

Modifying Lines-of-Flight in the Planning Process for Improved Maintenance Robustness

Marcial Lapp and Amy Cohn
Industrial & Operations Engineering Dept.
University of Michigan
1205 Beal Ave.
Ann Arbor, MI 48109-2117
{mlapp,amycohn}@umich.edu

January 14, 2010

Abstract

As part of their tactical planning process, airlines construct *lines-of-flight*, which are day-long sequences of flights to be flown by a single aircraft. Subsequently, airlines assign specific aircraft (*tails*) to these lines-of-flight. Such assignments create multi-day aircraft routings and thus determine when each aircraft will remain overnight at a maintenance station. Maintenance overnights, in turn, provide opportunities for aircraft to undergo required routine maintenance checks. Disruptions, however, often result in changes to these planned routings, which can in turn prevent an aircraft from reaching its planned maintenance event. Our goal is to develop lines-of-flight that are better able to recover from such disruptions. In order to do so, we first introduce a new metric, which we call *maintenance reachability* (MR). We claim that plans with higher levels of *maintenance reachability* provide greater opportunities for maintenance recovery when the planned routings are disrupted. Next, we develop models to increase the *maintenance reachability* of a given set of lines-of-flight. We then conclude with computational results, based on data from a major U.S. carrier, that demonstrate both the tractability of our approach and the potential improvements in maintenance robustness that can be achieved.

1 Introduction

Before presenting our approach to developing planned *lines-of-flight* (LOF) that are more robust in responding to operational disruption, we first describe the planning process at a major U.S. carrier, with whom we worked on this research. This overview will provide the necessary background to illustrate why robustness in maintenance schedules is essential.

After an airline has determined a set of flights to offer over a given time period, a process usually completed by the marketing/capacity planners of the airline, an optimization process takes place. The objective of this process is to build cost-effective lines-of-flight, each of which specifies a day-long sequence of flights to be assigned to a single aircraft (Gopalan and Talluri 1998). Subsequently, seven days prior to departure, *maintenance routings* are built comprised of multi-day assignments of aircraft (*tails*) to lines-of-flight. An example of a line-of-flight is shown in Figure (1). Here, a LOF directs an aircraft to fly DTW→BOS→FLL→DFW over the course of a day. It is these lines-of-flight that we will examine and enhance in this paper.

Although the planned LOFs and corresponding routings assure that maintenance can be achieved, these plans are often disrupted in practice. In such cases tail assignments, routings, and even the LOFs themselves may need to be modified to regain maintenance feasibility. These disruptions tend to manifest themselves in the form of unexpected equipment changes and other measures to counter delay propagation. For example,




Figure 1: Example of Lines-of-Flight

an operation controller may choose to alter the tail assignment to mitigate a delay propagation effect by performing an aircraft substitution. Such a scenario is presented as a numerical example in §2.

In our work, we focus on frequent maintenance checks that are mandated by the Federal Aviation Administration (FAA). These checks, known as A-Checks, must be carried out roughly every seven days (Sriram and Haghani 2009). From an airline perspective, this means that on its seventh day, each aircraft must end its day-of-operation at a maintenance base to be cleared for operation the next day.

On any given morning, each station in the flight network will have a number of outgoing LOFs, some of which will terminate in a maintenance station. We will refer to such lines as *maintenance lines-of-flight* or MLOFs. If a station has at least as many MLOFs as it does *day-7 aircraft* (aircraft starting the day at the station, which require maintenance at the end of the day), then the schedule can execute as planned. However, should the station hold more day-7 aircraft than MLOFs, potentially costly, over-the-day swaps with other lines-of-flight might be required to ensure that all aircraft reach their required maintenance. We use the term *maintenance misalignments* to refer to a day-7 aircraft that does not have an available MLOF. From a *planning* perspective, it is of interest to airlines to ensure that when an aircraft begins its 7th day of operation, there is a high likelihood that an MLOF exists, which it can be assigned to, allowing it to end the day at a maintenance station.

In this paper, we advocate strategic decisions during the planning process in which these MLOFs are created, in order to maximize the likelihood that an aircraft due for maintenance is able to reach a maintenance station. As intuition may suggest, a station that has a larger number of LOFs should have a corresponding larger number of MLOFs. Constructing the “optimal” set of MLOFs across stations is the objective of our proposed optimization problem for use within the schedule planning process (when LOFs are first constructed).

The contributions of this paper are two-fold. The first contribution is a new robustness metric called *maintenance reachability* (MR), with which we are able to analyze a set of LOFs and quantify its ability to withstand unplanned changes from a maintenance feasibility standpoint. Maintenance reachability in contrast with maintenance misalignment, is the expected number of day-7 aircraft that can reach maintenance at the end of the day without modifying the existing lines-of-flight. The second contribution is the development of an optimization model that is able to construct improved LOFs so as to maximize maintenance reachability. We do so by minimizing the total number of maintenance misalignments.

In §2, we provide a detailed problem description, illustrating different operational scenarios with numerical examples of maintenance reachability calculations. Next, in §3, we review the maintenance routing literature, specifically emphasizing routing issues which focus on maintenance robustness. In §4, we present a model for optimizing the maintenance reachability of a planned set of LOFs. In §5, we provide results from our model using schedules based on actual U.S. carrier data. Finally, §6 offers concluding thoughts and motivates future work to be done in the area of maintenance robustness.

2 Problem Description

In this section, we provide numerical examples to demonstrate *why* airlines modify a scheduled plan on the day-of-operations. In addition, we illustrate the implications of changing planned lines-of-flight (LOFs) with

respect to pre-planned maintenance. Finally, we define (numerically) what we mean by the *maintenance robustness* of a given flight plan, the objective over which we optimize in §4.

