Preemptive Rerouting of Airline Passengers under Uncertain Delays

July 15, 2015

An airline's operational disruptions can lead to flight delays that in turn impact passengers, not only through the delays themselves but also through possible missed connections. Since the length of a delay is often not known in advance, we consider *preemptive rerouting* of airline passengers before the length of the delay is realized. Our goal is to reaccommodate passengers proactively as soon as it is known that a flight will be delayed instead of waiting until passengers have missed connections. We consider the simplified version of the real-world problem in which only a single flight is delayed. We model this problem as a two-stage stochastic programming problem, with first-stage decisions that may preemptively assign passengers to new itineraries in anticipation of the delay's impact, and second-stage decisions that may further modify itineraries for any passengers who still miss connections after the delay has been realized. We present a Benders Decomposition approach to solving this problem and give computational results to demonstrate its effectiveness. This research lays the groundwork for the more-realistic case in which multiple flights in the network may experience concurrent delays.

1 Introduction

Airlines often face unexpected disruptions of their scheduled flight times, frequently caused by inclement weather, maintenance problems, and system congestion. A delay in one flight often propagates to other flights through delayed aircraft and crews. These operational disruptions can lead to flight delays that in turn impact passengers, not only through the delays themselves but also through resulting missed connections.

According to http://www.transtats.bts.gov (n.d.), in the period from January 2014 to February 2015, about 24% of flights in the US were delayed or cancelled (based on reporting carriers). Specifically, in the first two months of 2015, which represent the current available data for this year, 21.26% of flights were delayed and 3.61% were cancelled. It is clear that this high percentage of flight delays and cancellations causes passengers to miss many flight connections.

Passenger reaccommodation is commonly handled on an ad hoc basis, where each passenger is considered separately and only after a connection has been missed. Some airlines may watch for passengers who will be most affected when flights are delayed, especially those who have flights that are international or are at the end of the day, but reaccommodation is often done manually.

Unlike many other airline recovery problems, little research has been done on passenger reaccommodation. In this paper, we propose an approach to improve the recovery of passengers impacted by delays. In particular, we focus on doing so proactively, anticipating and addressing a possible missed connection, rather than waiting until after the missed connection has occurred.

Before looking at an example, we explain the idea of an itinerary. An itinerary is a passenger's plan consisting of one of more flights in a relative short time period that will end in the desired destination. Thus, an itinerary may be just one flight, and normally not more than 2 or 3. Also, this means that passengers on a given flight usually represent at least several different itineraries and therefore possible missed connections.

In particular, the length of a delay is often not known in advance, presenting challenges in making decisions for passengers, as illustrated by the following example. Suppose that a passenger has an itinerary consisting of two flights:

- Flight A from EWR to DTW, departing at 12:00pm and arriving at 2:10pm
- Flight B from DTW to LAX, departing at 3:40pm and arriving at 6:25pm

Further suppose that there is a mechanical delay on Flight A. The delay will either be 45 minutes, if the necessary part is available on site, or three hours, if the part must be flown in from another airport.

Other itineraries that have available capacity are the following:

- Itinerary 1:
  - Flight C from EWR to CMH, departing at 12:30pm and arriving at 2:30pm
  - Flight D from CMH to LAX, departing at 5:30pm and arriving at 8:00pm
- Itinerary 2:
  - Flight A from EWR to DTW, departing at 12:00pm and arriving at 2:10pm
  - Flight E from DTW to LAX, departing at 7:10pm and arriving at 9:55pm

Consider two possibilities. In the first, the passenger remains on her original itinerary. If the delay on her first flight turns out to be 45 minutes, then she will make her second flight and reach her destination on time. Conversely, if the delay is three hours, then she will miss her second flight. In this case, she will receive a new itinerary consisting of Flight E arriving at 9:55pm, 3 hours and 30 minutes later than planned.

Alternatively, if the passenger proactively switches to Flight C when the delay on her first flight is discovered, she will arrive at her destination at 8:00pm, 1 hour and 35 minutes later than planned.

In this example, keeping the original itinerary with the delayed flight leads to the passenger arriving either on time or 3.5 hours late, depending on the duration of the delay. Alternatively, changing itineraries proactively when her first flight's mechanical delay is discovered ensures that the passenger will arrive one hour and 35 minutes late. The optimal choice depends on the probability of each possibility for the length of delay and on the passenger's individual preferences. A passenger's priorities may depend on her destination, reason for travel, schedule once at her destination, and other considerations.

In actuality, airlines often have passengers keep their original itineraries until the delay length is realized, instead of considering the option of proactively moving the passengers onto different flights. If that is the case, after the delayed flight reaches its destination, recovery decisions are made and disrupted passengers are reaccommodated one-by-one.

This method has limitations in that, since passengers are not given new itineraries until they have already missed connections, alternate flight possibilities may be lost. Conversely, if reaccommodated proactively, passengers may be able to take itineraries that would otherwise no longer be available. This is the focus of our approach.

The remainder of the paper is arranged as follows. In Section 2, we review the literature, discussing work in passenger routing and, more broadly, airline recovery in general. Next, in Section 3, we present the *PRP (Preemptive Rerouting of Passengers)* Model. In Section 4, we discuss the implementation of the model, including how Benders Decomposition is used to find solutions in tolerable run times. Last, Section 5 contains a summary and discussion of future research plans.

The contribution of our research is in introducing a new approach to passenger reaccommodation that proactively handles passenger delays before misconnections occur. We present the PRP model and investigate a Benders Decomposition-based approach to finding tractable solutions. Our computational results show the approach's effectiveness for reducing the length of passenger delays. Thus, the model lays the groundwork for studying the more realistic problem in which multiple flights may experience concurrent delays.

2 Literature Review

When a flight is delayed or canceled, its passengers must be reaccommodated and the flow balance for the aircraft and crew must be preserved. Although there has been substantial research on recovery methods in passenger aviation, not much research has been done specifically in the area of passenger reaccommodation after a delay, and we did not find any other papers addressing preemptive rerouting of passengers. Below, we briefly discuss some models that route passengers, recovery of other factors, and recovery of multiple system resources including passengers.

2.1 Models that Route Passengers

There have been several papers that address routing of passengers, though not specifically re-routing after delays.

For example, within the fleet assignment model (FAM), it is necessary to consider not only assignment costs but also passenger revenues (often modeled through *spill costs,* the revenue loss associated with an aircraft that is too small to accommodate all demand for the flight). In Rexing (1997) and Rexing et al. (2000), the authors allow small changes in an original flight schedule in order to give better options while solving the fleet assignment problem and minimize operating costs and spill costs. To capture this, the model must explicitly consider the assignment of passengers to itineraries.

In Barnhart et al. (2002), the authors create the Itinerary-Based Fleet Assignment Model that approximates spill costs and recapture of passengers, producing better solutions for fleet assignments. The Passenger Mix Model is also described, where decisions are made as to what fraction of passengers from each itinerary to spill to each other itinerary given a solution to FAM. The authors of Jacobs et al. (2008) use origin and destination network effects and expected passenger flows in their fleet assignment model.

Finally, we note that in Rosenberger et al. (2002), the authors consider passenger inconvenience as one way to measure the quality of a schedule by counting the number of passenger misconnections from flights being canceled or delayed as they make decisions about which flights to cancel after disruptions.

2.2 Recovery of Factors Other Than Passengers

Although recovery models for passengers do not appear heavily in the literature, there is much published research on other types of recovery. In particular, recovery of aircraft, crews, and flight schedules has been studied in great depth. In addition, important research has been done on the topic of creating more robust schedules that respond better under delay situations in order to create schedules that perform better in practice.

