

**INDUSTRY STUDIES ASSOCIATION
WORKING PAPER SERIES**

Decreasing Airline Delay Propagation by Re-Allocating Scheduled Slack

By

Amy Cohn
Global Airline Industry Program
Massachusetts Institute of Technology
Cambridge, MA 02139
and
University of Michigan, Ann Arbor
Ann Arbor, MI 48109
amycohn@umich.edu

and

Shervin AhmadBeygi
University of Michigan, Ann Arbor
Ann Arbor, MI 48109
shervin@umich.edu

and

Marcial Lapp
University of Michigan, Ann Arbor
Ann Arbor, MI 48109
mlapp@umich.edu

2008

Industry Studies Association
Working Papers

WP-2008-14
<http://isapapers.pitt.edu/>

Decreasing Airline Delay Propagation By Re-Allocating Scheduled Slack

Shervin AhmadBeygi, Amy Cohn and Marcial Lapp
University of Michigan

April 4, 2008

Abstract

Passenger airline delays have received increasing attention over the past several years as airspace congestion, severe weather, mechanical problems, and other sources cause substantial disruptions to a planned flight schedule. Adding to this challenge is the fact that each flight delay can propagate to disrupt subsequent downstream flights that await the delayed flights' aircraft and crew. This potential for delays to propagate is exacerbated by a fundamental conflict: slack in the planned schedule is often viewed as undesirable, as it implies missed opportunities to utilize costly perishable resources, whereas slack is critical in operations as a means for absorbing disruption. In this paper, we show how delay propagation can be reduced by redistributing existing slack in the planning process, making minor modifications to the flight schedule while leaving the original fleet and crew scheduling decisions unchanged. We present computational results based on data from a major U.S. carrier, showing that significant improvements in operational performance can be achieved without increasing planned costs.

1 Introduction

Airline plans are made up of several costly and constrained resources such as aircraft and crews. These resources link flights across the network, with each resource flowing from one flight to another. Adding to this complexity is the fact that although each flight needs each type of resource, individual resources do not necessarily stay linked throughout the network. For example, an aircraft and crew might be assigned to a common flight at a particular point in the schedule, but assigned to separate flights at a later point.

One ramification of this linkage is the potential for delays to propagate. If one flight is delayed (for example, because of a mechanical problem with the aircraft assigned to that flight), then a subsequent flight might also be delayed because it is awaiting that inbound aircraft. The fact that resources can “split” compounds this. In Figure 1 (explained in detail in [3]) we see how a single flight delay can spread to delay several other flights as well.

Airline delays have increased substantially in the past 5 years (see Figure 2). The cost impact of these delays is substantial, including excess fuel costs (from idling aircraft), overtime pay for crew members, costs associated with re-accommodating misconnecting passengers, as well as the lost productivity of delayed passengers. Furthermore, the Air Transport Association has estimated that there were a total of 116.5 million delay minutes in 2006, resulting in a \$7.7 billion increase in direct operating costs to the U.S. airline industry (see Table 1).

There are many sources for flight delays, such as mechanical problems, weather delays, ground-hold programs, and air traffic congestion. But the secondary delays that propagate from such *root delays* are also quite substantial. For example, in November 2007, more than one-third of the delays at major U.S. airports were the result of a late-arriving aircraft (Figure 3). Furthermore, there is a natural conflict stemming from the fact that slack is typically viewed as negative from the planning perspective (i.e. a waste of resources), but as positive from the operational perspective (i.e. an opportunity to absorb disruption rather than allowing

it to propagate). The focus of our research is therefore on determining how to incorporate the operational issues associated with delay propagation into the airline planning process.




Figure 1: An example of a delay propagating. In this example, a 180-minute delay in Flight 1 results in 430 minutes of delay in downstream flights.

Direct Operating Costs (2006)	\$ Per Minute	Annual Delay Costs (\$ millions)
Fuel	28.31	3,296
Crew	14.25	1,659
Maintenance	10.97	1,277
Aircraft Ownership	9.18	1,069
Other	3.1	361
Total	65.80	7,663

Table 1: Direct costs of delays in the U.S. airline industry. Costs based on data reported by U.S. passenger and cargo airlines with annual revenues of at least \$100 million. Source: Air Transport Association

A key challenge in this research is the difficulty in trading off between *planned costs* (the cost of an airline plan under the assumption that all flights occur as scheduled and without disruption) and *operational costs* (the realized cost associated with the modified plan that is implemented in response to disruptions). Given two different plans with varying planned costs, it is difficult to determine which of the two plans will perform better operationally. Furthermore, it is also difficult to determine whether improvements in operational performance outweigh increases in planned costs, given that the plan will be operated several times (often, daily) over the planning horizon, and that potential disruptions may or may not occur during any given day.




Figure 2: The increasing trend in delayed flights. Source: Bureau of Transportation Statistics.

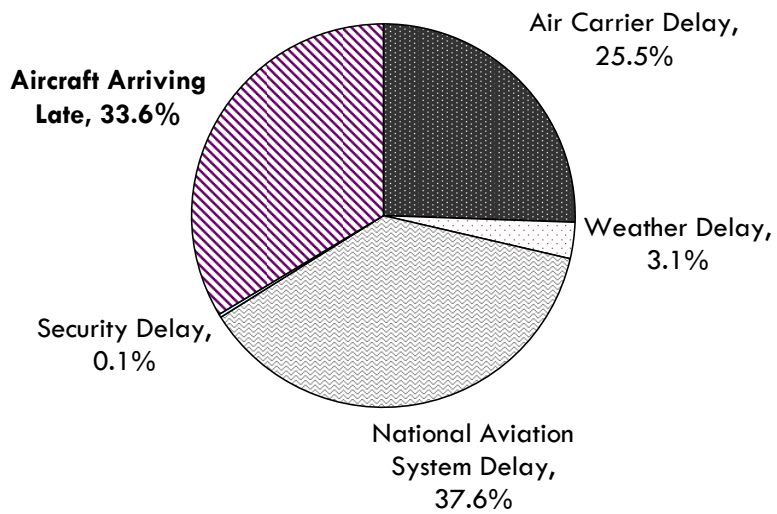


Figure 3: Delay causes among all major US airports in November 2007. Source: Bureau of Transportation Statistics.

Thus, even determining metrics for “robustness”, and then quantifying the value of these metrics (i.e. how much planned cost a carrier should be willing to incur to improve these metrics), are challenging research topics themselves that have yet to be adequately solved. The fact that the planning and operations processes are functionally separate within most carriers, with each group’s incentives aligned with different objectives, only serves to exacerbate the problem.

We therefore propose, as an interim step, to develop an approach that does not increase planned costs, but can nonetheless improve operational performance. Specifically, we propose to modify flight departure times so as to re-allocate the existing slack in the network. By re-timing flights, slack can be re-distributed to those flight connections that are most sensitive to disruption and thus delay propagation. We limit the time windows in which flights can be re-timed, so as to maintain existing revenue projections. Furthermore, we restrict flight re-timings such that crew pairings remain feasible and do not change in cost. Finally, we require that the same aircraft rotations be maintained. Our computational results, based on data from a major U.S. carrier, demonstrate that this approach leads to significant improvements in expected delay propagation without any associated increase in planned cost.

