers' requests remain the two most critical operational decisions for FMCs. Several researchers and practitioners have studied different aspects of the scheduling problem for FMCs, including scheduling of the aircraft only [11, 4, 5, 6] and combined aircraft and crew scheduling [16, 21]. Martin et al. [13], Hicks et al. [9], Karaesmen et al. [10] and Yang et al. [20] introduce decision support tools (DST) for FMCs. Several models and methods are proposed in these papers to obtain optimal or near-optimal schedules assuming perfect information about customers' requests, availability of aircraft and the crew. These models play an important role in meeting the FMCs need to prepare an advance schedule, often called a *Master Schedule*, that is required by some airports 24-hours in advance, and also needed for logistical planning of scheduled maintenance activities and crew swaps.

Given that the nature of the business and customer characteristics make it almost impossible to predict the demand in advance, FMCs modify the Master Schedule as new information (e.g., trip cancellations, new trip requests, aircraft break-downs, early completion of scheduled maintenance for aircraft) becomes available [10]. Based on the operations of a FMC, 20% of trips are requested with as little as 4 hours notice, and 30% of the trips requested are changed at least once within 48 hours of flight time [9]. A FMC is faced with uncertainty not only in demand, but also in supply. A total of 49 mid-day and 104 overnight unscheduled maintenance events were recorded during a one-month period for a FMC that has 35 aircraft [21]. It is often desirable that a modified schedule remains close to the Master Schedule while keeping the costs low.

Although rescheduling is of utmost importance for FMCs, so far it has received limited attention in the literature [9, 13]. For a slightly different business model for on-demand air transportation, a local search heuristic is suggested to modify a Master Schedule [6]. Our goal is to find scheduling solutions that achieve a good tradeoff between optimality and *persistence*, i.e., the modified solution should be nearly optimal with respect to standard criteria without deviating too much from the previous solution. Rather than penalizing deviations from a given schedule (as in [9]), we propose an approach that is pro-active in dealing with schedule persistence: Our heuristic enforces idleness of the aircraft in creating a Master Schedule in order to retain scheduling flexibility to deal with future supply/demand issues.

There are other industries in which deterministic optimization and/or scheduling is carried out despite uncertainty in demand and supply. For example, commercial airlines face supply disruptions due to aircraft or crew unavailability and rely on emergency or recovery tactics (e.g., delaying flight departures, cancelling flights, re-routing aircraft or passengers, re-assigning crews) to get their operations back inline with their original schedules [2, 7, 22]. Traditionally, commercial airlines have first created a schedule in order to minimize costs or maximize aircraft utilization, then tried to retain persistency via recovery efforts. More recently, the concept of *robustness* of airline schedules has received attention; i.e., airlines can strategically plan for disruptions while creating their schedules via robust fleet assignment [18], robust aircraft routing [12] and robust crew assignment [17]. In contrast to commercial airlines, FMCs face both supply and demand disruptions and need to consider optimization and persistence simultaneously in scheduling the aircraft and the crew. On the up side, since FMCs operate on-demand (they can be viewed as "unscheduled airlines") they have more flexibility in terms of aircraft and crew re-assignments.

Schedule perturbations and obtaining robust schedules are also of concern in manufacturing [1, 19]. The idea of *under-capacity scheduling* [1] is in spirit close to what we are proposing for a FMC. Unfortunately, the models and solution methods developed for production scheduling are not adequate for a FMC because of the differences in the objectives and the operational characteristics.

In this paper, we propose a heuristic to obtain persistent and cost-effective schedules for FMCs. Our approach is *unique* in the way that it aims to create a schedule that is flexible enough to deal with future supply and demand changes. We achieve such scheduling flexibility by enforcing idleness of the aircraft in creating a Master Schedule and call this the *reserve-fleet heuristic*. We discuss how the reserve-fleet size can be determined, what the composition of the reserve-fleet should be, and how the original schedule is modified using the reserve-fleet heuristic. We also report results of computational experiments that quantify the cost-effectiveness vs. persistence of our heuristic. To the best of our knowledge, our work is *the first* that studies the trade-off between optimization and persistence for FMCs. In the remainder of the paper, we first define the scheduling problem in Section 2. Our heuristic is introduced in Section 3, and results of computational experiments are reported in Section 4. We conclude and discuss future research directions in Section 5.