To assist in defining the concepts presented in this paper, we begin with a series of illustrations of a small flight plan. This flight plan, depicted in Table (1), is indicative of properties commonly found in actual, major U.S. carrier plans. In this example, flight sets $\{1, 4\}$, $\{2, 3\}$, $\{5\}$, $\{6\}$ and $\{7\}$ each form a LOF. Furthermore, we feature three aircraft, A, B and C , each of which has been assigned to a particular maintenance route during the tail-assignment phase.

Table 1: Example Aircraft Schedule

Flight	Day	Origin	Dest.	STD	STA	Tail
1	1	SEA	TUS	08:10	10:30	A
2	1	DTW	TUS	09:05	10:35	B
3	1	TUS	BOS	11:05	15:00	B
4	1	TUS	ORD	11:20	15:30	A
5	2	BOS	DAL*	09:00	14:30	B
6	2	ORD	SEA	09:30	12:45	A
7	2	BOS	MCO*	10:30	13:45	C

*MCO and DAL are maintenance stations.

2.1 Aircraft Swap Advantages

Situations can arise where it may be advantageous for an operations controller to deviate from the scheduled plan. For example, referring back to Table (1), suppose that Flight 2 from DTW to TUS is delayed by 15 minutes. Incorporating a 30 minute turn-time, the earliest the subsequent flight, Flight 3, is be able to depart due to the delay is at 11:20, 15 minutes past its scheduled departure time. Suppose also that Flight 1 landed on time and thus, even with a 30 minute turn-time, could satisfy an 11:05 departure. The operation controllers now have the choice of performing an aircraft swap (assuming the aircraft are swap-compatible). If an aircraft swap is performed, then the planned schedule is changed as illustrated in Table (2). The result of this aircraft swap is that the delay is absorbed, and both Flights 3 and 4 are able to depart on time. On the other hand, the maintenance plan for each aircraft has now been disrupted, because the swap caused the aircraft to exchange their respective routings. In other words, if it turns out that aircraft B was destined for a maintenance station, DAL in this instance, to receive its regularly-timed maintenance check, but now follows a different route that ends in SEA, potential maintenance feasibility issues can result. Such problem situations are explored explicitly in the next section.

Table 2: Example Aircraft Schedule *After* Performing Swap

Flight	Day	Origin	Dest.	STD	STA	Tail
1	1	SEA	TUS	08:10	10:30	A
2	1	DTW	TUS	09:05	10:50	B
3	1	TUS	BOS	11:05	15:00	A
4	1	TUS	ORD	11:20	15:30	B
5	2	BOS	DAL*	09:00	14:30	A
6	2	ORD	SEA	09:30	12:45	B
7	2	BOS	MCO*	10:30	13:45	C

*MCO and DAL are maintenance stations.

2.2 Swap-based Maintenance Implications

Performing swaps, as presented in the previous section, has the advantage of mitigating delay-propagation; however, such advantages come with a hidden cost. Pre-planned maintenance schedules are no longer followed, which can drive costs and further disruptions as demonstrated in this section.

In the following example based on Table (1) and also depicted in Figure (2), we feature two routes, examined over a two-day period. The first route departs from DTW and performs an intermediate stop in TUS. In addition, another route starting in SEA also arrives in TUS, at the same time as Route 1. As seen in the example in Table 2, the two aircraft assigned to these routes are potential candidates for a swap. We suppose that a situation, similar to the previous example, exists, where aircraft are swapped to mitigate delay propagation, but now also suppose the aircraft flying Route 1 is a *day-6* aircraft, referring to an aircraft that is on day 6 of its maintenance rotation. Given that aircraft must be maintained every 7 days, the aircraft in this example must receive maintenance at the end of the day *tomorrow*.

Performing a swap in this situation may have costly consequences for scheduled maintenance. The (day-6) aircraft flying Route 1 now ends the day in ORD. It turns out that the following day, day 7, the only outgoing line is destined for SEA which is *not* a maintenance base, i.e. no MLOF exists. This example is seen graphically in Figure (2). In this situation, we have no possible maintenance outlets for the aircraft, and thus we must find a swap opportunity over the course of the day to ensure this aircraft enters required maintenance. In other words, the intermediate stations indicated in Figure (2) by circles 1 and 2 must intersect with another line that does end in a maintenance station. Furthermore, if an intersection with a maintenance line is found, then this MLOF must not already be required to bring another tail to a maintenance station, such that we can perform a swap and allow the aircraft departing ORD on the beginning of day 7 to enter maintenance at the end of the day.




Figure 2: Aircraft Routing on Day 6 without Maintenance Opportunity

By using swap opportunities at either Station 1 or Station 2 in Figure (2) to bring the aircraft to its required maintenance station, we are changing our operational schedule again, i.e. forcing another aircraft to change its route, causing further network effects. Additionally, we may be breaking a *through-flight*, causing connecting passengers and crew to be disrupted. In either case, we end with a situation that may be sub-optimal and costly. In the next section we conclude the motivation for this problem with U.S. carrier data that exposes the significance of this problem in actual flight plans and proposes our solution approach.

2.3 Mitigation Approach

At the carrier with which we worked, on an average day at least 20% of all LOFs incur the type of swap illustrated in the previous examples. Such swaps result in aircraft ending their day-of-operations at random

stations in the network with various amount of flight-time left before maintenance must be performed. It is these significant network impacts that motivate our search for a schedule that is resistant (*robust*) to such unplanned changes.

In our approach, we aim to combat the effects of these unplanned, operational schedule changes, such that day-7 aircraft, starting the day at an unplanned station can still be assigned to an MLOF. This enables the aircraft to end its seventh day at a maintenance station without having to perform additional day-of-operation swaps. Instead, we examine the schedule from a planning perspective and attempt to “pre-plan” for recovery from operational, over-the-day swaps. To motivate the problem approach, we consider a numerical example next.