The primary contribution of our research is in developing models that can diminish delay propagation in operations, without any increase in planned costs. Our proposed models can take into account delay propagations caused by aircraft, crew members, connecting passengers, as well as other shared resources. By demonstrating that the integrality of these models can be relaxed (i.e. that the models can be solved as linear, rather than integer, programs), we are able to consider all down-stream impacts without sacrificing tractability. The paper is outlined as follows: In Section 2 we review the related literature. We present models for re-allocating slack in Section 3, as well as a simulation model to help validate the results. Computational experiments are presented and analyzed in Section 4. Finally, Section 5 offers conclusions and suggested areas for future research.

2 Literature Survey

In this section, we first briefly introduce the different steps in the airline planning process. Then we focus on different approaches to address airline delays and present related literature.

2.1 Airline Planning Problems

The airline planning process is classically decomposed into four sub-problems: schedule generation, fleet assignment, maintenance routing, and crew scheduling.

The objective of the *schedule generation* problem is to determine what markets, frequencies, and times to fly in a given period of time, taking into account both forecasted demand data and available resources (Berge [6]). The *fleet assignment* problem determines which type of aircraft should be assigned to each flight, considering the demand and capacity constraints (Abara [1]). The *maintenance routing* problem is primarily a feasibility problem that assigns specific aircraft to flights to ensure adequate opportunities for required maintenance checks (Barnhart and Talluri [5]). The objective of the *crew scheduling* problem is to find the most cost-effective assignment of cockpit or cabin crews to flights (Barnhart et al [4]).

These airline planning problems are complex and large-scale by nature. Therefore, they are often treated as deterministic in order to achieve tractability, not taking into account the impact of delays and disruptions. In the following subsection, we highlight some of the approaches taken to incorporate disruptions in the airline planning process.

2.2 Robust Planning

Airline operations are subject to significant uncertainties. Disruptions often occur as a result of weather conditions, unplanned maintenance issues, safety checks, security concerns, and more. The goal of *robust planning* is to generate schedules that are less sensitive to these disruptions. This is a demanding task, given

both the size and the uncertainty of the networks. There have been two major types of approaches to robust airline planning in the literature.

In the first approach, a stochastic programming model is used to capture uncertainties. Yen and Birge [26] developed a stochastic model for the airline crew scheduling problem and adopted a delay branching algorithm to solve the resulting problem. Alternatively, Rosenberger et al [21] used a simulation model to implement a stochastic approach for the crew scheduling problem and to evaluate recovery policies. Fuhr [10] developed a stochastic model to evaluate the on-time performance of a given schedule.

The second approach to modeling disruptions is to define surrogate problems that inherit the stochastic nature of the original problem. For example, Rosenberger et al [20] showed that fleet assignments with less hub connectivity and more short cycles perform better in operational circumstances. Schaefer et al [22] modeled the crew scheduling problem under uncertainty using approximated expected pairing costs. They also defined a lower bound on the expected cost of the pairings. Klabjan et al [11] defined a *regularity* measure as a way to capture robustness and considered maximizing this measure in addition to minimizing total cost.

2.3 Recovery Models

While robust planning models try to avoid disruptions proactively, *recovery models* seek the best way of reacting when disruptions do occur, so as to minimize their impact on the system and prevent propagation. Therefore, recovery models are often studied within the robust planning context [21]. Clarke et al [9] explained the role of an airline’s operations control center in mitigating the impact of irregularities on operations. Yan and Yang [25] developed a framework to handle schedule perturbations caused by aircraft breakdowns. Abdelghany et al [2] developed a decision support tool to automate crew recovery during irregular operations. Lettovsky et al [15] developed a real-time recovery plan to restore a disrupted crew schedule. Yu and Qi [27] have studied the recovery models used by United Airlines in the context of disruption management. More extensive surveys of the literature on recovery models can be found in Kohl et al [12] and Bratu and Barnhart [7].

2.4 Flight Re-timing Models

Perhaps the most closely related research to this paper is that of Stojkovic et al [23] and Lan et al [13], both of whom also consider the use of flight re-timings to improve schedule performance.

In [23], the primary focus is on day-of-operations recovery activities. In particular, they focus on how to modify an existing plan in order to recover from a set of minor disruptions. They require that crew connections, rest requirements, aircraft connections, maintenance requirements, passenger connections, etc. all be maintained. They permit not only changes to flight departure times (specifically, increases), but also allow activities to be expedited. For example, the amount of time required to off- and on-load passengers might be relaxed. The objective function then seeks to minimize the costs associated with extra resource utilization (for example, as needed when expediting activities) and passenger inconveniences.

Lan et al [13] consider two problems, both related to our research. First, they consider how changes in aircraft routings can be used to reduce the potential for delays to propagate via connecting aircraft. In this case, they hold flight departure times constant but allow the assignment of aircraft (i.e. tail numbers) to flights to change so as to better utilize the slack in the system to absorb disruption. In a separate problem, they keep the aircraft routings fixed but allow flight times to vary within a limited time window. The objective in this problem is to decrease the impact of delay on passengers’ ability to make flight connections.

Both of these papers serve to demonstrate how even minor schedule modifications can have significant impact on system performance under disruptions, and help motivate our research.

The idea of using time windows in the airline planning context was first introduced by Levin [16]. Rexing et al [19] allowed scheduled flight departure times to vary within a given time window to improve

flight connection opportunities and the cost-effectiveness of the fleet assignment. Stojkovic and Soumis [24] considered the problem of simultaneously modifying the existing flight departure schedule and planning individual work days (duties) while keeping planned aircraft itineraries unchanged. Mercier and Soumis [17] considered an integrated aircraft routing and crew scheduling model which allows the departure time of the flights to vary within a given time window. Burke et al [8] constructed a multi-objective optimization model to improve schedule reliability as well as schedule feasibility by re-timing flights and permitting minor changes in aircraft rotations while keeping the fleet assignments fixed.

We seek to extend this research by developing a flight re-timing model that focuses on minimizing the propagation of root flight delays. Our approach permits the simultaneous consideration of interactions between multiple resources (aircraft, crews, passengers, etc.) and follows the downstream propagation of delays until absorbed. In particular, we are able to do so in a linear (rather than integer) program, which has significant benefits in terms of computational performance. Because our model relies on a surrogate objective function to approximate delay propagation, we also develop a simulation-based approach to mimic the propagation of delays operationally. This simulation assists us in assessing the quality of our surrogate objective function.

3 Models for Re-Allocating Planned Slack

3.1 Main Idea

Our goal is to improve the expected *operational performance* of a planned airline schedule without increasing its *planned cost*. In particular, we want to re-time flight departures so as to re-distribute existing *slack* (the scheduled connection time between two flights sharing a common resource minus the minimum turn time between these flights) in the network, to make this slack available where it is most needed operationally. Note that by moving a flight’s departure earlier, we increase the slack in its outbound connections, but decrease the slack in its inbound connections (see Figure 4). Therefore, given a fixed amount of slack in the schedule, we want to re-distribute this slack to where it can best be utilized, taking into account both the current connection times (e.g. adding slack to a long connection is unlikely to provide significant benefit) and also the likelihood of root delays, which determine how frequently the slack will be needed.

We limit the extent to which each flight can be re-timed in three ways. First, we maintain the feasibility of the existing crew and aircraft assignments. Second, we protect connecting passenger itineraries. Third, we do not permit flight times to be modified so substantially that the projected demand levels would change. Although these restrictions limit the extent to which the schedule can be improved, we are nonetheless able to provide an interim level of improvement (which we show to be non-trivial) that can be achieved immediately, while researchers continue to study the broader challenges of how to quantify, value, and increase schedule robustness through more drastic changes.