2 Problem Definition

Given a Master Schedule created at time $t = 0$ (Stage-0) for a pre-defined planning horizon, consider the state of the system at time $t > 0$ (Stage-1). Part of the schedule may be activated by that time (i.e., some aircraft may already be serving trips or in the process of repositioning), and we may have new information about trip arrivals, changes or cancellations, as well as changes in aircraft availability. Given the new information, our goal is to create a modified schedule which stays close to the original one (in terms of trip-aircraft pairings) as much as possible while still minimizing repositioning and subcontracting costs.

Throughout the paper, index i denotes aircraft and j represents trips. We denote $N_0 = \{1, ..., n_0\}$ and $M_0 = \{1, ..., m_0\}$ as the set of all available aircraft and requested trips, respectively, at time $t = 0$. We further use the following notation:

Ω_i : the set of feasible routes for aircraft i, $i \in \bar{N}_0$,
p : a feasible route, $p \in \Omega_i$,
c_{ip} : the total cost of taking route p by aircraft i, $i \in \bar{N}_0$ and $p \in \Omega_i$,
a^i_{jp} = 1, if trip j is on route p; 0, otherwise, $p \in \Omega_i$, $i \in \bar{N}_0$ and $j \in M_0$,
θ_{ip} = 1, if aircraft i takes route p; 0, otherwise, $p \in \Omega_i$ and $i \in \bar{N}_0$

where $\bar{N}_0 = \{1, 2, ..., n_0 + m_0\}$ is the set of all available aircraft such that $\{1, ..., n_0\}$ represents the fleet and $\{n_0 + 1, ..., n_0 + m_0\}$ represents charter aircraft for subcontracting the trips at Stage-0. For each $i \in \{n_0 + j : j = 1, ..., m_0\}$, Ω_i consists of a single route that takes only trip j with the route cost being the cost of subcontracting trip j.

In our earlier research, we proposed the following set partitioning formulation (SP) to create the Master Schedule for the combined aircraft and crew scheduling problem, where each aircraft is coupled with its crew [20] (i.e., aircraft availability is constrained not only by maintenance related activities but also allowable duty and rest periods for the crew):

$$(SP) \quad Z_0^{SP} = \min \quad \sum_{i \in \bar{N}_0} \sum_{p \in \Omega_i} c_{ip}\theta_{ip} \tag{1}$$

$$s.t. \quad \sum_{i \in \bar{N}_0} \sum_{p \in \Omega_i} a^i_{jp}\theta_{ip} = 1, \qquad j \in M_0 \tag{2}$$

$$\sum_{p\in\Omega_i} \theta_{ip} \leq 1, \qquad i \in \bar{N}_0 \tag{3}$$

$$\theta_{ip} \in \{0,1\}, \qquad i \in \bar{N}_0, \text{ and } p \in \Omega_i. \tag{4}$$

Constraints (2) ensure that each trip is covered and constraints (3) ensure that an aircraft is assigned to at most one route. Note that the problem is always feasible because the trips can be subcontracted. In [20], we used a branch and price method which combines column generation with branch-and-bound to obtain a near-optimal integral solution to SP; we refer to that procedure as BP. Note that the routes for each aircraft can easily be created in this problem by pre-processing the input data to identify the aircraft-trip compatibility and the availability of an aircraft to take a trip or a pair of trips back-to-back (see [10, 20] for pre-processing aircraft- and trip-related information and [13] for pre-processing crew-related information). Our heuristic takes advantage of the SP formulation and the BP method; details are provided below.