For example, Lapp et al. (2008) study how robust a given flight schedule is, determining how delays can be propagated from one flight to others. In AhmadBeygi et al. (2010), the authors redistribute slack already existing in the system in order to lessen delay propagation. They allow small changes in the flight departure times, but do not allow changes in the fleet assignment solution or the crew scheduling solution, so that planned costs do not change, but operational performance can improve.

Delays still occur even with well-designed robust schedules. The Crew Recovery Model presented in Lettovsky (1997) and Lettovsky et al. (2000) gives solutions to the problem of creating new schedules for crews after flight delays. In Abdelghany et al. (2004), the authors present a tool for decision-making that proactively handles reaccommodation of crews. Their goal is to minimize cost from reassignments and delays. An example of a stochastic integer programming problem with recourse is given to solve the airline crew scheduling problem in Yen and Birge (2006). In the objective function of the model is the cost if the problem were deterministic, and there is also a term for the expected cost of recourse in case of disruptions.

In Rosenberger et al. (2003), the authors present a model where the decisions are what time the flights should be and how to reroute aircraft, and the goal is to minimize a function of rerouting and cancellation costs. In addition, in their section about future research, they discuss using a two-stage stochastic programming problem, taking into account the weather possibilities, to solve the aircraft recovery problem. The authors suggest using Benders Decomposition to solve this model to lessen the run time.

There are some similarities between aircraft and crew recovery and passenger recovery, since some of the same information is used for all three problems and resources are shared. Unfortunately, we cannot easily adapt the models for aircraft and crew recovery to help with passenger recovery. The objectives and the types of constraints are very different. For crew recovery, there are complicated work-time rules that must be followed, but there are not counterparts to these rules for passengers. Interestingly, Abdelghany et al. (2004) and Yen and Birge (2006) discuss preemptive rerouting of crews, which is similar to our work concerning passengers in some ways, but the models cannot easily be used for passengers. Similarly, in Rosenberger et al. (2003), the authors briefly suggest a method for preemptive rerouting for aircraft recovery in the future research section.

2.3 Recovery of Multiple System Resources Including Passengers

In the two models in Bratu and Barnhart (2006), the goal is to balance airline costs and passenger delay costs while several decisions are made in order to choose flight departure times and cancellations and reschedule aircraft, crew, and passengers after a delay or disruption has occurred. The first model minimizes operating costs and passenger disruption costs and the second minimizes operating costs and passenger delay costs. One of the models can be solved in real-time, and the models give noticeable reductions in passenger delays and disruptions.

In Eggenberg et al. (2010), the authors give a modeling framework, presenting a "recovery network," which can be used for different factors to be reaccommodated. Each feasible recovery method corresponds to a path through the nodes and arcs in the network, which has a two-dimensional coordinate system. The authors specifically discuss handling the passenger recovery problem using their approach.

In Jafari and Hessameddin Zegordi (2010), a model is solved after a disruption to reroute both aircraft and passengers simultaneously. The decisions include re-timing flights, switching aircraft, and giving passengers new itineraries. The authors use aircraft rotations and itineraries instead of flights for passengers, and include an option of reaccommodating passengers with other airlines or other transportation methods.

In Petersen et al. (2010), the authors solve simultaneously the recovery problems for the flight schedule, aircraft routings, crew schedules, and passengers' itineraries. This is advantageous, because a solution for a previous step may not produce the optimal solution for a subsequent step. Since solving this large MIP problem takes too long in most cases, the authors use Benders Decomposition, delayed column generation, and heuristics to decide which flights can be ignored as they are not affected by the delays, as well as other methods. The passenger recovery problem is modeled as a multi-commodity flow problem where flights correspond to arcs and passengers are assumed to be homogeneous. In using Benders Decomposition, the restricted master problem is the schedule generation problem, and the aircraft recovery model, the crew recovery model, and the passenger recovery model are the subproblems. Note that the three subproblems are independent of each other since they each depend only on the flight schedule. The authors present five types of Benders cuts, which are feasibility cuts for all three subproblems and optimality cuts for the crew and passenger recovery problems.

The papers presented in this section all discuss passenger reaccommodation after delays. While a goal for the authors is correcting passengers' itineraries when schedules fail, our strategy is recognizing uncertainty in how a plan will be corrected based on possibilities in delays, and to preemptively rebook passengers accordingly. In particular, we focus on doing so proactively, anticipating and addressing possible missed connections, rather than waiting until after the missed connections have occurred, giving more potential opportunities for passengers' arrival times at their destinations.

For a more comprehensive review of recovery models for flight schedules, aircraft, crew, and passengers, see Bratu and Barnhart (2006), Kohl et al. (2007), and Petersen et al. (2010). The book Yu and Qi (2004) is a good resource for better understanding disruptions and recovery in the airline industry. Belobaba et al. (2009) further discusses many of these topics in depth.

3 Rerouting of Airline Passengers

3.1 Problem Statement and Assumptions

We consider the problem of preemptive rerouting of airline passengers before the length of the delay of a flight is known. Our goal is to proactively reaccommodate passengers as soon as it is known that a flight will be delayed instead of waiting until passengers have missed connections.

For this problem, we assume that we are given one delayed flight with an unknown length of delay. There are a finite number of possible lengths of delay, and the probability of each candidate delay length is known. Note that we assume the delay distribution is fixed, so that our knowledge of the distribution does not change over time before the length of delay becomes known for certain. This happens shortly before the flight departs. We assume that the departure times for all other flights are known with certainty.

The decisions for this problem (in the first stage) are which updated itinerary to give to each passenger on the delayed flight when the delay is discovered. In addition, we must choose (in the second stage) alternate itineraries for any passengers who subsequently miss connections after the delay duration has been realized.

Our objective is to minimize the sum of the delay costs for all passengers on the delayed flight. We let the cost of an itinerary be the difference between a passenger's originally-scheduled arrival time at the desired destination and the actual arrival time, which depends on both the realization of the delay length and the reassignment of itineraries.

Since the passengers' originally scheduled arrival times vary, the cost assigned to a given itinerary depends on which passenger receives that particular itinerary. Note that more weight can be given to certain passengers, such as those who fly frequently or paid more for their tickets. The way we have chosen to measure the delay cost for a passenger is a function of only the arrival time in the final destination. No other things, such as flight time or layover location, are considered.

In defining the problem, we make the following assumptions:

- All passengers must be given first-stage itineraries that are scheduled to arrive by a given end time, and any itineraries assigned in the second stage must also arrive by this end time.
- No passenger can be assigned more than three flights in an itinerary.
- The available capacity of each flight cannot be exceeded.
- A passenger may be given a different first flight in the first-stage problem only when one of the following is true:
  - The passenger has two or more flights in his or her original itinerary and has some non-zero chance of being disrupted in at least one possible scenario.
  - The passenger has only the original flight in the itinerary. Since that flight will be delayed, it makes sense to have the option of moving to another itinerary.

3.2 Two-Stage Model and Solution Approach

In our approach, we first assign an itinerary to each passenger at the time that the delay is discovered. In some cases, passengers remain on their original itinerary. In other cases they are assigned to a new itinerary. Next, once the delay length is realized, we assign new itineraries to passengers whose (possibly new) stage-one itineraries have been disrupted. Our objective is to minimize the sum of all passengers' expected delays (possibly weighted to recognize preferential treatment of some passengers), where the delay is the additional time until landing compared to original itineraries.

We formulate this problem as a mixed-integer linear program with decision variables representing the two stages. For the second stage decisions, there is a set of variables for each of a finite set of possible outcomes for the delay. The costs for the second-stage problem in the different scenarios for the length of the delay are weighted by their probabilities. We then re-formulate and solve this large-scale integer programming problem using Benders Decomposition.