In this example, we feature two stations, BOS and ORD, each with 10 outgoing lines-of-flight. Out of the 10 lines-of-flight exiting station ORD, two end in a maintenance station, i.e. are MLOFs. On the other hand, of the 10 stations leaving BOS, zero are MLOFs. Furthermore, like (Sriram and Haghani 2009), we suppose that the probability of a day-7 aircraft appearing anywhere in the network is $1/7$. That is, of all aircraft at the start of each day, $1/7$ will require an A-Check at the end of that day. Derived from this assumption, we note that of the 10 flights that leave BOS, on average $1/7$ or ≈ 1.43 will require maintenance. However, these aircraft will not have the opportunity to be routed to a maintenance station directly due to the fact that none of the LOFs leaving BOS actually end in a maintenance station. This forces manual intervention (precisely the scenario we aim to avoid) over the course of the day to force a swap with another line, such that a maintenance opportunity is created.

On the other hand, applying a binomial approximation to the 10 LOFs out of ORD, of which 2 are MLOFs, we see that on average only 0.21 day-7 aircraft will be unable to reach a required maintenance station. [This calculation is also detailed in §5]. Alternatively, it would be beneficial to reallocate one of the maintenance opportunities out of ORD to BOS so as to decrease the overall expected number of maintenance misalignments. For this example, using the binomial approximation from §5, moving one of the two maintenance opportunities at ORD to BOS would reduce the expected number of maintenance misalignments from 1.64 ($0.21 + 1.43$) to 1.28 ($0.64 + 0.64$).

In our approach, we seek changes to the original plan that can improve maintenance reachability without substantially altering the initial lines-of-flight. We do so by identifying candidate LOFs for “splicing opportunities”. That is, when lines-of-flight, departing from differing stations, intersect at some point in space and time, these lines become candidates for a line-splice. In effect, at the splice, the route of the respective aircraft is changed such that each aircraft now flies the remaining route of the other. As a result, if one line-of-flight ends in maintenance, while the other does not, a re-distribution of maintenance opportunities has occurred.

For the example above, if one of the 10 LOFs that depart BOS and one the 10 LOFs leaving ORD intersect at some point in time and space, we can perform a line-splice. This *line-splice* involves changing the routing of a particular tail, as seen in the line-splice at TUS in Figure (2). In this example, such a line-splice results in two changes: the line leaving BOS gains a maintenance opportunity, while the line leaving ORD loses one. In other words, we have performed the desired re-balancing of the network without costly schedule changes.

In fact if we have m stations in the flight network with n flights ending the day in a maintenance station, and n is evenly divisible by m , then maintenance reachability is optimized by assigning exactly $\frac{n}{m}$ maintenance opportunities to each station. This optimal allocation is formally proven in Theorem (6.1) and the subsequent Corollary (6.2) in the appendix.

This example demonstrates an important property about the problem situation. A more *balanced* allocation of maintenance opportunities can produce an overall reduction in expected number of aircraft requiring additional, disruptive LOF swaps. In effect, each additional maintenance opportunity has a smaller probability of being needed and thus provides incrementally less value. Therefore, for two stations with the same

number of outgoing lines, the improvement of going from $n - 1$ to n at one station outweighs the loss of going from $n + 1$ to n at the other. Of course not all stations will have the same number of total outgoing LOFs, meaning that an equal allocation may not be optimal. The goal of our research is to determine the optimal allocation of maintenance opportunities in such cases.

This re-allocation is performed by identifying a set of line-splices, each of which re-distributes a maintenance opportunity from one station to another, such that the optimal balance is achieved.

3 Literature Review

Airline operations have benefited greatly from the advent of applied optimization techniques. Such optimization techniques have addressed various processes of the airline operations portfolio, including schedule planning, fleet assignment, crew rotation design and aircraft (tail) assignment. Relevant literature specifically related to our work falls into two inter-linked categories: maintenance route planning and maintenance recovery.

Maintenance route planning, a proactive optimization technique, strictly focuses on performing the tail assignment, such that required maintenance events are met from a planning perspective. On the other hand, maintenance recovery, a reactive optimization technique, emphasizes the need to re-evaluate maintenance plans on a daily basis, as unexpected changes can and do occur in an airline plan.

In the area of maintenance route planning, (Grönkvist 2004) provides a detailed overview of the tail assignment problem. In addition to various formulations, multiple solution approaches are also discussed. Additionally, it is noteworthy that the author discusses a set of objectives for the tail assignment problem. As noted in the introduction, the tail assignment is a feasibility problem, as opposed to an optimization problem. (Clarke, Hane, Johnson, and Nemhauser 1996) introduce new solution approaches to the maintenance routing problem. Specifically, these approaches draw analogies between the classic Traveling Salesman Problem (TSP) and maintenance routing (tail assignment) problem. Finally (Sriram and Haghani 2009) revisit the maintenance routing problem and provide a combination of random- and depth-first search heuristics to solve instances with large planning horizons.

In our work, we focus on maintenance routing from the planning perspective. More specifically, we address the process by which lines-of-flight are created and subsequently assigned to tails so as to route them to a maintenance station. (Gopalan and Talluri 1998) first use the concept of a line-of-flight (LOF) and detail this assignment process. The authors provide a polynomial-time algorithm to solve a three-day routing problem for a fixed set of lines-of-flight.