3 Reserve-Fleet Heuristic

As demand and supply changes are observed, the Master Schedule needs to be modified to remain feasible and cost-effective. Re-optimizing the schedule by re-solving SP may result in a drastically different output, even with a small change to the input data. This phenomenon is quite common in mathematical programming [3], and techniques for incorporating persistence into optimization models include (i) fixing a subset of variables to preferred values, (ii) converting certain variables into persistent variables that incur linear penalties for deviating from their preferred values (as in Hicks et al. [9]), and (iii) using persistent constraints to ensure that aggregate quantities achieve preferred values or ranges of values and to allow penalized deviations to occur. Notice that all these approaches deal with persistence only at the time of rescheduling (Stage-1). In contrast, we propose a pro-active approach (at Stage-0) to achieve schedule persistence.

We propose to create a Master Schedule at Stage-0 which may be slightly suboptimal in terms of costs but is flexible enough to accommodate future changes without significantly affecting the existing trip-aircraft assignments. The main idea is to use only a fraction of the fleet in creating the Master Schedule and keep the remaining aircraft idle as a *reserve-fleet* to respond to future demand/supply changes. Recall that a typical FMC only knows about 75-80% of the trips while creating its Master Schedule [20], and the reserve-fleet can help reduce the number of future schedule perturbations. While this is a simple idea, there are three key issues: What is the right size for the reserve-fleet? Which aircraft should be in the reserve? How does one adjust/modify a schedule that was created using the reserve-fleet idea? We provide the answers below assuming the FMC owns a homogeneous fleet.

Determining the Reserve-Fleet Size: Let r be the size of the reserve as a fraction of the fleet, i.e., $(1-r)n_0$ is the number of aircraft considered for scheduling at Stage-0. Due to the complexity of the FMC's problem and the need for practical methods, we favor a systematic analysis to determine the value of r. The company can use historical data (or use simulations to generate random data that is representative of company's operations) to see how different values of r would have affected the scheduling performance in the past (or in the simulations). Given demand and supply information, one can re-solve the scheduling problems using the reserve-fleet idea and compute estimates of costs and number of schedule perturbations given a value of r. Experimenting with different values of r enables the decision-makers to choose the reserve-fleet size that provides the best trade-off between optimization and persistence. We illustrate the systematic approach on randomly generated data in Section 4. However, we first need to decide which aircraft goes into the reserve for a given r.

Determining the Composition of the Reserve-Fleet: Suppose r is given. We can solve SP using BP to create a Master Schedule if we know a priori which aircraft is kept for the reserve. This can be done by randomly choosing the aircraft that goes into the reserve. While too simple, this randomization approach may provide reasonable results due to the complexity of the problem. We call this the *randomized-reserve-fleet (RRF)* approach. Another approach is to add $\tilde{m} = \lceil rn_0 \rceil$ (the smallest integer exceeding the reserve fleet size) *dummy trips* to set M_0. These trips can be taken by any of the aircraft in set $\bar{N}_0$, and their repositioning and flight costs are zero. However, they cannot be on the same route with any other trip, and they have extremely high subcontracting costs. In other words, $\tilde{m}$ feasible routes with extremely high costs are added to the set Ω_i for each $i \in \bar{N}_0$. Using these expanded set of trips and routes, we re-formulate SP. The optimal solution to this modified set-partitioning problem assigns $\tilde{m}$ company aircraft to the dummy trips to avoid the high cost of subcontracting. The aircraft assigned to these dummy trips provide the optimal composition for the reserve fleet. We call this approach the *optimized-reserve-fleet (ORF)*. We test the effectiveness of both RRF and ORF using computational experiments.

Rescheduling at Stage-1 Using the Reserve-Fleet: The goal of the reserve fleet heuristic is to give flexibility to schedulers in the future. The new schedule can be computed by re-solving SP with the updated information at Stage-1, by (partially) fixing trip-aircraft assignments from Stage-0 *except for* subcontracted trips, trip cancellations, or aircraft breakdowns. Note that the fixed assignments only change the column generation step at BP. Feasible routes can be obtained easily if the fixed assignments are included in pre-processing in order to determine feasible aircraft-trip assignments [10].