For both stages, we solve for an itinerary to give to each passenger. In the first-stage problem, we use an itinerary-based formulation, where binary variables represent the assignment of specific itineraries to specific passengers. Note that we can easily identify off-line, for each itinerary and each scenario, whether the itinerary would be disrupted under that scenario. This allows us to easily determine, for the second stage, what capacity remains available, i.e. which seats are reserved in the first stage but then become available in the second stage under a given scenario.

Conversely, we cannot use an itinerary-based formulation in the second-stage problem, because the integrality requirements would conflict with our intended use of a Benders Decomposition solution approach. We are able, however, to formulate this second stage as a pure minimum cost flow problem. Note that for the second stage, we do not have the same issue as in the first stage of determining capacity, as this is fixed from the first stage.

Before presenting the formulation of the PRP model, we describe its format. We define the following:

- Ω is the set of possible outcomes for the length of delay.
- $\omega \in \Omega$ is a possible outcome for the length of delay.
- ρ^{ω} is the probability of scenario ω.

- $\mathbf{x}$ is the vector of first-stage decisions.

- $\mathbf{y}$ is the vector of second-stage decisions.

- $f(\mathbf{x}, \omega)$ is the total delay associated with passengers completing their first-stage itineraries (which may arrive at a later time than their originally-scheduled itineraries) under scenario ω. The delay for an itinerary is weighted by the sum of the probabilities of the scenarios where the itinerary will not be disrupted. Note that if passengers have been assigned a new first-stage itinerary and those itineraries can be successfully completed under a given *delay scenario* (situation corresponding to outcome $\omega \in \Omega$), those passengers do not receive second-stage itineraries and there is no contributed second-stage delay for these passengers under these scenarios.

- $g(\mathbf{y}, \omega)$ is the total delay associated with passengers flying assigned second-stage itineraries. The delay from each second-stage itinerary is weighted by the probability of the scenarios under which the first-stage itinerary cannot be completed and the second-stage itinerary is assigned. Note that $g(\mathbf{y}, \omega)$ is 0 for a scenario ω where the length of delay allows the first-stage itineraries to be successful for all the passengers.

The formulation then takes the general form:

$$\min \sum_{\omega} \rho^{\omega} \big(f(\mathbf{x}, \omega) + g(\mathbf{y}, \omega) \big) \tag{1a}$$

$$\text{s.t. } \mathbf{x} \in F \tag{1b}$$

$$\mathbf{y} \in G(\mathbf{x}, \omega) \qquad \forall \omega \in \Omega \tag{1c}$$

The sets F and $G(\mathbf{x}, \omega)$ determine the constraints on $\mathbf{x}$ and $\mathbf{y}$. In particular, constraints (1b) say that every passenger on the delayed flight must be assigned a first-stage itinerary (possibly his or her original itinerary) and that no available flight capacities can be exceeded in the first stage. Constraints $(1c)$ provide an alternate itinerary to each passenger disrupted in a given scenario, those whose first-stage itinerary cannot be completed. Available flight capacity must again not be exceeded. Note that the second-stage decisions depend on the first-stage decisions, which determine which passengers are disrupted, as well as the realization of the delay length, which determines the set of viable second-stage itineraries.

3.3 First-Stage Formulation

For the first-stage part of the PRP model (the first term in $(1a)$ and constraints $(1b)$ in Section 3.2), we describe the sets, parameters, and variables used.

- P is the set of passengers on the delayed flight.
- F is the set of all flights of potential value for recovery options for passengers.
- I_p is the set of all possible itineraries for passenger $p \in P$.
- $F_i \in F$ is the set of all flights from itinerary $i \in I_p$ for $p \in P$.
- h_f is the available capacity of flight $f \in F$.
- c_{pi} is the delay for passenger $p \in P$ if reassigned to itinerary $i \in I_p$, relative to the original itinerary, and assuming that the new first-stage itinerary is not disrupted. In the objective function, this parameter is weighted by the probability of itinerary i not being disrupted.
- x_{pi} is a decision variable that takes the value 1 if passenger $p \in P$ is assigned to itinerary $i \in I_p$ and 0 otherwise.

The objective function for the first-stage problem, $f(\mathbf{x}, \omega)$ from Section 3.2, begins with $\sum_{p \in P} \sum_{i \in I_p} c_{pi} x_{pi}$. As described in Section 3.2, we subtract the cost for any passengers who will not reach their destinations with their first-stage itineraries. For all $p \in P$ and $i \in I_p$, let ϵ_i^ω be 1 if itinerary i will be disrupted in outcome $\omega \in \Omega$. Otherwise let it be 0. Then

$1 - \sum_{\omega \in \Omega} \epsilon_i^\omega \rho^\omega$ is the probability that itinerary i will not be disrupted. Thus, $f(\mathbf{x}, \omega)$ from the general form from Section 3.2 becomes $\sum_{p \in P} \sum_{i \in I_p} \Big(1 - \sum_{\omega \in \Omega} \epsilon_i^\omega \rho^\omega\Big) c_{pi} x_{pi}$.

The following make up constraints $(1b)$ from Section 3.2.

$$\sum_{i \in I_p} x_{pi} = 1 \qquad \forall p \in P \qquad (2a)$$

$$\sum_{p \in P} \sum_{i \in I_p : f \in F_i} x_{pi} \le h_f \qquad \forall f \in F \qquad (2b)$$

$$x_{pi} \in \{0, 1\} \qquad \forall p \in P,\ \forall i \in I_p \qquad (2c)$$

Constraint set $(2a)$ ensures that every passenger is assigned to exactly one itinerary. Set $(2b)$ says that we cannot assign more passengers to a flight than there are available seats, and set $(2c)$ makes all decision variables be binary.

3.4 Second-stage Formulation

The second-stage problem for each scenario ω consists of $g(\mathbf{y}, \omega)$ in the second term from the objective function $(1a)$, and the set of constraints $(1c)$ corresponding to ω from the general form in Section 3.2. Recall that $g(\mathbf{y}, \omega)$ from the objective function is the total delay associated with passengers flying assigned second-stage itineraries under scenario ω. The constraints provide an alternate itinerary to each passenger disrupted in a given scenario, those whose first-stage itinerary cannot be completed. Available flight capacity must not be exceeded.

We formulate the second-stage model for each ω as a minimum cost flow formulation by creating a separate end node for each passenger who is disrupted after the first stage under scenario ω. In addition, to enforce the final destinations of passengers, we strategically place arcs to these end nodes only where appropriate. With this formulation, we have the benefit of all variables being continuous, allowing us to use Benders Decomposition.

We describe the nodes, arcs, and notation for the network in each scenario ω, then provide an example. Note that the networks for the scenarios differ in only two ways, the number of arcs beginning at the start node and the demand at the end nodes, as described more fully below. First, N is the set of nodes in the network, and there are three types of nodes.

- A start node n_1 representing the original location, that has a supply equal to the total number of passengers disrupted under their new first-stage solutions and a given disruption scenario ω.
- A pair of flight nodes for each flight in the data set. Each of these nodes has a supply of 0.
- An end node for each passenger $p \in P$. This set is denoted N_P, and the end node for passenger $p \in P$ is denoted by $n_p \in N_P$. For any passenger who is disrupted after the first stage for the given first-stage decisions and disruption scenario ω, the corresponding end node has a demand of 1 in scenario ω.

Next, A^ω is the set of arcs in the network in outcome $\omega \in \Omega$. There are several types of arcs, all of which have infinite capacity and a cost of 0 unless otherwise specified.