While operations research techniques have been applied to airline operations for a over two decades, it was only recently that optimization techniques were integrated to provide robustness and recovery decisions of aircraft operations. (Filar, Manyem, and White 2001) provide a survey of various models that provide recovery methods under uncertainty. When dealing with maintenance routing issues under uncertainty, (Teodorovic and Stojkovic 1995) provide a model that alerts the dispatcher of inconsistencies in a particular maintenance routing plan. It is then the responsibility of the dispatcher to find a line-splicing opportunity and route aircraft accordingly. In our approach, we aim to increase maintenance reachability, a robustness measure, such that under unexpected changes to the master schedule, aircraft are still able to reach required maintenance stations without changing the planned LOFs.

4 Optimization Model

We next present an optimization model to improve maintenance reachability. This approach takes an existing plan, designed to ensure maintenance feasibility, provide desirable through-flights, keep crew with aircraft, etc. We then identify limited changes that improve maintenance reachability without major deviations from this original plan.

4.1 MLOF Assignment Problem

We begin by finding a lower bound on the expected number of maintenance misalignments. We do so by optimizing the distribution of maintenance opportunities, without regard for whether this distribution can be achieved given the existing schedule.

Sets

S	Set of stations in the flight network.
T	Number of flights ending the day in a maintenance station.

Parameters

p_s^n	Expected number of maintenance misalignments at station s given n MLOFs, $\forall s \in S, \forall n \in \{1 \dots T\}$. [The derivation of this parameter is detailed in §5].
---------	---

Variables

x_s^n	Binary variable that is 1 if station s has n MLOFs assigned and 0 otherwise, $\forall s \in S, \forall n \in \{1 \dots T\}$. For example, $x_1^5 = 1$ refers to station 1 having 5 MLOFs assigned.
q_s	Integer variable that represents the number of MLOFs assigned to station s . For example, $q_1 = 5$ refers to station 1 having 5 MLOFs assigned.

We now formulate the lower-bound problem by making the decision to assign a particular *number* of maintenance opportunities to each station in the flight network.

$$\min \sum_{s \in S} \sum_{n=0}^T p_s^n x_s^n \quad (1)$$

subject to the following constraints:

$$\sum_{n=0}^T x_s^n = 1 \quad \forall s \in S \quad (2)$$

$$\sum_{n=0}^T n x_s^n = q_s \quad \forall s \in S \quad (3)$$

$$\sum_{s \in S} q_s = T \quad (4)$$

$$x_{sn} \in \{0, 1\} \quad \forall n \in \{1 \dots T\}, \forall s \in S \quad (5)$$

$$q_s \geq 0 \quad \forall s \in S \quad (6)$$

In the above formulation, the objective, seen in Equation (1), minimizes the expected number of maintenance misalignments across all stations in the flight network. Constraint (2) ensures that only one numeric value of maintenance opportunities is assigned to a particular station. Constraint (3) provides a link between the number of maintenance opportunities assigned to station s and the numeric equivalent q_s . Finally constraint (4) accounts for all maintenance opportunities available for assignment in the network.

4.2 Determining Splice Opportunities

The previous section provided a lower-bound formulation on the allocation of maintenance opportunities to stations, based on the total number of outgoing lines-of-flight. This is done without considering the feasibility of such an allocation. In this section, we identify an optimal allocation based on modifying the existing schedule.

We begin by finding line-splice opportunities. That is, given two lines exiting from differing stations, we determine if these lines can be spliced, such that at the point of interchange, the aircraft routing may be changed, if so desired by the optimization model. A re-assignment of aircraft at a line-splice opportunity effectively re-distributes the maintenance opportunity from the line that ends in a maintenance opportunity to the one that does not.

In this model, we restrict ourselves to simple line-splices. That is, at any point, only two lines are considered candidates for splicing. [In section 6, we discuss more complex line-splices].

- **Space** : The lines must intersect at a common station.
- **Time** : The lines must intersect at this station at the same time.
- **Aircraft** : The aircraft type must match (assuming crews are not interchangeable).
- **Benefit** : A splice is only worth considering if one of the two lines gains a maintenance opportunity while the other loses one. (Splicing two MLOFs or two non-MLOFs, provides no operational benefit).
- **Line-Locks** : Certain connections within a line may be “locked”, eg. they represent a high-value one-stop itinerary that should not be altered.




Figure 3: Overlap Diagram illustrating Swap Opportunities

It is the overlap in the above-mentioned factors, and as illustrated in Figure (3), that provide a *splice-opportunity* during the pre-planning process. Thus, our solution approach is now two-fold. First, we determine all splice-opportunities that exist in a schedule. This information then serves as input to the optimization model (presented next) to determine which splicing opportunities should be implemented to minimize the expected number of maintenance misalignments.

4.3 Restricted MLOF Assignment Problem

Given a set of candidate line-splice opportunities, derived from the criteria as introduced in the previous section, we now pose an optimization model that selects the optimal set of line-splices to perform so as to maximize maintenance reachability.

$$\min \sum_{s \in S} \sum_{n=0}^T p_s^n x_s^n \quad (7)$$

Sets

P	Set of all possible line-splices.
L	Set of all lines in the flight network.
S	Number of stations in the system.
T	Number of flights ending the day in a maintenance station.