4 Computational Results

The purpose of our computational experiments is to quantify the optimization and persistence trade-off, to determine the appropriate reserve-fleet size based on that trade-off, and to benchmark the performance of the reserve-fleet heuristic. We use randomly generated data and work in a simplified environment. We assume that both Stage-0 and Stage-1 problems are solved before the Master Schedule is activated; this is typical of FMCs. We assume the only uncertainty is on the demand side, and that the entire fleet (of size n_0) is available at both stages. We also assume there are no cancellations. Hence, the only new information at Stage-1 is regarding additional trips. As a comparison to the performance of the reserve-fleet heuristic, we define an approach called *re-optimization* which solves SP at Stage-0 using the entire fleet, and resolves SP with updated information at Stage-1 without fixing any of the prior trip-aircraft assignments. The Stage-1 solution of this approach is the *hindsight optimal* solution that minimizes the total costs. However, re-optimization at Stage-1 is expected to lead to a high number of schedule perturbations.

We use the following parameters: The scheduling horizon is 36 hours. Only 80% of the trips are known at Stage-0. Aircraft locations and trip origins/destinations as well as departure times are generated randomly: We first choose 100 cities that are uniformly distributed on a 100 by 100 grid. The travel time between two adjacent points on the grid is assumed to be 3 minutes, therefore the longest flight is about 7 hours. We randomly select a departure location and a destination location for each trip with flight time greater than 30 minutes. Setting the time point for optimization at zero, we are planning the events between now and 2160 minutes later. The earliest available time of each aircraft is uniformly distributed between 0 and 1200 minutes. The departure time for a trip is uniformly chosen between 0 and 1800 minutes. We denote each experiment by $[r, (1-r)n_0, m_0, n_0, m_1]$ where $(1-r)n_0$ is the number of aircraft used in scheduling at Stage-0 and $m_1 = \frac{m_0}{0.8}$ is the total number of trips at Stage-1. We generate 10 random instances for each experiment. All instances in an experiment have the same m_0 and m_1, but trips differ in trip departure times and origins/destinations from one instance to another.

We compute two measures in each experiment: (i) The *average optimality gap* of the reserve-fleet heuristic is the percentage difference between the hindsight-optimal costs and the cost of the reserve-fleet solution. (ii) The *average persistence* is the average of the ratio of number of trips that have the same assignment at Stage-0 and Stage-1 (i.e., if a trip j is assigned to an aircraft $i \in N$, then j remains assigned to i at Stage-1, or if j is chartered at Stage-0, it remains chartered at Stage-1) to number of trips known at Stage-0 (m_0); this measure is an indication of the FMC's 'ability to schedule in advance' and is based on the total number of schedule perturbations. We also report the average number of trips subcontracted by re-optimization and reserve-fleet solutions at Stage-1.

Table 1 summarizes the results of our experiments. We implemented both the randomized procedure RRF and optimization-based procedure ORF. We only report results of experiments where $n_0 = 50$ and $m_1 \in \{50, 100, 150\}$. Results of our experiments with higher fleet sizes are omitted because the load factor (relative size of the fleet to the number of trips) impacts our results more than the actual fleet size. Here are the main observations from Table 1: (i) RRF and ORF have similar performance; neither one dominates the other. (ii) The reserve-fleet approach is most appropriate for moderately high load factors ($l = 2$ when $n_0 = 50, m_1 = 100$). This is expected because when the load factor is very low, there are unused aircraft and demand/supply changes can be handled fairly easily with this excess capacity. When the load factor is high, a significant number of trips are chartered. In that case, the reserve fleet further limits the capabilities at Stage-0, leads to a higher number of subcontracted trips. Consequently, the solution in Stage-1 tries to reassign those subcontracted trips to company aircraft, adversely affecting the schedule persistence. (iii) Average persistence decreases as r increases in both RRF and ORF. However, the average optimality gap fluctuates with no apparent pattern. This is expected: First of all, when $r = 1$, the optimality gap of the reserve-fleet heuristics is 0% because Stage-1 solution of reserve-fleet is the same as the hindsight optimal solution in that case. Hence, as $r \to 1$, the optimality gap gets smaller. Second, there is a delicate balance between the inefficiency at Stage-1 due to limited availability of aircraft because of prior assignments, and the inefficiency at Stage-0 due to limited number of aircraft. Depending on the size of the reserve fleet, either source of inefficiency becomes the dominant factor in the optimality gap. The randomness in the data sets make it difficult to detect when Stage-1 or Stage-0 inflexibility is more important, and how the size of the reserve fleet affects that. Note that the differences in the optimality gaps can also be explained by the total number of charter aircraft at Stage-1. The number of charters used also fluctuate with r.