- An arc from the start node to the first node of the pair associated with any flight that meets the following conditions: First, the departure city of the flight must be the same as the destination of the delayed flight. Second, the flight has to depart after the delayed flight has landed, plus time for passengers to change aircraft. Last, the flight must arrive before the end of the problem instance's time horizon. Note that the set of flights in the network varies by scenario ω. Also, in the second stage, we only consider passengers who flew on the delayed flight and missed a connection. Any other passengers originally on that flight have first-stage itineraries that will not be disrupted. This set of arcs beginning at the start node is denoted $B_1^\omega \subset A^\omega$.
- Arcs $A_F \subset A^\omega$ between two flight nodes in a pair with capacity equal to the available capacity of the relevant flight $f \in F$. For each $i \in I_p$ for all $p \in P$, $A_{pi} \subset A_F$ is the

set of arcs in the network that correspond to itinerary i for passenger p in the first stage. Also, for $a \in A_F$, u_a is the capacity on the flight corresponding to arc a after the first-stage decisions have been made.

- Arcs between the second of one pair of flight nodes and the first of another pair of flight nodes if it is possible to connect from the first flight to the second flight.

- Arcs $A_e \subset A^\omega$ going from the second flight node in a pair to an end node. For a pair of flight nodes, the second node in the pair connects to an end node corresponding to a passenger if the arrival location for the flight is the same as the passenger's desired destination. We assign a cost c_a to using arc $a \in A_e$, which represent landing in the desired location after taking the last flight in a passenger's itinerary. This cost is specific to each passenger, since passengers have different planned arrival times.

For the second-stage model, we need a few more pieces of notation.

- $B_n \subset A^\omega$ is the set of arcs that begin at any node $n \in N$ except the start node.

- $E_n^\omega \subset A^\omega$ is the set of arcs that end at node $n \in N$ in outcome $\omega \in \Omega$.

- s_n is the supply at node $n \in N$.

- ϵ_i^ω is 1 if first-stage itinerary $i \in I_p$ for some $p \in P$ will be disrupted in outcome $\omega \in \Omega$. The only situation where this can happen is when the first flight in itinerary i is the original flight, since all other flight times are known. Otherwise ϵ_i^ω is 0.

Last, the variables for the second-stage model are defined.

- y_a^ω is the number of units of flow (passengers) sent over arc $a \in A^\omega$.

An example is provided in figure 1. Note that above each arc between two corresponding flight nodes for flight $f \in F$ is u_a, which is the capacity of the arc between the pair of nodes corresponding to flight f.

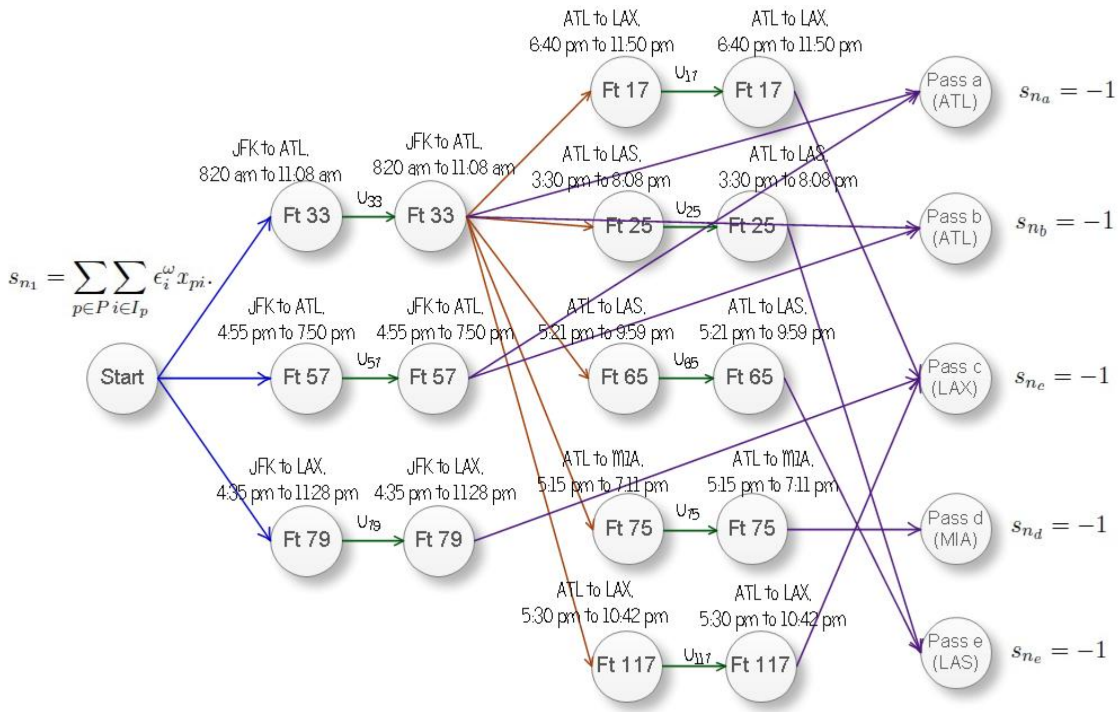

Figure 1: Minimum Cost Flow Network

The following constraints make up constraints $(1c)$ for a given $\omega \in \Omega$ from the general form in Section 3.2.

$$\sum_{a \in B_1^\omega} y_a^\omega = \sum_{p \in P} \sum_{i \in I_p} \epsilon_i^\omega x_{pi} \tag{3a}$$

$$\sum_{a \in E_{n_p}} y_a^\omega = \sum_{i \in I_p} \epsilon_i^\omega x_{pi} \qquad \forall p \in P \tag{3b}$$

$$\sum_{a \in B_n} y_a^\omega - \sum_{a \in E_n^\omega} y_a^\omega = 0 \qquad \forall n \in N \setminus (\{n_1\} \cup N_P) \tag{3c}$$

$$y_a^\omega \leq u_a \qquad \forall a \in A_F \tag{3d}$$

Constraints $(3a), (3b)$, and $(3c)$ ensure that all passengers disrupted with their first-stage itineraries under a given scenario ω are given second-stage itineraries that end at their desired destinations. Set $(3d)$ guarantees that the available capacities are not exceeded, taking into account first-stage solutions. Thus, $u_a = h_f - \sum_{p \in P} \sum_{i \in I_p : a \in A_{pi}} (1 - \epsilon_i^\omega) x_{pi}$ for all $a \in A_F$, where

h_f is the available capacity of the corresponding flight $f \in F$ before the first-stage decisions were made.

The objective function for the second-stage problem, $g(\mathbf{y}, \omega)$ from Section 3.2, is $\sum_{a \in A_e} c_a y_a^\omega$, since the only arcs with non-zero cost are those in A_e.

Note that the x variables are constants in the second-stage problem. Thus, this is a true minimum cost flow formulation. Constraints ($3a$), ($3b$), and ($3c$) are flow balance constraints, and ($3d$) are capacity constraints.

4 Solution Approach

For even fairly small data sets, it is not possible to solve the PRP model in reasonable run times using a direct implementation of branch-and-bound. We instead apply Benders Decomposition as introduced in Benders (1962), and seen in airline planning problems such as Cordeau et al. (2001), Mercier et al. (2005), Petersen et al. (2010), and Rosenberger et al. (2003). Branch and bound is still used, due to the integrality of the first-stage decision variables, since we must solve an integer program at each iteration of Benders Decomposition.

In this section, we first introduce and discuss Benders Decomposition. We explain how our model can be solved using it and then we present some computational results on the solutions to our model, varying the possible lengths of delay. In particular, the run time in solving our model in different situations is evaluated, and we specifically discuss how using Benders Decomposition affects the run time with different sizes of the data set. Last, we compare the delay experienced by the passengers from using our method versus reaccommodating passengers one-by-one after misconnections have occurred.