Parameters

m_s	Original number of MLOFs associated with station s , $\forall s \in S$
p_s^n	Expected number of maintenance misalignments at station s given n MLOFs, $\forall s \in S, \forall n \in \{1 \dots T\}$. [The derivation of this parameter is detailed in §5].
δ_{sp}^+	Binary parameter that is 1 if splice p increases the number of maintenance opportunities at station s by 1 and 0 otherwise, $\forall s \in S, \forall p \in P$
δ_{sp}^-	Binary parameter that is 1 if splice p decreases the number of maintenance opportunities at station s by 1 and 0 otherwise, $\forall s \in S, \forall p \in P$
ω_{pl}	Binary parameter that is 1 if line l is part of splice p , and 0 otherwise, $\forall p \in P, \forall l \in L$

Variables

x_s^n	Binary variable that is 1 if station s has n MLOFs assigned and 0 otherwise, $\forall s \in S, \forall n \in \{1 \dots T\}$.
y_p	Binary variable that is 1 if swap p is performed and 0 otherwise, $\forall p \in P$

subject to the following constraints:

$$\sum_{n=0}^T x_s^n = 1 \quad \forall s \in S \quad (8)$$

$$\sum_{n=0}^T n x_s^n = m_s + \sum_{p \in P} \delta_{sp}^+ y_p - \sum_{p \in P} \delta_{sp}^- y_p \quad \forall s \in S \quad (9)$$

$$\sum_{p \in P} \omega_{pl} y_p \leq 1 \quad \forall l \in L \quad (10)$$

$$x_s^n \in \{0, 1\} \quad \forall n \in \{1 \dots T\}, \forall s \in S \quad (11)$$

$$y_p \in \{0, 1\} \quad \forall p \in P \quad (12)$$

As before, the objective as seen in Equation (7) is to minimize the total number of expected maintenance misalignments. Constraint (8) ensures that only one number of maintenance opportunities is assigned to each station. Constraint (9) performs the numerical assignment based on the splice choices (y_p) made. That is, for every splice that adds an MLOF to station s , we increase the total count by 1. On the contrary, when a splice is made that removes an MLOF from station s , we decrease the total number of MLOFs for station s by 1. Finally, constraint (10) requires that all lines are spliced at most once.

4.4 Minimization of Schedule Changes

There may be multiple equivalent solutions to the mathematical model presented in §4.3. For example, suppose we have two stations, BOS and SEA, both with two lines-of-flight. For each station, one of the LOFs ends in a maintenance station while the other does not. Further suppose that both LOFs from BOS can splice with both LOFs from SEA. The optimization approach presented thus far could splice the MLOF from BOS with the LOF from SEA and the LOF from BOS with the MLOF from SEA. The net-effect of these splices is zero with respect to maintenance reachability, unnecessarily changing the original planned LOFs.

To minimize the impact in terms of the number of changes to the schedule, we formulate a second-stage optimization model. This optimization model minimizes the total number of swaps required to achieve the objective function value as derived during the first-stage optimization problem presented in §4.3.

Parameters

C Optimal objective value from the first-stage optimization.

$$\min \sum_{p \in P} y_p \quad (13)$$

subject to the following constraints:

$$\sum_{n=0}^T x_s^n = 1 \quad \forall s \in S \quad (14)$$

$$\sum_{s=1}^S \sum_{n=0}^T p_s^n x_s^n = C + \epsilon \quad (15)$$

$$\sum_{n=0}^T n x_s^n = m_s + \sum_{p \in P} \delta_{sp}^+ y_p - \sum_{p \in P} \delta_{sp}^- y_p \quad \forall s \in S \quad (16)$$

$$x_s^n \in \{0, 1\} \quad \forall n \in \{1 \dots T\}, \forall s \in S \quad (17)$$

$$y_p \in \{0, 1\} \quad \forall p \in P \quad (18)$$

In the above formulation, objective (13) seeks to minimize the total number of swaps that are made. We are re-using the notation from the previous formulation, such that y_p is a $\{0, 1\}$ variable allowing for a summation of the total number of splices performed. Constraints (14) ensure that we still only assign a single number of maintenance opportunities to each station. Constraint (15) requires that our new assignment must yield the same objective value (C) as the value obtained during the previous optimization step. Note, we add a small factor to account for precision errors. Finally, constraints (16) perform the required accounting of maintenance stations, as in the first-stage optimization model.

5 Computational Results

To validate the optimization model introduced in §4, we provide results from computational experiments that: demonstrate the tractability of our approach, provide a detailed summary of the quality of the solution and its impacts, and assess the sensitivity of the model to the probabilistic objective coefficients.

To do so, we first introduce an approximation of the expected number of day-7 aircraft at a station requiring maintenance at the end of the day-of-operations. We then refine this approximation using carrier-based data, showing the impact of more precise coefficients on the overall solution quality. Finally, we conduct a sensitivity analysis to determine how much impact errors in the objective coefficients have on the quality of the solution that is found.

The input data for our model was sourced from a major U.S. carrier. The carrier provided several daily tail assignments from which LOFs were built with maintenance and non-maintenance stations indicated. Additional details regarding the input plan can be found in Table (3).

Number of Flights:	3353
Number of LOFs:	512
Number of Stations:	64
Number of MLOFs:	206

Table 3: Characteristics of Input Data for Optimization Model

5.1 p_s^n Distributed as Binomial with $p = \frac{1}{7}$

As mentioned in §4, the optimization model’s objective function requires the expected number of maintenance misalignments. That is, coefficient p_s^n in Equation (7) is the *expected* number of day-7 aircraft at station (s) unable to reach a maintenance station at the end of the day-of-operations without over the day swaps, given that station s has n MLOFs. In the first part of this analysis, this coefficient is derived using an approximation based on an even maintenance rotation of all aircraft through the network.

