Integrated Disruption Management and Flight Planning to Trade off Delays and Fuel Burn

Lavanya Marla

iLab, Heinz College, Carnegie Mellon University, lavanyamarla@cmu.edu

Bo Vaaben

Jeppesen Commercial and Military Aviation, bo.vaaben@jeppesen.com
Department of Management Engineering, Technical University of Denmark, bova@man.dtu.dk

Cynthia Barnhart

Massachusetts Institute of Technology, cbarnhar@mit.edu

In this paper we present a novel approach addressing airline delays and recovery. Airline schedule recovery involves making decisions during operations to minimize additional operating costs while getting back on schedule as quickly as possible. The mechanisms used include aircraft swaps, flight cancelations, crew swaps, reserve crews and passenger rebookings. In this context, we introduce another mechanism, namely *flight planning*, which enables flight speed changes. Flight planning is the process of determining flight plan(s) specifying the route of a flight, its speed and its associated fuel burn. Our key idea in integrating flight planning and disruption management is to adjust the speeds of flights during operations, trading off flying time and fuel burn, and combining with existing mechanisms such as flight holds; all with the goal of striking the right balance of fuel costs and passenger-related delay costs incurred by the airline. We present models for integrated aircraft and passenger recovery with flight planning, both exact and approximate. From computational experiments on data provided by a European airline, we estimate approximately that reductions in passenger disruptions on the order of 66-83%, accompanied by small increases in fuel burn of 0.152 - 0.155% and total cost savings of 5.7 - 5.9% for the airline, may be achieved using our approach. We discuss the relative benefits of two mechanisms studied - specifically, flight speed changes and intentionally holding flight departures, and show significant synergies in applying these mechanisms. The results, compared to recovery without integrated flight planning, are increased swap possibilities during recovery, decreased numbers of flight cancelations, and fewer disruptions to passengers.

Key words: airline schedule recovery, flight planning, enhanced disruption management

1. Introduction

Inherent uncertainty in airline operations makes delays and disruptions inevitable. Because the airline system operates as a closely interconnected network, it is subject to 'network effects', that is, a disruption in one place can quickly propagate to multiple other parts of the network. Therefore, managing these delays as they arise is crucial. *Disruption management* is the process by which, on the day of operation, when a disruption occurs, airlines try to bring operations back on schedule

1

as quickly as possible, while incurring minimal costs. Measures such as flight cancelations, flight holds, aircraft swaps, crew swaps, reserve crew and passenger reaccommodation are used as part of the disruption management process. In this work, we integrate *disruption management* and *flight planning*. *Flight planning* is the process of determining, at the pre-departure stage of each flight, its three-dimensional trajectory, involving its path, altitude(s), speed and fuel burn as the aircraft flies from its origin to its destination. Our goal is to reduce flight delays and disruptions to passengers using disruption management combined with flight planning, to achieve the appropriate trade-off of passenger service with fuel burn and additional operating costs incurred during recovery. To our knowledge, this is the first work that integrates these aspects of airline operations.

The ability to change a flight's speed directly impacts its block (or flying) time, and thus, its arrival time; which in turn can impact network connectivity of the flight's aircraft, crew and passengers to downstream flights. Therefore, through changes to block times, to trade-off the costs of changing a flight's arrival time (including, for example, network connectivity costs capturing the costs associated with resulting delays and disruptions to the flight's aircraft, crews and passengers) with the change in fuel burn costs (associated with the flight's block time adjustment).

To illustrate our *integrated disruption management and flight planning* approach, consider, for example, a flight experiencing a departure delay at its origin. The choices are to: (1) operate the flight at increased speeds (and increased fuel burn), and employ techniques such as aircraft swaps and flight cancelations as necessary, to absorb delays at the flight destination and to decrease costs associated with passenger delays and misconnections; or (2) reduce the flight's speed using flight planning to decrease fuel burn and emissions if connectivity is unaffected. Our overarching goal is to decrease costs incurred during airline operations by identifying the operational trade-offs between (i) aircraft and passenger delay costs; and (ii) fuel burn costs.