The fleet size can be chosen by detailed examination of the results in Table 1. Different rules need to be developed for different load factors. For instance, when $m_1 = 150$, there is an improvement in optimality gap when r increases from 0 to 10%. However, further increases in r do not yield significant decreases in optimality gap and lead to low schedule persistence. In that respect, r can be 10-20% for high load factors. This choice guarantees 5% gap in costs and 50% schedule persistence. When $m_1 = 100$, $r = 20\%$ guarantees about 68-70% schedule persistence and 6-9% optimality gap. When $m_1 = 50$, r can take any value from 0 to 30% to guarantee 4.8-6% optimality gap and 78-80% schedule persistence.

5 Conclusion

While there has been recent research on scheduling problems arising in the on-demand travel industry, the issue of schedule persistence has often been neglected by researchers. In this paper, we focused on the rescheduling problem of FMCs. We developed a heuristic to deal with both

Experiment	RRF			ORF			Re-optimization		
$r, (1-r)n_0, m_0, n_1, m_1$	Opt. Gap (%)	Pers. (%)	No. Charters	Opt. Gap (%)	Pers. (%)	No. Charters	Opt. Gap (%)	Pers. (%)	No. Charters
0.00, 50, 40, 50, 50	4.78%	80.00%	11.4	4.57%	80.00%	11.3	0.00%	62.60%	10.2
0.10, 45, 40, 50, 50	5.81%	79.40%	11.6	4.57%	80.00%	11.3	0.00%	62.60%	10.2
0.20, 40, 40, 50, 50	4.80%	78.00%	10.7	4.57%	80.00%	11.3	0.00%	62.60%	10.2
0.30, 35, 40, 50, 50	7.23%	77.60%	11.4	4.58%	80.00%	11.3	0.00%	62.60%	10.2
0.40, 30, 40, 50, 50	4.47%	76.80%	10.3	4.61%	80.00%	11.3	0.00%	62.60%	10.2
0.50, 25, 40, 50, 50	6.78%	75.20%	10.9	5.12%	79.80%	11.3	0.00%	62.60%	10.2
0.60, 20, 40, 50, 50	6.94%	68.00%	10.8	6.07%	77.40%	11.4	0.00%	62.60%	10.2
0.70, 15, 40, 50, 50	6.79%	57.60%	10.5	4.41%	66.40%	11.1	0.00%	62.60%	10.2
0.80, 10, 40, 50, 50	5.65%	45.40%	10.3	0.98%	53.00%	10.7	0.00%	62.60%	10.2
0.00, 50, 80, 50, 100	8.78%	73.63%	33.0	8.29%	70.75%	32.5	0.00%	50.70%	29.5
0.10, 45, 80, 50, 100	6.96%	72.22%	32.8	8.17%	70.63%	32.4	0.00%	50.70%	29.5
0.20, 40, 80, 50, 100	5.53%	68.38%	31.9	8.21%	70.50%	32.3	0.00%	50.70%	29.5
0.30, 35, 80, 50, 100	5.88%	62.75%	32.6	7.85%	68.75%	32.2	0.00%	50.70%	29.5
0.40, 30, 80, 50, 100	6.95%	55.97%	31.8	7.79%	64.38%	32.2	0.00%	50.70%	29.5
0.50, 25, 80, 50, 100	6.65%	49.88%	32.3	7.38%	56.75%	31.8	0.00%	50.70%	29.5
0.60, 20, 80, 50, 100	5.51%	40.38%	32.1	8.56%	47.75%	32.6	0.00%	50.70%	29.5
0.70, 15, 80, 50, 100	5.31%	29.00%	31.8	4.91%	37.50%	31.5	0.00%	50.70%	29.5
0.80, 10, 80, 50, 100	4.17%	19.63%	31.8	1.49%	25.13%	30.9	0.00%	50.70%	29.5
0.00, 50, 120, 50, 150	6.15%	64.58%	69.1	6.01%	58.50%	69.1	0.00%	48.53%	67.4
0.10, 45, 120, 50, 150	5.01%	59.50%	69.5	5.48%	57.42%	68.9	0.00%	48.53%	67.4
0.20, 40, 120, 50, 150	4.11%	53.67%	68.8	5.15%	55.75%	68.4	0.00%	48.53%	67.4
0.30, 35, 120, 50, 150	4.58%	47.75%	70.2	4.94%	51.83%	68.5	0.00%	48.53%	67.4
0.40, 30, 120, 50, 150	4.02%	40.33%	69.9	4.87%	46.75%	69.4	0.00%	48.53%	67.4
0.50, 25, 120, 50, 150	3.69%	34.67%	69.5	4.62%	40.75%	68.9	0.00%	48.53%	67.4
0.60, 20, 120, 50, 150	3.53%	28.58%	69.3	3.81%	33.58%	69.0	0.00%	48.53%	67.4
0.70, 15, 120, 50, 150	3.27%	20.50%	69.9	3.04%	25.25%	69.1	0.00%	48.53%	67.4
0.80, 10, 120, 50, 150	2.79%	14.33%	68.9	0.02%	16.83%	68.2	0.00%	48.53%	67.4