The decisions in our model are how much earlier or how much later to re-time each flight’s departure. The constraints needed to enforce the limitations on how flights can be re-timed (so as to meet the three restrictions outlined previously) are fairly straightforward, as is seen in the following sections. The challenge lies in determining how to represent the objective function, so as to maintain tractability while providing a solution that does in fact reduce delay propagation.

To explain our objective function (which is a surrogate for “robustness”), first consider a single flight f . For this flight, we also have a set of relevant connections. For example, the aircraft assigned to flight f might next connect to flight g , while the cockpit crew connects to flight h . In addition, the cabin crew might connect to flight i , and flights j and k may represent key passenger connections. Each of these connections has some (non-negative but possibly zero) scheduled slack. In our constraints, we enforce that the change in departure time for flight f does not violate the minimum turn times for these connections – in other words, this scheduled slack must remain non-negative.

But what happens if flight f is delayed in operations? For example, suppose its assigned aircraft has a mechanical problem that takes 35 minutes to repair. In the absence of a recovery intervention, any

available slack will be used to absorb this delay, with the *residual delay* (the delay beyond the available slack) propagating to the connecting flights. In our example, if the slack between flights f and g is 40 minutes, then flight g will not be delayed, but if the slack between flights f and h is only 20 minutes, then h will experience a 15 minute departure delay. If we move the departure time of flight f to be 15 minutes earlier, then a 35 minute root delay to flight f will no longer propagate to impact flight h .

As a surrogate objective function, we propose to consider the sum of the impacts of each individual root delay as it propagates downstream, weighting these delays by their relative probabilities. There are of course limitations to this approach. First, we are not considering recovery decisions, which can impact how delays propagate. Incorporating recovery is a sizeable challenge, largely due to the fact that recovery decisions are not pre-defined but rather based on individual personnel’s prior experience and intuition. We nonetheless suggest that incorporating some baseline rules for recovery interventions would be an important extension of our model to consider. Second, we do not consider the interactions of root delays and, as a result, we over-count propagation. For example suppose that the aircraft of flight f connects to flight g , while the crew of flight h also connects to flight g . When we consider the impact of a root delay to flight f , we may capture a downstream impact on flight g . The same will occur when we consider the impact of a root delay to flight h . In reality, should these two root delays occur concurrently, then their effect on flight g will typically *not* be additive. Of course, this over-counting will occur in both the original schedule and our proposed alternative.

While these limitations will impact the quality of our results, we nonetheless suggest that our surrogate function can improve over the existing schedule without increasing planned costs, enabling carriers to see immediate improvements in their operational performance while the research community continues to seek ways to provide further benefits. To support this hypothesis, we have developed a discrete-event simulation model (presented in Section 3.4) to help assess the quality of our proposed solutions.

3.2 Single-Layer Model

We begin by presenting a *single-layer model (SLM)* for redistributing slack. This model only considers the impact of disruptions one layer “downstream” - that is, on flights directly connected to the flight experiencing the root disruption.

We present this model for two reasons. First, it provides a simple framework that will facilitate understanding of the more complex multi-layer model. Second, as demonstrated in section 4, even this simple model can yield non-trivial benefits.

3.2.1 Notation

Sets

\mathcal{F}	set of flights
\mathcal{A}	set of all considered connections
\mathcal{M}	set of possible delay values (minutes) - here we assume that \mathcal{M} is a discrete set

Parameters

$k_f^+ \geq 0$	$\forall f \in \mathcal{F}$	the amount by which the departure time of flight f can be moved later
$k_f^- \geq 0$	$\forall f \in \mathcal{F}$	the amount by which the departure time of flight f can be moved earlier
$0 \leq p_f^m \leq 1$	$\forall m \in \mathcal{M}, \forall f \in \mathcal{F}$	the probability that flight f experiences a root delay of m minutes
$s_{f_1, f_2} \geq 0$	$\forall (f_1, f_2) \in \mathcal{A}$	the slack between flights f_1 and f_2 in the original schedule

Decision Variables

$$\begin{aligned}
-k_f^- &\leq x_f \leq k_f^+ && \forall f \in \mathcal{F} \\
y_{f_1, f_2} &\geq 0 && \forall (f_1, f_2) \in \mathcal{A} \\
d_{f_1, f_2}^m &\geq 0 && \forall (f_1, f_2) \in \mathcal{A}, \forall m \in \mathcal{M}
\end{aligned}$$

x_f is the change in the departure time of flight f . This change is restricted to the specified range $[-k_f^-, k_f^+]$. Note that a negative value for x means the flight is shifted earlier and a positive value for x means that the flight is shifted later.

y_{f_1, f_2} corresponds to the new slack between flights f_1 and f_2 , according to the modified schedule.

d_{f_1, f_2}^m corresponds to the delay that will propagate from flight f_1 to flight f_2 in the new schedule if there is a root delay of m minutes imposed on flight f_1 .

3.2.2 Formulation

$$\text{(SLM)} \quad \min \sum_{m \in \mathcal{M}} \sum_{(f_1, f_2) \in \mathcal{A}} p_{f_1}^m d_{f_1, f_2}^m \tag{1}$$

s.t.

$$y_{f_1, f_2} = s_{f_1, f_2} - x_{f_1} + x_{f_2} \quad \forall (f_1, f_2) \in \mathcal{A} \tag{2}$$

$$d_{f_1, f_2}^m \geq m - y_{f_1, f_2} \quad \forall (f_1, f_2) \in \mathcal{A}, \forall m \in \mathcal{M} \tag{3}$$

$$d_{f_1, f_2}^m \geq 0 \quad \forall (f_1, f_2) \in \mathcal{A}, \forall m \in \mathcal{M} \tag{4}$$

$$-k_f^- \leq x_f \leq k_f^+ \quad \forall f \in \mathcal{F} \tag{5}$$

$$y_{f_1, f_2} \geq 0 \quad \forall (f_1, f_2) \in \mathcal{A} \tag{6}$$

The objective function (1) of SLM minimizes the expected value of the delay propagation by weighting the probability of an m -minute delay on flight f_1 times its propagation to f_2 , summed over all flight connections (f_1, f_2) and delay lengths m . [Recall that, in this objective function, we consider only one layer of propagation.]

Constraints (2) calculate the new slack (y_{f_1, f_2}) between two flights. This is the old slack (s_{f_1, f_2}) minus the change in the departure of the first flight (x_{f_1}), plus the change in the departure of the second flight (x_{f_2}). Note, as illustrated in Figure 4, that if flight f_1 is moved earlier, then x_{f_1} will have a negative value and the slack y_{f_1, f_2} will therefore *increase* because we are subtracting this value. Constraints (6) ensure that the connection stays feasible, i.e. that the slack remains non-negative.

Constraint sets (3) and (4) calculate how much delay would propagate from flight f_1 to its outbound connection f_2 if f_1 were to experience a root delay of m minutes. Specifically, the propagated delay is m minus y_{f_1, f_2} (the root delay minus the new slack between the flights) unless this is negative, in which case the delay is zero.

Finally, constraints (5) limit the amount by which flight times can be changed. Note that this is flight-specific and can be used not only to restrict changes so that market share is not impacted, but also to recognize gate limitations, slot restrictions, hourly departures (in which case the time window would be zero), etc. Key passenger itineraries could be protected as well, by placing a lower value on the slack time (y) between two connecting flights.