An aircraft that ferries passengers on a daily basis is required to receive a routine “A-Check” based on the block or flying time (Sriram and Haghani 2009). For the major U.S. carrier considered in this analysis, the maintenance requirement can be translated into an aircraft requiring an A-Check every seven days. Based on this type of maintenance schedule and the assumption that aircraft routing does not heavily influence the terminating station of any particular aircraft, we begin by assuming that any aircraft in the fleet on any given day has a $1/7$ probability ($p = 1/7$) to being a day-7 aircraft, i.e. an aircraft that requires maintenance at the end of the day. Based on this probability, we turn to the binomial distribution to derive the expected number of maintenance misalignments at a given station using Equation (19).

$$p_s^n = \sum_{i=n+1}^{L_s} \binom{L_s}{i} (p)^i (1-p)^{L_s-i} (i-n) \quad (19)$$

In Equation (19), the parameter L_s indicates the total number of LOFs exiting from a given station and (n) represents the number of MLOFs, i.e. how many LOFs provide access to a maintenance facility. The first piece of this equation provides the probability of day-7 aircraft requiring maintenance, which is then multiplied by the number of day-7 aircraft without a maintenance opportunity for each discrete case, resulting in the expected value calculation. It should be noted that we sum the expected day-7 aircraft starting at $(n+1)$, where n represents the number of maintenance opportunities associated with a particular station, because all aircraft up to n are able to reach a maintenance station as required.

5.1.1 Experimental Results

We begin by first solving the lower-bound formulation presented in §4.1, followed by the schedule modification model presented in §4.3. Using the seven-day maintenance assumption with p_s^n approximated using the binomial distribution presented in Equation (19) with $p = 1/7$, we obtain the following results in terms of the expected number of maintenance misalignments.


Schedule without Optimization:	7.33 misalignments
Lower-bound Optimization:	0.84 misalignments (-88%)
Optimized Schedule (with $p = 1/7$)	0.84 misalignments (-88%)
Runtime:	< 5 seconds*

*Computation performed on an Intel Duo-Core 2.2GHz processor with 4GB of memory.

The optimization model is able to reduce the expected number of maintenance misalignments by 88%. This means that when the newly formed LOFs are operated, on average, our plan will incur 0.84 day-7 aircraft (as opposed to 7.33) that will require manual re-routing to end their day in a maintenance station. Note that the restricted optimization model is able to achieve the same objective value as the lower-bound, indicating that the lowest number of network-wide maintenance misalignments has been reached. Further detail as to the implications of these results on the actual aircraft plan and the airline are detailed in §6.


Figure (4(a)) and Figure (4(b)) illustrate the allocation of maintenance opportunities in the original and optimized sets of LOFs. Overall, we notice a general re-balancing of maintenance opportunities, with much lower peaks but nonetheless higher numbers of maintenance opportunities at stations with larger numbers of LOFs. Evident from these figures, the variability in terms of the number of maintenance opportunities is lower than pre-optimization. Correspondingly, Figure (5) shows an overall decrease in the expected number of maintenance misalignments at every station. In this figure, we provide the original expected number of maintenance misalignments (dashed-line), as well as the expected number of maintenance misalignments under the new LOF configurations (solid-line).

Figure 4: Maintenance opportunity assignment using ($p = 1/7$) approximation



(a) Maintenance opportunity assignment - Original Plan (b) Maintenance opportunity assignment - Optimized Plan

Figure 5: Expected number of maintenance misalignments using $p = 1/7$



5.2 Probability Scenarios

Having first considered the simple case where all aircraft at all stations are equally likely to be day-seven aircraft (with a probability of $1/7$), we next conduct experiments using empirical data from a major U.S. carrier to provide insights into whether using more finely-tuned values for the probability of a day-7 aircraft will change the results of the optimization. To evaluate the effectiveness of our optimization model under this data set, we evaluate five probability vectors, as provided in Table (4). Each vector provides empirical estimates, from different periods of time, of the probability of a day-7 aircraft appearing at a maintenance station (m), a large station (l) and a small (s) station. A “large” station differentiates itself from a “small” station by the number of LOFs that depart each day. In this case, a large station is defined as one which has at least seven outgoing LOF. Probability vector P_0 denotes the case where all station types have probability $p = 1/7$ (i.e. 0.1429) as assumed in §5.1.

Using the parameters shown in Table (4), we solve six different instances of the problem, i.e. we generate six optimal sets of LOFs, one for each probability vector. We then evaluate how each of these sets of LOFs performs under all of the different probability vectors. In other words, we ask the question: If we assume one probability vector to represent the true parameters and optimize accordingly, but another probability

	MX	Large	Small
P_0	0.1429	0.1429	0.1429
P_1	0.1162	0.1926	0.1689
P_2	0.0462	0.1824	0.1181
P_3	0.0829	0.1585	0.1048
P_4	0.1632	0.1722	0.1571
P_5	0.0653	0.1216	0.0685

Table 4: Probability vectors based on Actual U.S. Carrier Data

vector is in fact the true data, how will our “optimized” LOFs perform? This question is answered in Table (5). Each row corresponds to a given set of LOFs — the original set of LOFs and the results of optimizing over each of the six probability vectors, as shown given in the left-most column. The remainder of each row then gives the objective value for that set of LOFs as evaluated under each of the six probability vectors. For example, if probability vector P_1 is used to optimize the LOFs but then P_2 is realized, the expected number of maintenance misalignments will be 0.45.

		Probability Vectors					
		P_0	P_1	P_2	P_3	P_4	P_5
LOFs	Original	7.33	8.98	5.64	5.01	8.69	3.02
	P_0	0.84	1.51	0.95	0.62	1.39	0.24
	P_1	1.14	1.10	0.45	0.35	1.89	0.14
	P_2	3.17	1.93	0.31	0.56	4.90	0.23
	P_3	1.33	1.15	0.42	0.34	2.20	0.13
	P_4	0.87	1.45	0.81	0.54	1.35	0.21
	P_5	1.35	1.16	0.42	0.34	2.25	0.13

Table 5: Objective Value Comparison for Scenario Data

In Table (5), the optimal values are highlighted for each set of LOFs and its corresponding probability vector. Of course, by definition, this value is the lowest value in each column. These values illustrate the impact of the optimization model under perfect information with each optimized solution offering a significant improvement in terms of maintenance reachability. In addition, the remainder of the columns show that even without perfect information, the optimization model is able to reduce the total number of maintenance misalignments.