Table 1: Average performance of re-optimization, randomized reserve-fleet and optimized reserve-fleet heuristics

persistence and optimization issues in scheduling. Our reserve-fleet approach creates a schedule that is flexible enough to deal with future supply and demand changes by enforcing idleness of the aircraft in developing a Master Schedule. We discussed different ways of implementing the reserve-fleet idea. Finally, we conducted computational experiments that quantify the cost-effectiveness vs. persistence of re-scheduling efforts. Based on our numerical results, we discussed how the reserve-fleet size can be determined in a systematic way.

Our novel idea of the reserve-fleet is very practical. We have shown its use for a homogeneous fleet in this paper. Our future work will involve extending this idea to a fleet that includes multiple types of aircraft (where customers can be upgraded/downgraded), and comparing the reserve-fleet heuristics to other standard methods that are used to achieve persistence in optimization (e.g., penalizing deviations from an initial solution). Investigating the effect of both supply and demand uncertainties on the reserve-fleet heuristic also warrants further study.

References

[1] H. Aytug, M.A. Lawley, K. McCay, S. Mohan and R. Uzsoy (2005), Executing production schedules in the face of uncertainties: A review and some future directions, *EJOR* **161**(1), 86 – 110.

[2] F. Bian, E.K. Burke, S. Jain, G. Kendall, G.M. Koole, J.D. Landa Silva, J. Mulder, M.C.E. Paelinck, C. Reeves, I. Rusdi and M.O. Suleman (2003), Making Airline Schedules More Robust, *MISTA 2003 Conference Proceedings*, 678 – 695.

[3] G. Brown, R. Dell and R. Wood (1997), Optimization and Persistence, *Interfaces* **27**(5), 15 – 37.

[4] A. Erdmann, A. Nolte, A. Noltemeier and R. Schrader (2002), Modeling and Solving an Airline Schedule Generation Problem *Annals of Operations Research,* **107**(1), 117 – 142.