Claim: The solution to (2)-(6) will yield integer values of x .




Figure 4: Visualizing constraint (2): In this example, the departure of the first flight is shifted earlier by x_{f_1} (which is negative) and the departure time of the second flight is pushed forward by x_{f_2} (which is positive). As a result, the new slack (y_{f_1, f_2}) equals the old slack (s_{f_1, f_2}) minus x_{f_1} plus x_{f_2} .

Proof: In order to prove the claim, we need to show that the coefficient matrix presented by (2)-(6) is totally unimodular. In that case, given that all elements in the right-hand-side vector of (2)-(6) are integer, all the extreme points of (2)-(6) will be integer [18].

In order to show that the coefficient matrix defined by (2)-(6) is totally unimodular, it is sufficient to prove that the coefficient matrix corresponding to constraints (2) and (3) – or equivalently, its transpose – is totally unimodular. The constraints (4)-(6) are upper and lower bounds, which do not impact the claim.

We can re-write constraint (3) substituting (2), which yields:

$$d_{f_1, f_2}^m \geq m - s_{f_1, f_2} + x_{f_1} - x_{f_2} \quad \forall (f_1, f_2) \in \mathcal{A}, \forall m \in \mathcal{M} \quad (7)$$

By transferring the variables to the left-hand-side we get:

$$d_{f_1, f_2}^m - x_{f_1} + x_{f_2} \geq m - s_{f_1, f_2} \quad \forall (f_1, f_2) \in \mathcal{A}, \forall m \in \mathcal{M} \quad (8)$$

Here we can see that the coefficient matrix (A) corresponding to constraint (8) has entries (a_{ij}) of only -1, 0 and 1. Furthermore, we can partition the columns of this matrix into two sets: \mathcal{C}_1 includes the columns corresponding to the d variables and \mathcal{C}_2 the columns corresponding to the x variables. We see that for each row of this matrix, the summation of all the coefficients in \mathcal{C}_1 equals 1 and the summation of all the coefficients in \mathcal{C}_2 equals 0. Therefore:

$$\left| \sum_{j \in \mathcal{C}_1} a_{ij} - \sum_{j \in \mathcal{C}_2} a_{ij} \right| \leq 1 \quad \forall i \quad (9)$$

Therefore, A' (the transpose of A) is totally unimodular, and thus A itself is also totally unimodular. \square

This claim tells us that it is sufficient to solve SLM as an LP, rather than an IP, while still yielding integer departure times. This has significant implications not only for the tractability of the model, but also for our ability to extend it to include indirect downstream effects, as we see in the next section.

3.3 Multi-Layer Model

The model presented in section 3.2 only considers the impact of a root delay on the flight's immediate connections. Of course, these delayed connections can in turn delay their outbound connections as well if there is not enough slack to fully absorb the disruption. In this section, we present a *multi-layer model (MLM)* in which the downstream propagation of delays continues until they are fully absorbed.

To formulate this model, we must first define the notion of a *propagation tree*. For each root flight f_0 and each delay value m , the propagation tree represents the set of all downstream flights that could potentially be delayed as a result of the root delay propagating. Note that because we construct these trees *before* re-timing

the flights, we do not know for sure which flights will receive propagated downstream delays. Therefore the propagation tree considers the “worst-case” scenario.

To construct a propagation tree for root flight f_0 and delay value m , we look at each of its outbound connections. For example, let f_1 be the flight awaiting flight f_0 's aircraft. We assume that f_0 is moved to its latest permitted departure time and f_1 is moved to its earliest departure time – this creates the minimum possible slack between each flight pair. Clearly, if a delay of length m does not propagate across this slack, then it will not propagate across any re-timing of the flight pair. On the other hand, if it does, then the connecting flight is added to the propagation tree. For each outbound connection that is added to the propagation tree, we then look at all of its outbound connections, again assuming the minimum amount of slack in the connection. Observe that for a given root flight f_0 , the propagation tree for a delay of length m_1 will be a subset of the propagation tree for the delay of length $m_2 > m_1$. Figure 5 presents two such propagation trees, where the potential propagation from f_i to f_j is calculated as f_i 's root delay minus the slack between f_i and f_j plus the maximum amount by which the departure time of f_i can be moved later ($k_{f_i}^+$) plus the maximum amount by which the departure time of f_j can be moved earlier ($k_{f_j}^-$).

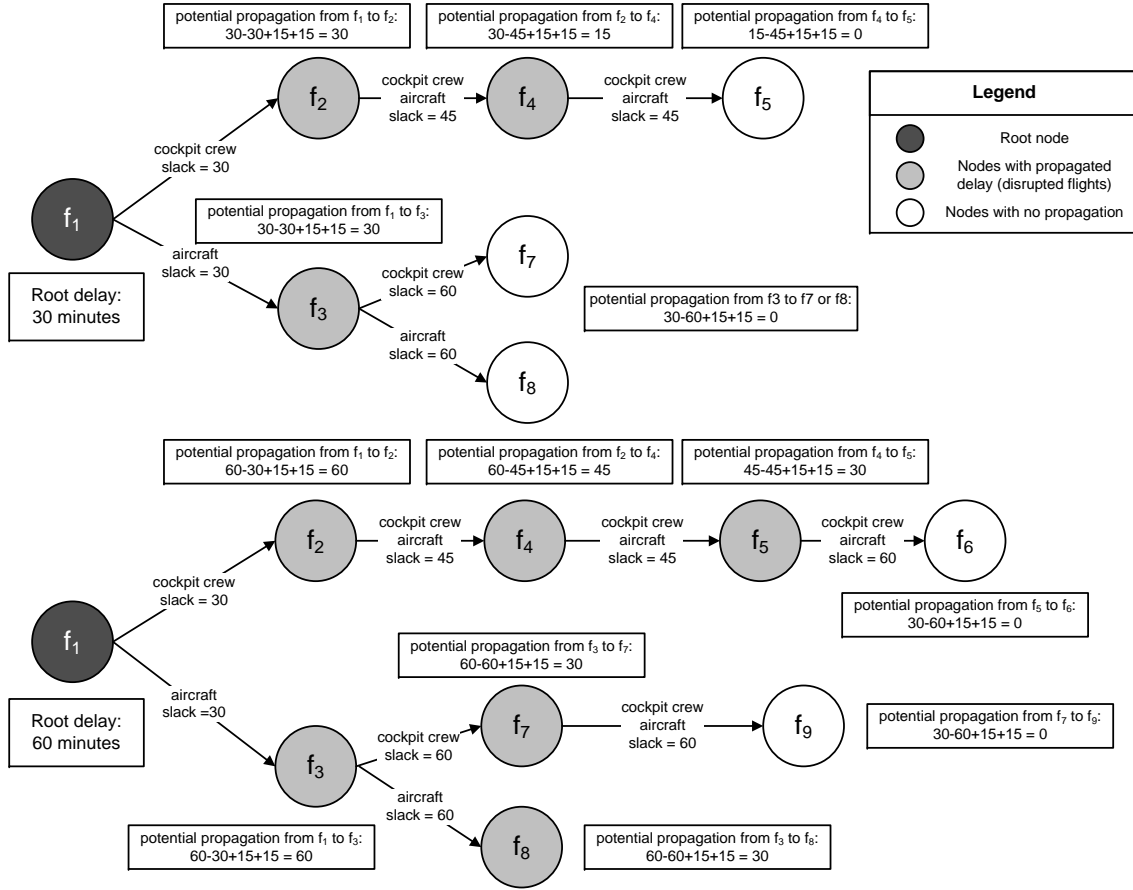


Figure 5: Visualizing propagation trees: The propagation tree with root flight f_1 and root delay 30 ($\mathcal{T}_{f_1}^{30}$, above) has fewer nodes compared to the propagation tree with root flight f_1 and root delay 60 ($\mathcal{T}_{f_1}^{60}$, below).