4.1 Benders Decomposition for the PRP Model

We consider the application of Benders Decomposition to the PRP model. Recall that the second-stage problem for each outcome $\omega \in \Omega$ is

$$
\begin{aligned}
\text{minimize} \quad & \sum_{a \in A_e} c_a y_a^\omega \\
\text{subject to} \quad & \sum_{a \in B_1^\omega} y_a^\omega = \sum_{p \in P} \sum_{i \in I_p} \epsilon_i^\omega x_{pi} \\
& \sum_{a \in E_{np}} y_a^\omega = \sum_{i \in I_p} \epsilon_i^\omega x_{pi} && \forall p \in P \\
& \sum_{a \in B_n} y_a^\omega - \sum_{a \in E_n^\omega} y_a^\omega = 0 && \forall n \in N \setminus (\{n_1\} \cup N_P) \\
& y_a^\omega \le u_a && \forall a \in A_F \\
& y_a^\omega \ge 0 && \forall a \in A^\omega
\end{aligned}
$$

To form the dual of the second-stage problem (the subproblem), we need the following notation for each $\omega \in \Omega$:

$$
\delta_a^\omega = \begin{cases} 1 & \text{arc}\, a\, \text{begins at the start node in outcome}\, \omega \\ 0 & \text{otherwise} \end{cases}
$$

$$
\zeta_{n,a}^\omega = \begin{cases} 1 & \text{arc}\, a\, \text{begins at flight node}\, n\, \text{in outcome}\, \omega \\ 0 & \text{otherwise} \end{cases}
$$

$$
\lambda_{n,a}^\omega = \begin{cases} 1 & \text{arc}\, a\, \text{ends at flight node}\, n\, \text{in outcome}\, \omega \\ 0 & \text{otherwise} \end{cases}
$$

$$
\sigma_{p,a}^\omega = \begin{cases} 1 & \text{arc}\, a\, \text{ends at end node for passenger}\, p\, \text{in outcome}\, \omega \\ 0 & \text{otherwise} \end{cases}
$$

$$
\kappa_{f,a} = \begin{cases} 1 & \text{arc}\, a\, \text{is between the two nodes for flight}\, f \\ 0 & \text{otherwise} \end{cases}
$$

Using this new notation and suppressing the ω for simplicity, our second-stage problem for each $\omega \in \Omega$ can be written

$$
\begin{aligned}
\text{minimize} \quad & \sum_{a \in A} c_a y_a \\
\text{subject to} \quad & \sum_{a \in A} \delta_a \cdot y_a - \sum_{p \in P} \sum_{i \in I_p} \epsilon_i x_{pi} = 0 \\
& \sum_{a \in A} (\zeta_{n,a} - \lambda_{n,a}) \cdot y_a = 0 && \forall n \in N \setminus (\{n_1\} \cup N_P) \\
& \sum_{a \in A} \sigma_{p,a} \cdot y_a - \sum_{i \in I_p} \epsilon_i x_{pi} = 0 && \forall p \in P \\
& \sum_{a \in A} \kappa_{f,a} \cdot y_a \leq h_f && \forall f \in F \\
& y_a \geq 0 && \forall a \in A^\omega
\end{aligned}
$$

Thus the dual of our second-stage problem for each $\omega \in \Omega$ is

$$
\begin{aligned}
\text{maximize} \quad & \Big(\sum_{p \in P} \sum_{i \in I_p} \epsilon_i x_{pi} \Big) \cdot \alpha + \sum_{p=1}^{|P|} \Big(\sum_{i \in I_p} \epsilon_i x_{pi} \Big) \cdot \pi_p + \sum_{f=1}^{|F|} h_f \cdot \mu_f \\
\text{subject to} \quad & \delta_a \cdot \alpha + \sum_{n=1}^{2|F|} (\zeta_{n,a} - \lambda_{n,a}) \cdot \beta_n + \sum_{p=1}^{|P|} \sigma_{p,a} \cdot \pi_p + \sum_{f=1}^{|F|} \kappa_{f,a} \cdot \mu_f \ \leq \ c_a, \quad \forall a \in A \\
& \boldsymbol{\mu} \ \leq \ 0
\end{aligned}
$$

Now that we have our dual subproblem, the original model can be written as

$$
\begin{aligned}
\text{minimize} \quad & \sum_{p \in P} \sum_{i \in I_p} \Big(1 - \sum_{\omega \in \Omega} \epsilon_i^\omega \rho^\omega \Big) c_{pi} x_{pi} + \sum_{\omega \in \Omega} \rho^\omega \theta^\omega \\
\text{subject to} \quad & \sum_{i \in I_p} x_{pi} = 1 && \forall p \in P \\
& \sum_{p \in P} \sum_{i \in I_p : f \in F_i} x_{pi} \leq h_f && \forall f \in F \\
& \theta^\omega = \max \Big\{ \Big(\sum_{p \in P} \sum_{i \in I_p} \epsilon_i^\omega x_{pi} \Big) \cdot \alpha^\omega + \sum_{p=1}^{|P|} \Big(\sum_{i \in I_p} \epsilon_i^\omega x_{pi} \Big) \cdot \pi_p^\omega + \sum_{f=1}^{|F|} h_f \cdot \mu_f^\omega \,| \\
& \qquad \delta_a^\omega \cdot \alpha^\omega + \sum_{n=1}^{2|F|} (\zeta_{n,a}^\omega - \lambda_{n,a}^\omega) \cdot \beta_n^\omega + \sum_{p=1}^{|P|} \sigma_{p,a}^\omega \cdot \pi_p^\omega + \sum_{f=1}^{|F|} \kappa_{f,a} \cdot \mu_f^\omega \ \leq \ c_a \\
& \qquad \forall a \in A^\omega, \ \boldsymbol{\mu} \ \leq \ 0 \Big\}
\end{aligned} \tag{4.2}
$$

We know that our dual subproblem is optimized at an extreme point, so the formulation (4.2), the master problem, becomes

$$
\begin{aligned}
\text{minimize} \quad & \sum_{p\in P}\sum_{i\in I_p}\Big(1-\sum_{\omega\in\Omega}\epsilon_i^\omega\rho^\omega\Big)c_{pi}x_{pi}+\sum_{\omega\in\Omega}\rho^\omega\theta^\omega && (4.3)\\
\text{subject to} \quad & \sum_{i\in I_p}x_{pi}=1 && \forall p\in P\\
& \sum_{p\in P}\sum_{i\in I_p:f\in F_i}x_{pi}\le h_f && \forall f\in F\\
& \theta^\omega \ge \Big(\sum_{p\in P}\sum_{i\in I_p}\epsilon_i^\omega x_{pi}\Big)\cdot\psi_j^\omega+\sum_{p=1}^{|P|}\Big(\sum_{i\in I_p}\epsilon_i^\omega x_{pi}\Big)\cdot\tau_{jp}^\omega+\sum_{f=1}^{|F|}h_f\cdot\nu_{jf}^\omega\\
& \qquad \forall\ \begin{bmatrix}\psi_j^\omega & \boldsymbol{\tau}_j^\omega & \boldsymbol{\nu}_j^\omega\end{bmatrix}\in Q^\omega,\ \omega\in\Omega
\end{aligned}
$$

where Q^ω for all $\omega\in\Omega$ is the set of extreme points of the polyhedron

$$
\Big\{\begin{bmatrix}\alpha^\omega & \boldsymbol{\beta}^\omega & \boldsymbol{\pi}^\omega & \boldsymbol{\mu}^\omega\end{bmatrix}\ \Big|\ \delta_a^\omega\cdot\alpha^\omega+\sum_{n=1}^{2|F|}(\zeta_{n,a}^\omega-\lambda_{n,a}^\omega)\cdot\beta_n^\omega+\sum_{p=1}^{|P|}\sigma_{p,a}^\omega\cdot\pi_p^\omega+\sum_{f=1}^{|F|}\kappa_{f,a}\cdot\mu_f^\omega\ \le\ c_a
$$
$$
\forall a\in A^\omega,\ \boldsymbol{\mu}^\omega\ \le\ \mathbf{0}\Big\}
$$

To avoid the need for feasibility constraints, we add an arc from each start node to each end passenger node. Flow on one of these arcs represents the passenger not arriving at the desired location. We assign a very high delay to these arcs, so that they are not chosen unless it is not possible to get all passengers to their destinations, and thus all instances have a feasible solution.