[5] D. Espinoza, R. Garcia, M. Goycoolea, G.L. Nemhauser, and M.W.P. Savelsbergh (2006), Per-Seat, On-Demand Air Transportation Part I: Problem Description and an Integer Multi-Commodity Flow Model,Working Paper, School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA.

[6] D. Espinoza, R. Garcia, M. Goycoolea, G.L. Nemhauser, and M.W.P. Savelsbergh (2006). Per-Seat, On-Demand Air Transportation Part II: Parallel Local Search. Working Paper, School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA.

[7] J.E. Filar, P. Manhem and K. White (2001), How Airlines and Airports Recover from Schedule Perturbations: A Survey, *Annals of Operations Research,* **108**, 315 – 333.

[8] D. Foust (2007). One Jet, 16 Owners, Big Problems. *BusinessWeek*, January 29, 2007.

[9] R. Hicks, R. Madrid, C. Milligan, R. Pruneau, M. Kanaley, Y. Dumas, B. Lacroix, J. Desrosiers and F. Soumis (2005), Bombardier Flexjet Significantly Improves Its Fractional Aircraft Ownership Operations, *Interfaces* **35**(1), 49 – 60.

[10] I. Karaesmen, P. Keskinocak, S. Tayur and W. Yang (2005), Scheduling Time – Shared Aircraft: Models and Methods for Practice, *Proceedings of MISTA 2005: The 2nd Multidisciplinary Conference on Scheduling: Theory and Applications,* 18 – 21 July, 2005, New York, NY, USA.

[11] P. Keskinocak and S. Tayur (1998), Scheduling of Time – Shared Jet Aircraft, *Transportation Science* **32**(3), 277 – 294.

[12] S. Lan, J.P. Clarke and C. Barnhart (2005), Planning for robust airline operations: optimizing aircraft routings and flight departure times to minimize passenger disruptions. to appear in *Transportation Science.*

[13] C. Martin, D. Jones and P. Keskinocak (2003), Optimizing On-Demand Aircraft Schedules for Fractional Aircraft Operators, *Interfaces* **33**(3), 22 – 35.

[14] National Business Aviation Association (2004), *NBAA Business Aviation Fact Book,* Washington DC, available at *http://web.nbaa.org/public/news/stats/factbook/2004/2004factbook.pdf.*

[15] NetJets (2006), *NetJets Fast Facts.* Woodbridge NJ, December 2006, available at *http//www.netjets.com/News%20and%20Info/pdfs/NetJets%20Fast%20Facts%2012 – 06.pdf.*

[16] D. Ronen (2000), Scheduling Charter Aircraft, *Journal of Operational Research Society* **51**, 258 – 262.

[17] Shebalov, S. and D. Klabjan (2004), Robust airline crew scheduling: move – up crews. *Transportation Science,* **40**, 300 – 312.

[18] Smith, B.C. and E.L. Johnson (2006), Robust Airline Fleet Assignment: Imposing Station Purity Using Station Decomposition. *Transportation Science,* **40**(4), 497 – 516.

[19] G.E. Vieira, J.W. Herrmann and E. Lin (2003). Rescheduling manufacturing systems: a framework of strategies, policies, and methods, *Journal of Scheduling,* **6**(1), 39 – 62.

[20] W. Yang, I. Karaesmen, P. Keskinocak and S. Tayur (2006), Aircraft and Crew Scheduling for Fractional Ownership Programs, Working Paper, Long Island University, Long Island, NY.

[21] Y. Yao, O. Ergun, E. Johnson, W. Schultz and J.M. Singleton (2004), Strategic Planning Problems in Fractional Aircraft Ownership Programs, Working Paper, School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA.

[22] G. Yu, M. Arguello, G. Song, S. N. McCowan and A. White (2003), A new era for crew recovery at Continental Airlines, *Interfaces,* **33**(1), 5 – 23.