We define the following notation:

$\mathcal{T}_{f_0}^m$ set of flights in the propagation tree associated with root flight f_0 and root delay m (excluding f_0)
 $r_{f_0}^m(f)$ the parent node of flight f in $\mathcal{T}_{f_0}^m$.

Based on this notation, we extend the single-layer model to the multi-layer model:

$$(MLM) \quad \min \sum_{m \in \mathcal{M}} \sum_{(f_0 \in \mathcal{F}, f \in \mathcal{T}_{f_0}^m)} p_{f_0}^m d_{f_0, f}^m \quad (10)$$

s.t.

$$y_{f_1, f_2} = s_{f_1, f_2} - x_{f_1} + x_{f_2} \quad \forall (f_1, f_2) \in \mathcal{A} \quad (11)$$

$$d_{f_0, f}^m \geq m - y_{f_0, f} \quad \forall (f_0 \in \mathcal{F}, f \in \mathcal{T}_{f_0}^m : r_{f_0}^m(f) = f_0), \forall m \in \mathcal{M} \quad (12)$$

$$d_{f_0, f}^m \geq d_{f_0, r_{f_0}^m(f)}^m - y_{r_{f_0}^m(f), f} \quad \forall (f_0 \in \mathcal{F}, f \in \mathcal{T}_{f_0}^m : r_{f_0}^m(f) \neq f_0), \forall m \in \mathcal{M} \quad (13)$$

$$d_{f_0, f}^m \geq 0 \quad \forall (f_0 \in \mathcal{F}, f \in \mathcal{T}_{f_0}^m), \forall m \in \mathcal{M} \quad (14)$$

$$-k_f^- \leq x_f \leq k_f^+ \quad \forall f \in \mathcal{F} \quad (15)$$

$$y_{f_1, f_2} \geq 0 \quad \forall (f_1, f_2) \in \mathcal{A} \quad (16)$$

There are two key differences between the single-layer model presented in section 3.2 and this model. First, in the objective (10), we include not only the delay minutes that a root delay of m minutes on flight f_0 imposes on f_0 's immediate outbound connections, but also the propagated impact on all subsequent flights in f_0 's propagation tree.

Second, constraints (13) enforce the fact that the delay propagated to a flight downstream from the root delay will be the amount of delay propagated to its parent minus the amount of slack between these flights.

This model is structurally quite similar to the single-layer model and exhibits the same integrality property. The main difference is in size. For each flight in any given propagation tree, we have to add both a new variable and a new constraint. As we observe in Section 4, however, the modified formulation remains highly tractable for networks of a realistic size.

Claim: The solution to (11)-(16) will yield integer values of x .

Proof: The proof is similar to the one presented in 3.2. Here we argue that the coefficient matrix presented by (11)-(13) is totally unimodular. (The constraints (14)-(16) again do not affect the integrality of the x variables.)

As in the earlier proof, we can re-write constraints (12) by substituting in constraints (11), yielding:

$$d_{f_0, f}^m - x_{f_0} + x_f \geq m - s_{f_0, f} \quad \forall (f_0 \in \mathcal{F}, f \in \mathcal{T}_{f_0}^m : r_{f_0}^m(f) = f_0), \forall m \in \mathcal{M} \quad (17)$$

Note that this constraint, which captures the relationship between the root flight and each of its ‘‘children’’, includes one d variable, with coefficient of 1, and two x variables, one with coefficient 1 and one with coefficient -1.

Next we consider constraint set (13). First, suppose that f is a ‘‘grandchild’’ of f_0 – in other words, its inbound connection is f_0 's outbound connection. Then by substituting (11) we get:

$$d_{f_0, f}^m \geq d_{f_0, r_{f_0}^m(f)}^m - s_{r_{f_0}^m(f), f} + x_{r_{f_0}^m(f)} - x_f \quad (18)$$

and then by substituting (12) we get:

$$d_{f_0, f}^m \geq m - s_{f_0, r_{f_0}^m(f)} + x_{f_0} - x_{r_{f_0}^m(f)} - s_{r_{f_0}^m(f), f} + x_{r_{f_0}^m(f)} - x_f \quad (19)$$

Canceling and moving the variables to the left hand side of the constraint yields:

$$d_{f_0,f}^m - x_{f_0} + x_f \geq m - s_{f_0,r_{f_0}^m(f)} - s_{r_{f_0}^m(f),f} \quad (20)$$

which again has one d variable with coefficient one and two x variables, one with coefficient 1 and one with coefficient -1.

More generally, when f is a descendant of f_0 in the propagation tree, substitution as above will yield:

$$d_{f_0,f}^m - x_{f_0} + x_f \geq m - \sum_{(f_1,f_2) \in \mathcal{P}_{f_0}^m(f)} s_{f_1,f_2} \quad (21)$$

as all the intermediate x variables cancel, with each intermediate flight on the path from f_0 to f playing the role of both parent and child. [In equation 21, $\mathcal{P}_{f_0}^m(f)$ is the set of all arcs (f_1, f_2) that form the path from the root node f_0 to f in the propagation tree $\mathcal{T}_{f_0}^m$.] Again, we see one d variable with coefficient 1 and two x variables, one with coefficient 1 and one with coefficient -1.

Thus, using the same partitioning of the columns of the A matrix as the earlier proof, the sufficient condition for total unimodularity follows directly. \square

3.4 Simultaneous Delays Model

The models presented in Sections 3.2 and 3.3 rely on the use of a surrogate objective function to achieve robustness. In particular, objective functions (1) and (10) attempt to minimize the total amount of propagated delay in the flight network by looking at how each individual flight propagates delay. The primary limitation of this approach is that it does not take into account the fact that multiple flight delays may occur in the network simultaneously. As a result, propagated delay will in some cases be estimated inaccurately. Figure 6 presents one such situation.

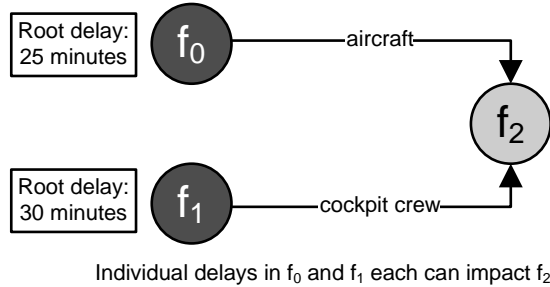


Figure 6: An example of simultaneous root delays where the surrogate objective function *overestimates* the delay propagations.

Given flights f_0 and f_1 , both of which are parent (preceding) flights of f_2 , the objective function includes the sum of the delays propagated from both f_0 and f_1 . In scenarios where both f_0 and f_1 experience a root delay simultaneously, however, f_2 's delay should be the *maximum* of the two upstream delays, rather than their sum.

Conversely, there are occasions where the surrogate objective function underestimates delay. For example, as illustrated in Figure 7, when simultaneous root delays happen consecutively, their overall results can be more disruptive than what the objective functions in (1) or (10) can estimate. In this figure, individual root delays in f_0 or f_1 are not enough to delay f_2 . However, when f_0 and f_1 both experience a thirty-minute delay, their delays *can* propagate to cause a ten-minute delay in f_2 .