1.1. Disruption Management

During operations, operational recovery procedures of dynamic scheduling, routing and disruption management vary among carriers. The first priority for most airlines facing disrupted operations is to bring operations back to the plan. For this, operations controllers re-assign the resources of the airline in order to minimize the costs associated with the disruption. Three types of decisions are made: (i) whether or not to cancel a flight, (ii) what the rescheduled departure time is of flights that are to be operated, and (iii) which aircraft and crew is assigned to each operated flight. Typically, following aircraft and crew recovery, passenger recovery and re-accommodation is performed.

For an in-depth study of airline recovery, we recommend the reader to Barnhart (2009), Yu and Qi (2004), Barnhart et al. (2006), Kohl et al. (2007) and Clausen et al. (2010). Other relevant studies are Dienst (2010) and Thengvall et al. (2001).

1.2. Flight Planning

A flight plan is a document prepared by an operator (usually an airline) indicating the movement of the concerned aircraft in time and space, from its origin to its destination. The flight plan specifies the route (ground track) of the aircraft, its profile (altitudes along the route), its speed (which varies along the route) and the fuel burned in operating the flight plan. For an example of a flight plan, we refer the reader to Altus (2007).

The relationship between fuel burn and flying time (and consequently, block time) for a given flight leg is highly non-linear. Figure 1 illustrates the relationship between flying time and fuel burn for a long-haul flight. The flexibility in speed changes is the highest for long-haul flights and least for short-haul flights. Each point on this curve represents a specific flight's flight plan with the associated flying time and fuel burn.
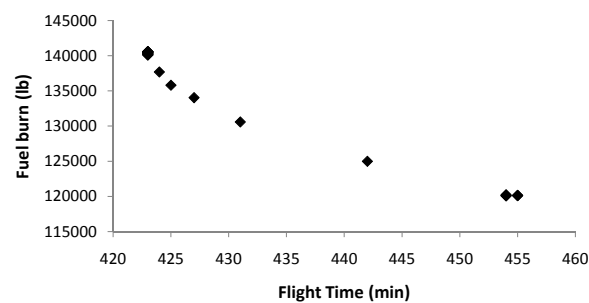

Figure 1   Relationship between flight time and fuel burn

1.3. The Problem

We briefly describe the problem setting in this section. We consider scenarios in which a flight is delayed at its origin due to a disruption in the network. Our decision time frame is from one hour to one half-hour prior to flight departure, when we know the expected departure time of the flight and are in a position to select the flight plan and satisfy the necessary fueling requirements related to the choice of flight plan. We consider disruption management techniques that combine flight planning with aircraft swaps, flight cancelations and passenger recovery. Through this process, we trade-off network connectivity costs and delay costs associated with flight arrival times, with the fuel costs associated with flight speed changes. The effect is to re-allocate slack in block and ground times by: (i) increasing aircraft speed to reduce block times and add ground time at the destination, thereby preserving connections; (ii) decreasing aircraft speed to increase block time and save on fuel costs if fuel costs dominate airline delay costs (especially those related to passengers); and (iii) intentionally delaying (or holding) downstream flights to preserve passenger connections, without increasing the speed of the arriving flight and incurring increased fuel costs.

1.4.   Contributions

The contributions of our research are as follows. First, we introduce an *enhanced disruption management* tool with integrated flight planning, and provide exact and approximate optimization models that combine flight planning with traditional disruption management models. In particular, we focus on two aspects of flight planning, speed changes and flight departure holding, and trade-off fuel costs and passenger delay costs. Our approach represents an integration of two aspects of airline operations hitherto studied separately, namely, disruption management and flight planning.

Second, through dialogue with multiple airlines, we provide a synopsis of the current state-of-the-practice with regards to flight planning approaches. We also discuss the current practices of flight planning and disruption management. We identify opportunities for improving disruption management through integration with flight planning and show the need for optimization-based decision support.