Next, we use delayed constraint generation to solve this model. We ignore the last set of constraints and solve the restricted master problem, using branch-and-bound to find integer solutions:

$$
\text{minimize} \quad \sum_{p\in P}\sum_{i\in I_p}\Big(1-\sum_{\omega\in\Omega}\epsilon_i^\omega\rho^\omega\Big)c_{pi}x_{pi}+\sum_{\omega\in\Omega}\rho^\omega\theta^\omega \qquad (4.4)
$$

$$
\begin{aligned}
\text{subject to} \quad & \sum_{i \in I_p} x_{pi} = 1 && \forall p \in P \\
& \sum_{p \in P} \sum_{i \in I_p : f \in F_i} x_{pi} \leq h_f && \forall f \in F
\end{aligned}
$$

From solving this, we get an optimal solution $(\mathbf{x}^\star, \boldsymbol{\theta}^\star)$. Next we check if $(\mathbf{x}^\star, \boldsymbol{\theta}^\star)$ satisfies the ignored constraints. If it does, then $(\mathbf{x}^\star, \boldsymbol{\theta}^\star)$ is the solution to model (4.3) with all the constraints included. In order to identify a violated constraint, we solve the dual subproblem (4.1) for each $\omega \in \Omega$, and add any violated constraints to the restricted master problem (4.4).

We re-solve this model, check for ignored constraints, and continue this process until no violated constraints are found. Then we have an optimal solution to (4.3). We can then use complementary slackness to find $\mathbf{y}^{1\star}, \mathbf{y}^{2\star}, ..., \mathbf{y}^{|\Omega|\star}$.

4.2 Computational Results

In this section, we present our computational results. First, we present our data sets. Next, we present results demonstrating the tractability of our approach and the PRP model, studying the run time in different situations. We compare solving the model directly as a MIP versus using Benders Decomposition, and show the benefits of using Benders Decomposition in many instances.

Finally, we analyze the improvement of solution quality based on the PRP Model. We contrast the PRP model with a method where passengers cannot replace the delayed flight, and any passengers that have missed connections after the delayed flight has landed are reaccommodated one-by-one, not necessarily in any particular order. We call this approach the "comparison method."

To simulate the comparison method, we use a greedy heuristic with the same possibilities and probabilities for the length of delay as in our model, but we force all passengers to keep the original flight, then reaccommodate them one-by-one after they land if necessary, giving each person the best itinerary still available. The order that we reaccommodate the passengers is randomly generated.

For the first six data sets, described in the next section, the model was implemented using C++ with CPLEX 12.1/Concert Technology as the underlying solver and run on an Intel Pentium(R) Dual-Core 2.10 GHz CPU 64-bit operating system. For the last three data sets, CPLEX 12.6.0 was utilized, along with an Intel(R) Core(TM) 2.40 GHz CPU 64-bit operating system.

4.2.1 Data

For our computational experiments, we consider nine data sets, described in Tables 1, 2, and 3. For each data set, we have 15 test instances. These have randomly generated data for two components, described below.

For all the data sets, there is a fixed set of 1144 flights from one airline from January 6, 2010 (http://www.transtats.bts.gov (n.d.)). In each data set, we allow the available capacity to range from 4 to 12 for all the non-delayed flights. According to http://www.transtats.bts.gov (n.d.), load factors in 2010, the year we took our flight data from, were roughly 82%. Our passenger data sets and capacity parameters were designed to mimic conservative estimates for available capacities.

In all the data sets, a specific flight from JFK to ATL, originally scheduled to depart at 8:20 am and arrive at 11:08 am, is delayed. The number of delay scenarios for this flight and the lengths of delay vary by the data set. Within each data set 1-7, the probabilities for all delay scenarios are equal. For the eighth and ninth data sets, the probabilities are shown in Table 3. Note that since this flight was scheduled to depart and arrive in the morning, there are more options for reaccommodating passengers than there would be for a flight later in the day.

In each data set, there are either 50, 100, or 200 passengers on this delayed flight. For the 15 test instances for each data set, we vary the final destination of each passenger on the delayed flight, by randomly selecting the destinations from the 15 or 30 chosen for the test instance. We distribute the passengers evenly over these locations, except that one third of passengers have only the original flight. For each test instance, we originally give each passenger the best or second best itinerary going to his or her final destination, based on the arrival times. The test instances also vary by having the order that the passengers who miss connections with their first-stage itineraries are reaccommodated when using the comparison method randomly generated.

Tables 1, 2, and 3 below summarize the different data sets, with the following pieces of information.

- Num. Pass.: The number of passengers on the delayed flight.
- Num. Final Dests.: The number of different destinations of the passengers on the delayed flight.
- Av. Caps.: The available capacity on all flights.
- Num. Delay Scens.: The number of possible delay lengths, which is the number of scenarios.
- Delays: The possible delay lengths in the different scenarios, in minutes.
- Probs: The probabilities of the scenarios (shown for data sets 7-9 only)

Table 1: Data Sets 1 to 3

Data Info	Set 1	Set 2	Set 3
Num. Pass.	50	100	100
Num. Final Dests.	15	15	30
Av. Caps.	5	5	4
Num. Delay Scens.	4	4	4
Delays	30, 90, 210, 300	30, 90, 210, 300	30, 90, 210, 300

Table 2: Data Sets 4 to 6

Data Info	Set 4	Set 5	Set 6
Num. Pass.	100	200	200
Num. Final Dests.	15	15	15
Av. Caps.	8	10	12
Num. Delay Scens.	16	16	40
Delays	30, 60,...,480	30, 60,...,480	10, 20,..., 400

Table 3: Data Sets 7 to 9

Data Info	Set 7	Set 8	Set 9
Num. Pass.	200		
Num. Final Dests.	15		
Av. Caps.	10		
Num. Delay Scens.	7		
Delays.	10, 15, 30, 45, 60, 90, 120		
Probs	$\frac{1}{7}, \ \forall\omega$	$0.1, 0.1, 0.1, 0.25, 0.2, 0.15, 0.1$	$0.1, 0.15, 0.2, 0.25, 0.2, 0.05, 0.05$

4.2.2 Run Time Results

For data sets 1-6, we solve our model in two different ways. The first way is to solve the complete mixed-integer model all at once, and the second is to use Benders Decomposition to break down the model as explained in Section 4.1. We present the average run time over 15 instances for each data set. We first describe the columns for the data given in Table 4.

- MIP - The total time, in seconds, it takes to create and solve the PRP Model by directly solving the mixed integer program.

- Bend. - The total time, in seconds, to create and solve the PRP Model using Benders Decomposition.

- B./MIP. - The ratio of the entry in column “Bend.” to the entry in column “MIP.”