Figure 7: An example of simultaneous root delays where the surrogate objective function *underestimates* the delay propagations.

Given that both the original and the re-timed schedules suffer from these inaccuracies, it is not clear how the overall reduction in delay propagation is affected. Therefore, we have developed a discrete-event simulation model to provide further analysis into the overall effects of this approximation. This enables us to compare how both the original and the optimized schedules perform when multiple flights in the network incur root delays concurrently.

Note that we assume root delays to be additive, meaning that if a flight has previously incurred a propagated delay, but additionally incurs a root delay, the total departure time is pushed back by the sum of these values. For example, consider the case when a flight is delayed because it is awaiting its (delayed) inbound aircraft. A weather delay or mechanical delay (i.e. a root delay) at the second flight would typically not arise until after the incoming aircraft arrived, making these delays additive. On the other hand, when two different upstream delays simultaneously affect a flight (e.g. its crew is delayed on one inbound flight and its aircraft delayed on another), then the propagation will be the maximum of the two propagated delays, rather than their sum.

To accurately simulate a given flight network using a discrete-event simulation, we require a source of randomness that will be used to determine the amount of root delays incurred by each flight. To model these delays, we have generated empirical distributions based on the origin airport of the departing flight, using actual delay data spanning a 12-month period. Specifically, we filtered out all delays that were propagated from an upstream root disruption. We then clustered the remaining delays by origin station and length of disruption. The corresponding probability mass functions were used in both the simulation, to randomly generate the initial root delays, and as coefficients in the objective functions of SLM and MLM.

Note that in a given flight schedule, there is no obvious “first flight” – because the schedule is repeating, any particular flight can appear as both the root of one propagation tree and a downstream flight in another. To address this complication, our simulation algorithm employs a recursive strategy that explores all flights in a given network without requiring a sequential ordering [14].

The delay that a flight experiences consists of two parts – propagated delay (resulting from the need to wait for delayed upstream resources such as cockpit crews, cabin crews and aircraft) and root delay (associated with the flight itself, such as a weather delay). The propagated delay is computed by taking the maximum delay across all of the inbound resources, then adding this to the root delay.

To compute the total propagated delay in the network, we first initialize all flights in the network to have a propagated delay of zero and a root delay of zero. After this initialization stage, the propagation algorithm proceeds by executing the following steps for each flight in the network.

We pick an arbitrary flight in the network to process and call upon the random delay generator to provide a root delay for this particular flight. If this root delay is non-zero, we update the departure time of the current flight to be the sum of the (current value of the) propagated delay and this additional root delay. We then consider all outbound connections from this particular flight. For each connection, we determine how much (if any) of the current flight’s delay would propagate to this child. We then check whether this

propagated delay is larger than the connecting flight’s current propagated delay. If it is, we update the connecting flight’s propagated delay. Next, we add this to the connecting flight’s root delay (which will be zero if that flight hasn’t been processed yet) to determine its new departure time. Finally, we recursively use this new departure time to update the propagated delay of its outbound connections, repeating until a flight delay does not propagate.

This process repeats until all flights in the network have been processed and therefore exposed to the possibility of incurring a root delay. Through the recursive nature of the algorithm, we are guaranteed to explore every flight and every propagation tree, updating them with the appropriate amount of propagated delay.

4 Computational Experiments

The objectives of our computational experiments were three-fold. First, we wanted to assess the run-time performance of our models and determine their tractability. Second, we wanted to evaluate the extent to which minor changes in flight times could impact the potential for delays to propagate – would the optimized schedule have significantly less delay propagation than the original schedule? Third, we wanted to use simulation to assess the accuracy of our surrogate objective function, i.e. our metric for delay propagation. When multiple delays are allowed to occur simultaneously, as is the case in reality, does the our new schedule still show improved performance over the original schedule?

Our computational experiments were conducted using data provided by a major U.S. carrier offering over 500 flights per day. We considered two different dates in 2007. For each of these dates, we were given the flight schedule, cockpit crew schedules, and aircraft rotations. Thus, we were able to consider propagation of delays associated with incoming aircraft and cockpit crews. [We did not include delays associated with cabin crews or connecting passengers, due to lack of data.]

As a default, we assumed that each flight was allowed to be moved up to fifteen minutes earlier or fifteen minutes later than its original departure time. Note that, in theory, this could impact crew costs and feasibility. For example, if a duty were currently at its maximum allowed elapsed time and we moved the departure time of the first flight of the duty earlier and/or moved the departure time of the last flight of the duty later, then we would violate the elapsed time limit.

To account for this, we considered four different variations. In the first and most restrictive case, we assumed that the first flight of any duty could not be moved earlier and the last flight of any duty could not be moved later. This greatly limits the flexibility of the network. In both of the data sets considered, most of the flights are either the first or last flights of their duty, and thus only about twenty-five percent of the flights were allowed to move freely. Furthermore, this overly restricts the system: in a duty whose cost is dominated by flying time, increasing the elapsed time by stretching the first and last flights slightly may have no impact on cost, and the feasibility of the duty may be unchanged as well.

Therefore, we considered three additional instances, each progressively less restrictive. In these, the first and last flights of the duty were only allowed to change by five minutes, by ten minutes, or by the same fifteen minutes as all other flights in the network.

All code was implemented and run on an Intel® Pentium® D 3.20 GHz CPU architecture using the C++ programming language. The optimization models were developed using CPLEX/Concert Technology and solved using CPLEX 11.0 solver.

4.1 Goal 1 – Tractability

Table 2 shows the size and run time of each problem instance solved. The first column of this table indicates which of the two data sets is being considered. The second column indicates the restriction on first and last flights in a duty (note that all other flights have a \pm fifteen minute time window in all instances). The third column indicates whether the single-layer or multi-layer model is being used. The fourth, fifth, and sixth columns give the size of the model (in number of constraints, number of variables and number of nonzero elements) and the final column gives the run time (in seconds).

Data Set	Duty Restriction	Model	No. of Constraints	No. of Variables	No. of Nonzeros	Run Time
1	0	SLM	1,319	1,592	2,869	4
1	5	SLM	2,278	2,724	4,938	3
1	10	SLM	2,498	2,975	5,420	4
1	15	SLM	2,701	3,217	5,867	4
1	0	MLM	4,745	5,048	12,328	7
1	5	MLM	5,874	6,356	15,626	7
1	10	MLM	5,995	6,494	15,853	7
1	15	MLM	6,108	6,636	16,080	8
2	0	SLM	1,041	1,251	2,245	3
2	5	SLM	2,081	2,505	4,509	3
2	10	SLM	2,269	2,714	4,906	3
2	15	SLM	2,444	2,923	5,289	2
2	0	MLM	3,932	4,178	10,035	6
2	5	MLM	5,131	5,573	13,526	7
2	10	MLM	5,273	5,733	13,835	6
2	15	MLM	5,376	5,865	14,081	6

Table 2: Size of the instances and their corresponding run times.

Note that all run times are less than 10 seconds, suggesting that there is no computational limitation on either model. Although the model size increases substantially when moving from the single-layer model to the multi-layer model, the fact that the problem can be solved as an LP implies that tractability is not sacrificed when using the more accurate model.