Third, we evaluate our approach on scenarios based on data from an international airline. Our experiments focus on hub operations and opportunities for improved trade-offs between passenger costs and fuel costs, with the goal of minimizing total realized costs. Based on our assumptions, we estimate approximately that in comparison with conventional disruption management, our integrated flight planning and disruption management strategy could result in decreases in passenger misconnections of about 66-83%, decreases in passenger-related delay costs for the airline of 60-73%, increases in fuel costs of 0.152-0.155%, and total cost savings of 5.7 - 5.9% for the airline under consideration. Additionally, passenger delay costs over the two-month period of our experiments are estimated approximately to decrease by $17.5-17.9M. By demonstrating the dynamic nature of the trade-off frontier between passenger costs and fuel burn costs and discussing this trade-off for different disruption scenarios, we make the case for dynamically selecting aircraft speeds during operations. We also discuss the relative benefits of the two types of mechanisms studied - that of flight speed changes and that of holding flight departures - and show significant synergies in applying the two mechanisms simultaneously.

1.5.   Organization of the paper

In §1.1, and §1.2, we presented an overview of disruption management and flight planning, respectively. In §2 we introduce some terms relevant to flight planning and a summary of the practice of operational flight speed changes via dialogue with six airlines. In §3 we illustrate, with an example, opportunities for integrating flight planning and disruption management to minimize costs; and indicate shortcomings in the state-of-the-practice that motivate our *integrated flight planning and disruption management* approach. We then present a diagrammatic description of our approach. In

§4, we present our modeling architecture to integrate flight planning with disruption management. Our models provide a way to trade-off passenger delay costs and fuel burn costs, and minimize total realized costs. We provide models in §5 that capture passenger connectivity costs exactly and approximately, thereby facilitating solution. We describe our experimental setup in §6. In §7, we present our results and compare them with the current state-of the practice to estimate cost savings to the airline under consideration.

2. Flight Planning: Current Practice

As illustrated in Figure 1, each point on the flight time-fuel burn curve for a given flight represents a flight plan from the origin to the destination of that flight. A flight plan on this curve may be identified based on multiple metrics, among which the two most common ones are: (i) a fixed flight cruise speed, and (ii) *cost-index* (CI). Flight cruise speed, as its name suggests, identifies the flight plan based on how fast the aircraft flies, and its associated flying time. It does not include a notion of fuel burn in its specification. CI, on the other hand, is a measure that has been introduced to capture explicitly both flight time (speed) and fuel burn in its definition; as detailed in the following subsection.

2.1. Cost Index (CI)-based flight planning

Cost Index (CI) is an assumed ratio of the time-related costs of a flight divided by the fuel cost; that is, it is the ratio of cost per unit time divided by the cost per mass unit fuel. Time-related costs are defined as those that are related to (i) the duration of the flight, examples include aircraft maintenance costs per minute and crew duty costs per minute; and (ii) the arrival time of the flight, examples include aircraft connectivity, crew connectivity, and passenger connection and delay costs per minute. CI is expressed in units of 100lb/hr (Boeing) or kg/min(Airbus) and can be interpreted physically as the amount of additional fuel worth burning (relative to the minimum fuel burn to operate the flight) to save one unit of time. CI thus captures within it a notion that time-related costs and fuel burn costs can be balanced. The use of CI is now standard practice in the industry, and is used as a rule-of-thumb, capturing the notion that associated with flight speed changes are both fuel impacts and network connectivity impacts.

Typically, an airline selects the 'right' CI value at which to build and operate its schedule by analyzing its historical operations, and computing the ratio of the total realized cost of fuel and the total realized cost of time-related effects (delays, connectivity, etc.). This can be done at a network, fleet or market level, resulting in 'network CI', or 'fleet CI', or 'market CI'. Airlines typically create their flight schedules such that flights are assumed to operate at the historically derived CI value

(referred to as the 'normal' CI) and its associated speed. The speed associated with the 'normal CI' is the speed for which the fuel burn rate equals the CI value. To the estimated flying time for the selected speed, additional time is added for taxiing, transiting, delays, etc. to finalize the block time, and schedule for the flight. Compared to flying by simply determining a speed, CI is a more balanced measure that is meant to account for delay- and fuel-related costs.

A CI value of zero means that relative to fuel costs, time-related costs are zero; or the additional fuel worth burning to save one unit of time relative to the minimum fuel burn speed is zero. In this case, the aircraft should be operated at its most fuel-efficient cruise speed, called the *maximum range cruise speed* (and minimum fuel burn speed). When operating at a high CI, the value of time is greater than the associated fuel burn cost, and to minimize the sum of fuel and connectivity costs, the aircraft is sped up, incurring higher fuel costs and lower delay costs.