Table 4: Average Run Times Over All Instances for Each Data Set (Seconds)

Data Set	MIP	Bend.	B./MIP
1	7.6	10.6	141%
2	7.0	19.6	282%
3	7.9	11.9	150%
4	28.0	36.6	131%
5	n/a	138.3	n/a
6	n/a	186.7	n/a

Note that memory was exceeded when trying to solve all instances of data sets 5 and 6 without using Benders Decomposition. In these data sets, there are a higher number of passengers and delay scenarios. From the results in Table 4, we draw several conclusions:

- The run times are relatively quick when Benders Decomposition is used, even for data set 6 with 200 passengers and 40 scenarios. For instances 1 to 3, the run times are on the order of several seconds.

- For small models, both approaches (with and without Benders Decomposition) are very effective with quick run times. For the smaller models, data set 1 through 4, using Benders Decomposition does not improve the run time.

- As noted, data sets 5 and 6 cannot be solved without Benders Decomposition.

Note that, while doing one iteration of Benders Decomposition, we have to solve $|\Omega|$ subproblems. Because the subproblems are independent of each other, they can be solved in parallel. In that case, the time to solve all of them is the highest time to solve any one of them. We compare the run times using this perspective of parallelization of the algorithm as well. We first describe the columns for the data given in Table 5.

- MIP - The total time, in seconds, it takes to create and solve the PRP Model without Benders decomposition.

- Bend., p. - The total time, in seconds, to create and solve the PRP Model using Benders Decomposition if we parallelize the subproblems. Note that to find this quantity, we calculate the time to solve each of the $|\Omega|$ subproblems, and take the maximum.

- B.,p./MIP - The ratio of the entry in column "Bend., p." to the entry in the column "MIP."

Table 5: Average Run Times Over All Instances with Parallelization (Seconds)

Data Set	MIP	Bend., p.	B.,p./MIP
1	7.6	6.3	83%
2	7.0	10.3	147%
3	7.9	7.7	97%
4	28.0	14.4	52%
5	n/a	62.7	n/a
6	n/a	87.1	n/a

Note that, when parallelized, a Benders Decomposition approach becomes faster than the direct MIP in most of the smaller test instances as well.

For data sets 7-9, we present the total time, in seconds, to create and solve the PRP Model using Benders Decomposition if we parallelize the subproblems, as in Table 5.

Table 6: Average Run Times Over All Instances with Parallelization (Seconds)

Data Set	Bend., p.
7	18.8
8	18.4
9	26.1

For some flight delays, airlines may want to consider many different possibilities for the outcome of the length of the delay. To accomplish this goal, more subproblems can be included while solving the PRP model with the Benders Decomposition approach. Since the algorithm can be parallelized, adding more subproblems has minimal, if any, effect on the run time in solving the model. Thus, the PRP model is adaptable to any number of possible delay lengths and would continue to have a reasonable run time.

4.2.3 Results on Delays

In this section, we present results that show the difference in passenger delays when the PRP Model is used versus the comparison method. Recall that in the comparison method, disrupted passengers from the delayed flight are reaccommodated one-by-one after missing a connection, and that the test instances for a specific data set vary by the destinations of the passengers and the order they are reaccommodated with the comparison method.

To see this difference, we look at the number of passengers experiencing a delay, the number not reaching their destination by the end time, the average delay experienced over all passengers on the delayed flight, and the highest delay for any one passenger. Note that we have assigned a delay of 4320 minutes (three days) for not being able to be reaccommodated. The metrics, given for both methods, are now described further.

- "AEND" - the average expected number of passengers delayed over all scenarios and all instances.
- "AENNR" - the average expected number of passengers not reaccommodated, i.e., not reaching their destinations by the end time, over all scenarios and all instances. Note that passengers not reaccommodated are considered to be delayed.
- "AED" - the average expected delay over all passengers, scenarios, and instances, in minutes.
- "Delay%" - The percentage that the average expected delay for the PRP Model is of the average expected delay for the comparison method.
- "AEHD" - the average expected value of the highest passenger delay over all scenarios and instances, in minutes.

Table 7: Average Values over all Instances and Scenarios for Each Data Set

Set	Method	AEND	AENNR	AED	Delay%	AEHD
1	PRP	25.4	0.0	85.7	27%	255.2
1	Comp.	29.3	2.7	340.2		2231.2
2	PRP	51.8	0.05	96.5	18%	513.6
2	Comp.	57.3	10.6	558.9		2706.7
3	PRP	53.6	0.0	92.7	27%	371.0
3	Comp.	58.4	5.5	348.7		2683.3
4	PRP	63.8	0.0	159.4	17%	390.1
4	Comp.	72.0	18.7	973.7		3681.7
5	PRP	126.0	0.03	157.9	13%	645.1
5	Comp.	136.3	49.4	1199.5		3739.5
6	PRP	116.2	0.0	123.9	12%	330.6
6	Comp.	134.0	42.3	1032.7		3758.9
7	PRP	83.2	0.2	35.2	26%	537.2
7	Comp.	87.4	8.6	222.3		2334.0
8	PRP	83.6	0.1	34.8	31%	426.0
8	Comp.	86.3	5.9	164.3		2520.8
9	PRP	80.3	0.03	27.8	23%	254.2
9	Comp.	83.4	5.9	157.5		2561.3

In Table 8, we present the same information, but specifically for the scenario with the highest delay.

Table 8: Average Values over all Instances and the Worst Scenarios for Each Data Set

Set	Method	AND	ANNR	AC	Delay%	AHD
1	PRP	31.1	0.0	164.4	21%	379.8
1	Comp.	41.5	7.2	826.1		4320
2	PRP	63.5	0.07	181.3	16%	444.6
2	Comp.	77.5	23.5	1182.9		4320
3	PRP	66.5	0.0	178.4	22%	472.9
3	Comp.	80.3	14.4	833.7		4320
4	PRP	68.9	0.0	276.7	13%	523.1
4	Comp.	91.9	42.8	2083.6		4320
5	PRP	136.5	0.1	273.3	12%	780.3
5	Comp.	176.4	94.3	2243.9		4320
6	PRP	134.5	0.0	241.1	12%	478.7
6	Comp.	174.5	83.9	1987		4320
7	PRP	110.3	0.5	91.2	22%	1640.1
7	Comp.	118.7	25.3	624.6		4040.0
8	PRP	103.8	0.3	80.8	23%	1177.3
8	Comp.	109.3	17.4	449.9		4320
9	PRP	106.2	0.3	80.9	16%	872.2
9	Comp.	115.1	21.5	541.9		4320

4.2.4 Evaluation and Comparison of Methods

In Section 4.2.3, we presented the results of the test instances for each of our data sets, showing differences between the PRP Model and the comparison method. Recall that to evaluate the potential impact of the PRP Model, we compared this to the case where there was no preemptive rerouting, but then disrupted passengers were reassigned one-by-one after missing a connection.

The computational results demonstrate that our proposed method yields better passenger outcomes in our instances. We now make some conclusions regarding the differences in delays, based on the results given in Section 4.2.3. There is a significant range in the average delays experienced, which is not surprising. Some destinations are much harder to get to than others, based on which markets have many flights going to and from them. As expected, our average expected delay is significantly less than the comparison average expected delay.

We have done some statistical analysis on the data for all the data sets. First is a hypothesis test comparing our average expected delay and the comparison average expected delay. Letting μ_c and μ_n be the average expected delays for the comparison method and our model, respectively, our null hypothesis is $H_0 : \mu_c - \mu_n = 0$. Thus our null hypothesis says that there is no difference between μ_c and μ_n. The alternative hypothesis is $H_1 : \mu_c - \mu_n > 0$. Note that this is a one-sided alternative hypothesis. Since our method has the added benefit of allowing passengers to get off the delayed flight and considering all passengers together, the delay from our method will never be more than the delay from the comparison method. Any solution that can be achieved using the comparison method can also be achieved using our method.