Finally, these results suggest that even if a significant number of additional connections were considered (e.g. incorporating cabin crews, key passenger itineraries, etc), run times would remain tractable.

4.2 Goal 2 – Impact

Our second goal was to evaluate the potential improvements in delay propagation to be gained by re-timing flights slightly. First, we began by computing two estimates of the propagated delay of the original schedules – one using the objective function from the single-layer model and one using the more accurate objective function from the multi-layer model. We then optimized the original schedules twice, once using each of the two objective functions. Again, for each of these solutions, we computed both the SLM and MLM objective estimates of delay propagation.

Table 3 presents these results. The first column specifies which of the two data sets is being analyzed. The second column indicates how many minutes the first/last flights of a duty are allowed to change. The next three columns present the objective function value of the original, single-layer optimal, and multi-layer optimal schedules relative to the single-layer objective function. For the optimized schedules, ($n\%$) indicates the percent improvement over the original schedule, relative to the single-layer objective function. Likewise, the following three columns represent the multi-layer objective function applied to the three schedules.

Observe that, as expected, the amount of propagation is larger for all scenarios when multiple layers of propagation are taken into account. In addition, the single-layer optimal solution of course performs better

Data Set	Duty Restr.	Evaluated w/ SLM Obj. Function			Evaluated w/ MLM Obj. Function		
		Orig. Sch.	SLM Opt. Sol.	MLM Opt.Sol.	Orig. Sch.	SLM Opt.Sol.	MLM Opt. Sol.
1	0	1,294.24	1212.89 (6.3%)	1231.07 (4.9%)	2,268.55	2156.22 (5.0%)	2104.95 (7.2%)
1	5		996.102 (23.0%)	1018.69 (21.3%)		1694.16 (25.3%)	1647.24 (27.4%)
1	10		859.65 (33.6%)	882.99 (31.8%)		1380.14 (39.2%)	1340.32 (40.9%)
1	15		756.49 (41.5%)	771.08 (40.4%)		1152.08 (49.2%)	1112.17 (51.0%)
2	0	1,192.81	1129.54 (5.3%)	1139.59 (4.5%)	2,037.05	1944.77 (4.5%)	1919.01 (5.8%)
2	5		912.88 (23.5%)	929.98 (22.0%)		1518.1 (25.5%)	1481.76 (27.3%)
2	10		782.05 (34.4%)	800.21 (32.9%)		1224.79 (39.9%)	1192.33 (41.5%)
2	15		678.61 (43.1%)	691.48 (42.0%)		1008.55 (50.5%)	973.816 (52.2%)

Table 3: Reduction in the objective function (potential delay propagation) – comparing different schedules

under the single-layer objective, while the multi-layer optimal solution performs better under the multi-layer optimization.

It is interesting to note that there is not a dramatic difference between the single- and the multi-layer schedules in their performance under the multi-layer objective function. This is presumably due to the fact that the delays dissipate fairly quickly and thus the immediate outbound connection plays a dominant role. [See [3] for an empirical analysis of the characteristics of propagation trees.] Therefore, minimizing the first layer of delays will capture much of the possible benefits.

Finally, we observe that flight re-timing *can* have substantial improvements on the delay propagation, ranging from approximately 5 percent for the most tightly restricted instances to approximately 50 percent for the least restricted instances. Of course, there are several caveats that must accompany these results. First, we did not take into account the delay propagation associated with cabin crews or any passenger itineraries for which the second flight of the itinerary would be “held” for connecting passengers (these could easily be incorporated, given available data, and would have little impact on run times). It is not clear what impact these additions would have on the quality of the optimal solutions, as the changes would impact both the original and the modified schedules. Second, we have not taken into account recovery decisions, which again will affect propagation within both the original and the optimized schedules. Finally, we re-iterate that our surrogate objective function over-counts propagation because it considers each delay one at a time. We address this limitation through the use of a discrete-event simulation model in the next section.

4.3 Goal 3 – Validity

In order to better assess the impact of the fact that the surrogate objective function does not incorporate the concurrency of delays in our optimization model, we simulated each of the schedules (the original schedule and the re-timed schedules based on the single-layer and multi-layer models) using the same probability distribution functions for the root delays as in the optimization models.

The results are summarized in Table 4. The first two columns describe the instance. The next three columns give the ratio of the expected amount of delay propagation (in minutes), as estimated by 2000 replications of the simulation model, divided by the value of the surrogate objective function under the original or re-timed schedule.

As expected, the simulated values differ from the surrogate values. However, the ratio is fairly consistent, suggesting that the impact of ignoring concurrent delays has comparable impact on both the original and the re-timed schedules. Thus, it is not surprising that the simulated value of the re-timed schedules still demonstrates a significant improvement over the simulated value of the original schedule. Table 5 demonstrates this. The first two columns describe the instance. The next three columns give the expected amount of delay propagation (in minutes), as estimated by the simulation model, for each of the three schedules. Columns four and five also provide the relative improvement over the original schedule. Although these improvements

Data Set	Duty Restriction	Simulation / Surrogate		
		Orig. Sch.	SLM	MLM
1	0	1.187	1.224	1.180
1	5		1.261	1.200
1	10		1.275	1.203
1	15		1.288	1.202
2	0	1.168	1.193	1.166
2	5		1.224	1.170
2	10		1.238	1.170
2	15		1.252	1.158

Table 4: Comparing the simulation results with the surrogate objective function value.

Data Set	Duty Restrictions	Orig. Sch. Average	SLM Average	MLM Average
1	0	2692.1	2576.3 (4.3%)	2543.9 (5.5%)
1	5		2076.6 (22.8%)	2033.7 (24.4%)
1	10		1708.6 (36.5%)	1660.1 (38.3%)
1	15		1432.1 (46.8%)	1384.8 (48.6%)
2	0	2378.6	2289.0 (3.8%)	2266.7 (4.7%)
2	5		1813.9 (23.7%)	1776.4 (25.3%)
2	10		1476.6 (37.9%)	1433.0 (39.7%)
2	15		1219.3 (48.7%)	1167.8 (50.9%)

Table 5: Reduction in the delay propagation (simulation results) – comparing different schedules

are lower than the surrogate objective values, they nonetheless demonstrate a substantial opportunity for reduction in delay propagation through minor flight re-timings.

4.4 Discussion

We conclude this section with a few final observations.

First, we note that changing flight times can make key passenger itineraries infeasible by decreasing the connection time between them beyond that which is feasible for a passenger to connect. Such itineraries can be protected by adding constraints of the form:

$$y_{f_1, f_2} = s_{f_1, f_2} - x_{f_1} + x_{f_2} \quad (22)$$

$$y_{f_1, f_2} \geq 0 \quad (23)$$

which says that, for a protected itinerary (f_1, f_2) , the time between the scheduled arrival of flight f_1 and the scheduled departure of flight f_2 must not decrease below the minimum passenger connection time. Observe, however, that we do not include delay propagation from flight f_1 to flight f_2 in the model, because we do not assume that flights are delayed to await connecting passengers.

Second, we observe that – as with any model – the re-timing solutions that result from our models will only be starting points, which will need to be fine-tuned before implementation. In particular, although the 0-minute model may be overly restrictive in terms of maintaining the current crew costs and feasibility, the 15-minute model may result in some crew violations. These violations would have to be corrected manually, but we suggest that substantial benefits can still remain even after these modifications.