	Probability Vectors					
	P_0	P_1	P_2	P_3	P_4	P_5
Original	7.33	8.98	5.64	5.01	8.69	3.02
Worst-Case Scenario	3.17	1.93	0.95	0.62	4.9	0.24
Best-Case Scenario	0.84	1.10	0.31	0.34	1.35	0.13
Max Improvement	6.49	7.88	5.33	4.67	7.34	2.89
	88.5%	87.8%	94.5%	93.2%	84.5%	95.7%

Table 6: Objective Value (Maintenance Misalignments) Improvement Results

5.3 Sensitivity of Expectation Coefficients

In the previous section we considered a variety of possible values for the probability vector P . We showed that even if the incorrect value of P was used to optimize the lines-of-flight, the resulting schedule (when evaluated under the true value of P) yields higher maintenance reachability than the original schedule. In this section, we further explore the relationship between P , the “optimized” LOFs, and the realized

maintenance reachability. We show that in most cases, even with large deviations from the actual probability, the “optimized” LOFs still gain relative to the original LOFs with respect to maintenance reachability.

As introduced in the previous section, we divide stations into three different types: maintenance (m), large (l) and small (s). We use a probability vector (p_m, p_l, p_s) that indicates the probability that an aircraft at each of the respective station types is a day-7 aircraft. In this analysis, we vary the probability of a day-7 aircraft from 0.04 up to 0.2 incrementing by 0.04 for each of the station types, resulting in a total of 125 scenarios. We optimize with respect to each of these probability vectors and then evaluate each of these 125 optimized plans relative to the other 124 probability vectors. We quantify the results from this analysis next.

To do so, we first define the notion of *perfect information*. By this, we consider the case where the probability vector under which the LOFs were optimized is in fact the realized probability vector. Clearly, this is the best case scenario, whereas the worst case scenario is when the LOFs are optimized relative to a probability vector that is very different from the realized probability vector. For the case of perfect information, in all 125 cases, the optimized LOFs performed *strictly* better than the original schedule. Table (7) shows the average improvement across all 125 probability vectors.

Average (across scenarios) of Original Plan:	6.27 misalignments
Average (across scenarios) of Optimized Plan:	0.57 misalignments
Average Improvement:	5.70 misalignments (-90%)

Table 7: Maintenance misalignments improvement across all probability vectors

As alluded to before, our optimization model also performs well in situations where the LOFs are optimized under imperfect information, i.e. relative to the wrong probability vector. Consider an arbitrary probability vector. Suppose that we optimize relative to this probability vector and then evaluate this set of LOFs under all other possible vectors. If we take the maximum of each of these objective function values *and* this maximum is lower than the original plan, then it must be the case that even if the probabilities of a day-7 aircraft are unknown, the optimization still outperforms the original plan. In our analysis, there are 71 out of 125 cases where the optimized LOFs out-perform the original LOFs, *independent of what probability vector is actually realized*.

In the remaining 54 cases the optimized LOFs are not guaranteed to perform as well as the original schedule in the worst case scenario, but only when the optimized plan is based on a probability vector that deviates largely from the actual realized probability vector used for execution. Even in these 54 cases, the optimized LOFs out-perform the original LOFs in most cases —on average, 118 out of 125 realized vectors yield higher maintenance reachability with the optimized rather than the original LOFs. That is, on average, in only 7 of the 125 planned LOFs, in only 51 out of 125 scenarios, does the original set of LOFs achieve higher maintenance reachability.

6 Impact & Conclusions

Aircraft routings are designed to meet strict FAA mandates for aircraft maintenance intervals. As emphasized in this paper, on any given day, these plans are routinely changed to mitigate unforeseen events, such as delay propagation, which disrupt airline schedules. To evaluate these day-of-operation changes and their impact on maintenance routing, we provide a metric, known as *maintenance reachability*. This metric captures a schedule’s ability to satisfy maintenance constraints. More specifically, for each station in the flight network, maintenance reachability quantifies the ability of an aircraft to reach a maintenance opportunity on its last day-of-operation.

When maximizing maintenance reachability, the goal is to minimize the expected number of aircraft that are *unable* to reach a maintenance station on their seventh day of operation using existing lines-of-flight. As

we have shown, this is done by correlating the number of maintenance opportunities with the total number of lines-of-flight for each station. In special cases, where each station has the same number of LOFs, the optimal solution is one where each station has the same number of maintenance opportunities. However, since airlines rarely have the same number of LOFs leaving each station, an optimization model is required to perform the assignment of maintenance opportunities to stations, so as to minimize the expected number of day-7 aircraft without a maintenance opportunity.

The optimization model presented in this paper provides the necessary network balance of maintenance opportunities to maximize maintenance reachability. That is, we strategically splice lines-of-flight to change the allocation of maintenance opportunities from stations with less need to stations with greater need. Using a two-stage linear-programming formulation, we are able to show significant improvements with respect to maintenance reachability by making no-cost changes to the flight plan. As a result, the new schedule is more robust to unexpected changes that may occur on the day-of-operation, reducing the need for costly interventions of manual re-routing aircraft to ensure maintenance feasibility. While our mathematical model depends on probability parameter estimates, we show through a sensitivity analysis that those parameter estimates can often deviate from the actual probability while still producing an improved schedule.

6.1 Future Work

To develop this model further, several avenues of extension exist.