The t-statistic is $t = \frac{\bar{y}}{s/\sqrt{n}}$, where $\bar{y}$ is the average value of $\mu_c - \mu_n$ over the instances, s is the sample standard deviation of the value of $\mu_c - \mu_n$ in the instances, and n is the number of instances. With the values obtained, for all data sets, the null hypothesis can be rejected even at the 0.1% level.

It is notable that the highest expected delay experienced for any one passenger decreases significantly in all but two instances in all data sets. These are the eleventh instance for data set 2, where the highest expected delay for our method and the comparison method are 3276.5 and 3247.5, respectively, a difference of 29 minutes, and the ninth instance for data set 7, where the highest expected delays from the two methods are equal. Recall that 4320 minutes (three days) is the cost of not being reacommodated. Also, in the scenario with the greatest delay for each data set, the highest delay for any one passenger using the PRP model is always less than or equal to that for the comparison method, but is usually much less.

The number of passengers delayed and the number delayed over 30 minutes decreases as well while using the PRP Model. It is notable that in every instance, all of the values both in expectation and in the scenario with the longest delay are better using the PRP Model. The only exception is a couple of instances, where the expected number of people delayed is slightly higher or the same using the PRP Model, since the average delay length is minimized instead, but the other values are significantly better with the PRP Model.

One question is how great the impact is of the preemptive part of the PRP model in this reduction in the delays passengers experience using the PRP model. In Table 9, we present for data sets 1-6, the average percentage of passengers over all instances who had the delayed flight in their original itineraries and were given new itineraries in the first-stage problem that did not include the delayed flight using the PRP model. We see that moving passengers off the delayed flight preemptively affects many passengers.

Table 9: Average Percentage of Passengers Moved Off Delayed Flight for Each Data Set

Set 1	Set 2	Set 3	Set 4	Set 5	Set 6
25.4%	17.1%	17.4%	37.1%	35.0%	26.1%

Note that in our instances, we assume a linear relationship between the delay a passenger experiences and the cost that should be assigned for that delay. For example, we assume that a two hour delay should be assigned twice the cost of a one hour delay. In actuality, a passenger may miss a meeting with a two hour delay but not with a one hour delay, so the two hour delay is much more than twice as costly. Also, it may be more accurate to assign a higher delay to one passenger being delayed two hours than two passengers waiting one hour each. We can generalize the model to a nonlinear relationship between the length of delay and cost assigned to the delay by simply changing the delay cost c_{pi} for an itinerary $i \in I_p$ for passenger $p \in P$ or changing the delay on an arc $a \in A_e$.

5 Conclusions and Future Research

Our contributions are a new approach to passenger reaccommodation, specifically in proactively handling passenger delays before misconnections occur and using known probabilities; the PRP Model where the second-stage problem is modeled as a minimum cost flow problem; and a Benders Decomposition-based approach to finding tractable solutions. Our strong computational and run time results show the framework and model's effectiveness for reducing the length of passenger delays.

In practice, there is often not only one delayed flight, but many different delays and cancellations all affecting each other during daily operations for an airline. For example, if there is a snowstorm or other bad weather, many flights may be affected. In addition, a delay in one flight, even as a result of a mechanical problem, could cause downstream delays in other flights. Thus concurrent delays are possible.

The PRP model lays the groundwork for studying the more realistic problem in which multiple flights may experience concurrent delays. In the future, we plan to expand our research to address multiple simultaneous delays. We hope to use a formulation that employs the idea of a "portfolio" of flights for each passenger. This would determine what flight the passenger should take from the current location, and also any flights to take in each possible disruption scenario based on the delayed flights' set of possible delay outcomes.

Acknowledgments

We would like to thank Divakar Viswanath from the Department of Mathematics at the University of Michigan. Dr. Viswanath was especially helpful during the implementation of the PRP model and also gave valuable feedback on several sections of the paper.

References

Abdelghany, A., Ekollu, G., Narasimhan, R. and Abdelghany, K. (2004). A proactive crew recovery decision support tool for commercial airlines during irregular operations, *Annals of Operations Research* **127**(1): 309–331.

AhmadBeygi, S., Cohn, A. and Lapp, M. (2010). Decreasing airline delay propagation by re-allocating scheduled slack, *IIE Transactions* **42**(7): 478–489.

Barnhart, C., Kniker, T. and Lohatepanont, M. (2002). Itinerary-based airline fleet assignment, *Transportation Science* **36**(2): 199.

Belobaba, P., Odoni, A. and Barnhart, C. (2009). *The global airline industry*, Wiley.

Benders, J. F. (1962). Partitioning procedures for solving mixed-variables programming problems, *Numerische mathematik* **4**(1): 238–252.

Bratu, S. and Barnhart, C. (2006). Flight operations recovery: New approaches considering passenger recovery, *Journal of Scheduling* **9**(3): 279–298.

Cordeau, J., Stojkovic, G., Soumis, F. and Desrosiers, J. (2001). Benders decomposition for simultaneous aircraft routing and crew scheduling, *Transportation Science* **35**(4): 375.

Eggenberg, N., Salani, M. and Bierlaire, M. (2010). Constraint-specific recovery network for solving airline recovery problems, *Computers & Operations Research* **37**(6): 1014–1026.

http://www.transtats.bts.gov (n.d.). Bureau of transportation statistics.

Jacobs, T., Smith, B. and Johnson, E. (2008). Incorporating network flow effects into the airline fleet assignment process, *Transportation Science* **42**(4): 514–529.

Jafari, N. and Hessameddin Zegordi, S. (2010). Simultaneous recovery model for aircraft and passengers, *Journal of the Franklin Institute* .

Kohl, N., Larsen, A., Larsen, J., Ross, A. and Tiourine, S. (2007). Airline disruption management–Perspectives, experiences and outlook, *Journal of Air Transport Management* **13**(3): 149–162.

Lapp, M., AhmadBeygi, S., Cohn, A. and Tsimhoni, O. (2008). A recursion-based approach to simulating airline schedule robustness, *Proceedings of the 40th Conference on Winter Simulation*, Winter Simulation Conference, pp. 2661–2667.

Lettovsky, L. (1997). Airline operations recovery: An optimization approach.

Lettovsky, L., Johnson, E. and Nemhauser, G. (2000). Airline crew recovery, *Transportation Science* **34**(4): 337.

Mercier, A., Cordeau, J. and Soumis, F. (2005). A computational study of benders decomposition for the integrated aircraft routing and crew scheduling problem, *Computers & Operations Research* **32**(6): 1451–1476.

Petersen, J., Solveling, G., Johnson, E., Clarke, J. and Shebalov, S. (2010). An Optimization Approach to Airline Integrated Recovery.

Rexing, B. (1997). *Airline fleet assignment with time windows*, PhD thesis, Massachusetts Institute of Technology.

Rexing, B., Barnhart, C., Kniker, T., Jarrah, A. and Krishnamurthy, N. (2000). Airline fleet assignment with time windows, *Transportation Science* **34**(1): 1–20.

Rosenberger, J., Johnson, E. and Nemhauser, G. (2003). Rerouting aircraft for airline recovery, *Transportation Science* **37**(4): 408.

Rosenberger, J., Schaefer, A., Goldsman, D., Johnson, E., Kleywegt, A. and Nemhauser, G. (2002). A stochastic model of airline operations, *Transportation science* **36**(4): 357–377.

Yen, J. and Birge, J. (2006). A stochastic programming approach to the airline crew scheduling problem, *Transportation Science* **40**(1): 3–14.