Third, to partially reduce this post-processing, we suggest that one way to reduce crew infeasibilities would be to identify the largest amount by which the length of the duty could increase without violating

elapsed time limits and/or changing cost the dominant cost from the flight time to the elapsed time. Then, for each duty, a constraint of form:

$$-x_{f'} + x_{f''} \leq e^{max} - e \tag{24}$$

could be used to ensure that the length of the duty not increase beyond this limit. In the above constraint, f' and f'' are the first and the last flights of a duty respectively, e is the elapsed time of that particular duty, and e^{max} is the maximum elapsed time in a duty.

5 Conclusions and Future Research

Airline delays have significant negative impact on airline costs, passenger convenience and productivity, and the environment. One major cause of delay is the down-stream propagation of initial delays to subsequent flights. This issue is exacerbated by the fact that slack is undesirable from a planning perspective, as it “wastes” costly resources, but is critical from an operational perspective, as it can be used to absorb disruptions and prevent their propagation.

Addressing operational concerns in the planning process can be quite challenging, however. First, metrics for evaluating the operational performance of a planned schedule must be developed. Second, cost functions must be developed to trade-off planned and (anticipated) operational costs. Finally, these cost functions must be incorporated in an already-challenging planning process. In particular, because delay propagation spans across multiple resources, schedule design, fleet, crew scheduling, and maintenance routing must all be considered concurrently.

As an intermediate measure to partially decrease the propagation of delay while these other challenging topics are studied by the research community, we propose to modify flight departure times within the framework of an existing airline plan. By re-allocating the existing slack to those flight connections that are most prone to delay propagation, we can reduce downstream impacts without changing planned crew or fleet costs and without changing revenue projections. Our computational results show significant opportunities for improvement without any increase in planned costs.

Future research in this area of course includes the three issues raised above – metrics for evaluating the robustness of a planned schedule, cost functions for computing the trade-off between planned costs and anticipated operational costs, and methods for incorporating these cost functions in the planning process. In the shorter time horizon, our research could be expanded by more explicitly addressing correlations between different root delays. Finally, we are interested in including recovery decisions (e.g. canceling flights, swapping aircraft, and calling in reserve crews) in our model to capture their impacts on the operational performance of a planned schedule under disruptions.

Acknowledgments:

We gratefully acknowledge the financial support of the Alfred P. Sloan Foundation Industry Studies Program and the MIT Global Airline Industry Program.

References

- [1] J. Abara. Applying integer linear programming to the fleet assignment problem. *Interfaces*, 19:20–38, 1989.
- [2] A. Abdelghany, G. Ekollu, R. Narasimhan, and K. Abdelghany. A proactive crew recovery decision support tool for commercial airlines during irregular operations. *Annals of Operations Research*, 127:309–331, 2004.

- [3] S. AhmadBeygi, A.M. Cohn, Y. Guan, and P. Belobaba. Analysis of the potential for delay propagation in passenger airline networks. Technical report, 2007. <http://www.umich.edu/~amy-cohn/PAPERS/proptreefinal.pdf>, to appear in Journal of Air Transport Management.
- [4] C. Barnhart, A. Cohn, E. Johnson, D. Klabjan, G. Nemhauser, and P. Vance. *Handbook of Transportation Science*, chapter Airline Crew Scheduling. Kluwer’s International Series, 2003.
- [5] C. Barnhart and K. Talluri. *Design and Operation of Civil and Environmental Engineering Systems*, chapter Airline Operations Research, pages 435–469. John Wiley and Sons, New York, 1997.
- [6] M. Berge. Timetable optimization: Formulation, solution approaches, and computational issues. pages 341–357, 1994. AGIFORS proceedings, Zandvcort, The Netherlands.
- [7] S. Bratu and C. Barnhart. Flight operations recovery - new approaches considering passenger recovery. *Journal of Scheduling*, 9:279–298, 2006.
- [8] E.K. Burke, P. De Causmaecker, G. De Maere, J. Mulder, M. Paelinck, and G. Vanden Berghe. A multi-objective approach for robust airline scheduling. *Computers and Operations Research*, (to appear), 2007.
- [9] M. Clarke, L. Lettovsky, and B. Smith. *Handbook of Airline Operations*, chapter The Development of the Airline Operations Control Center, pages 131–147. Aviation Week, 2000.
- [10] B. Fuhr. Robust flight scheduling - an analytic approach to performance evaluation and optimization. Technical report, 2007. Clausthal University of Technology - Lufthansa Airlines.
- [11] D. Klabjan, E. L. Johnson, and G. L. Nemhauser. Airline crew scheduling with regularity. *Transportation Science*, 35(4):359–374, 2001.
- [12] N. Kohl, A. Larsen, J. Larsen, A. Ross, and S. Tiourine. Airline disruption management perspectives, experiences and outlook. *Journal of Air Transport Management*, 13:149–162, 2007.
- [13] S. Lan, J.P. Clarke, and C. Barnhart. Planning for robust airline operations: Optimizing aircraft routings and flight departure times to minimize passenger disruptions. *Transportation Science*, 40(1), 2006.
- [14] M. Lapp, S. AhmadBeygi, A.M. Cohn, and O. Tsimhoni. A recursion-based approach to simulating airline schedule robustness. Technical report, 2008. To be submitted to proceedings of the 2008 Winter Simulation Conference.
- [15] L. Lettovsky, E.L. Johnson, and G.L. Nemhauser. Airline crew recovery. *Transportation Science*, 34(4):337–348, 2000.
- [16] A. Levin. Scheduling and fleet routing models for transportation systems. *Transportation Science*, 5:232–255, 1971.
- [17] A. Mercier and F. Soumis. An integrated aircraft routing, crew scheduling and flight retiming model. *Computers and Operations Research*, 34(8):2251–2265, 2007.
- [18] George L. Nemhauser and Laurence A. Wolsey. *Integer and combinatorial optimization*. Wiley-Interscience, New York, NY, USA, 1988.
- [19] B. Rexing, C. Barnhart, T. Kniker, A. Jarrah, and N. Krishnamurthy. Airline fleet assignment with time windows. *Transportation Science*, 34(1):120, 2000.
- [20] J. Rosenberger, E. Johnson, and G. Nemhauser. A robust fleet-assignment model with hub isolation and short cycles. *Transportation Science*, 38(3):357–369, 2004.
- [21] J.M. Rosenberger, A.J. Schaefer, D. Goldsman, A.J. Kleywegt, E.L. Johnson, and G.L. Nemhauser. A stochastic model of airline operations. *Transportation Science*, 36(4):357–377, 2002.

- [22] A. Schaefer, E. Johnson, J. Kleywegt, and G. Nemhauser. Airline crew scheduling under uncertainty. *Transportation Science*, 39(3):340–348, 2005.
- [23] G. Stojkovic, F. Soumis, J. Desrosiers, and M. M. Solomon. An optimization model for a real-time flight scheduling problem. *Transportation Research Part A*, 36:779–788, 2002.
- [24] M. Stojkovic and F. Soumis. An optimization model for the simultaneous operational flight and pilot scheduling problem. *Management Science*, 47(9):1290–1305, 2001.
- [25] S. Yan and D.H. Yang. A decision support framework for handling schedule perturbation. *Transportation Research Part B*, 30(6):405–419, 1996.
- [26] J. Yen and J. Birge. A stochastic programming approach to the airline crew scheduling problem. *Transportation Science*, 40(1):3–14, 2006.
- [27] G. Yu and X. Qi. *Disruption management: Framework, Models and Applications*. World Scientific Publishing, 2004.