6.1.1 Maintenance Capacity Constraints

While the optimization model presented in this paper does assign maintenance opportunities to stations, it does so without constraints on the actual maintenance station. Since maintenance stations have a finite capacity of the number of aircraft that can be serviced on any given evening, a further extension of this model could incorporate a maintenance balancing decision when making maintenance opportunity assignments. More specifically, the maintenance opportunities and their respective maintenance station should be balanced across each of the stations. Such an assignment would allow for a more equal work-balance in terms of the number of aircraft that will actually require maintenance at the end of a day-of-operations. Alternatively, capacity constraints could be assigned to each maintenance station, restricting the number of maintenance tasks at each facility in accordance with its capabilities.

6.1.2 Complex Line-Splicing

As mentioned in §1, this model was developed with input from a major U.S. carrier. Within this carrier's network, we were able to achieve the lower bound on maintenance reachability by employing single line splices. This may not be the case for other carriers' networks. For example, one can imagine that if a required maintenance opportunity cannot be obtained through single line splices, another level of splicing could provide additional maintenance access. That is, line *A* is spliced with line *B* for no gain in maintenance reachability. However, further down the line-of-flight, line *B* is spliced with line *C*, creating a maintenance opportunity for the aircraft flying *A*. Future research would investigate alternative levels of splicing, and the corresponding modification of the optimization model, to achieve higher levels of maintenance reachability.

6.1.3 Through-Flight Connection Cost

Finally, in our model, we assume that line-splices come at virtually no cost. In fact, there may be a revenue impact associated with breaking an existing through-flight connection, and/or an operating cost associated with changing crew/aircraft pairings. Conversely, new LOFs may create new revenue opportunities (by introducing new through itineraries) and new opportunities for operational savings. Thus, it may be worth extending the existing model to incorporate the costs and benefits of LOF modification.

References

- Clarke, L., C. Hane, E. Johnson, and G. Nemhauser. 1996. Maintenance and crew considerations in fleet assignment. *Transportation Science* 30 (3): 249–260.
- Filar, J., P. Manyem, and K. White. 2001. How airlines and airports recover from schedule perturbations: A survey. *Annals of Operations Research* 108:315–333.
- Gopalan, R., and K. Talluri. 1998, March. The aircraft’s maintenance routing problem. *Operations Research* 46 (2): 260–271.
- Grönkvist, M. 2004. *Integration of ai and or techniques in constraint programming for combinatorial optimization problems*, Chapter A Constraint Programming Model for Tail Assignment, 142–156. Springer Berlin / Heidelberg.
- Sriram, C., and A. Haghani. 2009. An optimization model for aircraft maintenance scheduling and re-assignment. *Transportation Research Part A* 37:29–48.
- Teodorovic, D., and G. Stojkovic. 1995. Model to reduce airline schedule disturbances. *Journal of Transportation Engineering* 121 (4): 324–331.

Appendix

Discrete Convexity Proof

Theorem 6.1. *The day-6 aircraft expectation coefficients used in the objective function are discrete-convex when the number of LOFs are equal for each station.*

Proof. To show convexity of the objective function coefficients, a discrete function, we must show that:

$$f(n-1) + f(n+1) \geq 2f(n) \quad (20)$$

We begin with the probability coefficient as seen in equation 21.

$$f(n) = \sum_{i=n+1}^{L_s} \binom{L_s}{i} \left(\frac{1}{7}\right)^i \left(\frac{6}{7}\right)^{L_s-i} (i-n) \quad (21)$$

Replacing the left-hand side of equation 20, we obtain the following expression:

$$\begin{aligned} & \binom{L_s}{i-1} \cdots \times (1) + \binom{L_s}{i} \cdots \times (2) + \binom{L_s}{i+1} \cdots \times (3) + \cdots + \binom{L_s}{L_s} \cdots \times (L_s - n - 1) \\ & \binom{L_s}{i+1} \cdots \times (1) + \cdots + \binom{L_s}{L_s} \cdots \times (L_s - n + 1) \end{aligned}$$

Combining terms from the equation above yields equation 22:

$$\binom{L_s}{i-1} \cdots \times (1) + \binom{L_s}{i} \cdots \times (2) + \binom{L_s}{i+1} \cdots \times (3+1) + \cdots + \binom{L_s}{L_s} \cdots \times ((L_s - n + 1) + (L_s - n - 1)) \quad (22)$$

Now, replacing the right-hand side of equation 20, we obtain the following expression:

$$2 \left[\binom{L_s}{i} \cdots \times (1) + \binom{L_s}{i+1} \cdots \times (2) + \cdots + \binom{L_s}{L_s} \cdots \times (L_s - n) \right] \quad (23)$$

Combining equations 22 and 23, eliminating common terms and re-writing in the form of equation 20, we obtain the following expression

$$\binom{L_s}{i-1} \left(\frac{1}{7}\right)^{i-1} \left(\frac{6}{7}\right)^{L_s-i-1} \times (1) \geq 0 \quad (24)$$

which is trivially true. From this, we have shown that our probability coefficient function is indeed discrete-convex. $\square$

Equal MLOF Assignment Proof

Corollary 6.2. *Given a situation where n MLOFs are to be assigned to m stations with the fact that n is evenly divisible by m , then the optimal allocation is one where an equal number of MLOFs is assigned to each station.*

Proof. We will prove this corollary by contradiction. Given an assignment whereby all MLOFs outlets are equally distributed, then changing this assignment will result in a *lower* number of day-6 aircraft that will be unable to reach maintenance. In other words, changing a station by increasing its number of MLOFs by one and therefore decreasing another station by one will result in a lower number of expected day-6 aircraft requiring maintenance.

In mathematical terms:

$$f(n+1) + f(n-1) \leq f(n) + f(n) \quad (25)$$

This result is clearly false as it contradicts the convexity argument in the previous proof. Therefore, it must be the case that the optimal allocation is one where an equal number of MLOFs are assigned to each station. $\square$