2.2. State-of-the-practice at airlines

In this section, we discuss the current state-of-the-practice involving operational flight planning, for six international carriers.

In practice, computing the CI values using historical data (as described in §2.1) is time-consuming, costly, and requires the use of dedicated software (Altus 2010). As a result, typically a single 'average' CI value (equal to the 'normal' CI value) is used to determine flight plans and the speeds at which to operate flights. Operationally, airlines also specify a *range* of CI values that serve as the operating bounds on a given flight. The dispatcher or pilot is allowed, at his or her discretion, to speed up or slow down within this range in the event of schedule disturbances. (The max CI in the range does not mean that further speed up is not physically possible, instead it is the allowable upper limit at which the flight can be operated at the pilot's discretion.) The min CI value in the range is 0, that is, the minimum fuel burn speed. The max CI value in the range is typically set as a percentage cap on excess fuel burnt beyond the 'normal' CI, which can differ from carrier to carrier. The max CI value is at least limited by the fuel tankering policies of the airline, which do not allow speed up to an extent that requires the use of emergency fuel (this should occur only in emergency situations.) The max CI value can also be more conservatively set, reflecting the fact that the marginal cost of fuel burn per minute of flying time saved is increasing. Yet another consideration in setting the maximum CI value is the objective of airline management to prevent pilots from 'flying too fast' to reach their destinations early and cut their work day short, without regard to the high fuel costs incurred in the process. These guidelines result in pilots or dispatchers filing a faster flight plan prior to departure (at higher CI) if delayed at departure and a slower flight plan if departing early. Note, again, that this is with the intention of minimizing the sum

of fuel and time-related costs, as operating at a higher CI (speeding up) means that time-related costs dominate and operating at a lower CI means fuel costs dominate.

Pilots are also given the latitude, if tailwinds are encountered or if the aircraft has an early start, to operate the flight at a CI value lower than the 'normal CI' value. And, in cases of headwinds or late starts, pilots may adjust the flight speed during the flight to operate at higher CI values within the range. Typically, these guidelines are issued with the caution that speeding up will consume excess fuel, and such decisions should be taken judiciously.

A trend that has been observed in the industry during the recent fuel price spike in 2007 is increased use of speed changes during flight. Operational speed changes were used as a mechanism to save on fuel costs by utilizing slack in the flight schedule. Associated Press articles (Associated Press May 1, 2008) and (Associated Press May 2, 2008) reported that airlines slowed down flights, resulting in longer flying times but lower fuel burn. As a result, airlines reported savings of about $20 million in one year.

The prevalence of CI as a measure for choosing flight speeds and plans indicates that airlines give significant consideration to the trade-off between time-related connectivity costs and fuel costs. Current practices, however, have some shortcomings. A major issue, one we address in this work, is that the CI values and ranges do not capture the dynamics of operations and thus do not model the true time-fuel tradeoff. To illustrate this point, consider for example, that the 'right' choice of aircraft speed can differ for the same flight on different days, based on the network state, and aircraft and passenger connectivity of that flight on that day. We will demonstrate this further using an example in §3. No airlines, to our knowledge, make flight speed decisions that optimize the trade-off in passenger delay costs and fuel costs using *current* flight *network* information, taking into account downstream impacts involving flight and passenger misconnections. In the following sections, we describe how we enhance and extend current practices to capture these dynamics and network effects.

3. Our Integrated Flight Planning and Disruption Management Approach

Figure 2 provides a schematic of our basic concept. Consider a flight a into hub H, delayed by Δ at departure. If the aircraft flies at the scheduled speed, flight a reaches H Δ time units later than scheduled. This decreases connecting time available to passengers at H by Δ time units, and results in disruptions to passengers with connection time consequently reduced to less than the minimum connecting time $MinCT$. Using the following mechanisms, however, passenger connections may be preserved.

Flight speed changes: By changing the speed at which flight a is operated, block time can be decreased and ground time at H increased or vice versa. Figure 2(ii) shows how using alternate flight plans that operate at different speeds can create different amounts of slack in the schedule, with faster speeds on a allowing passengers adequate time to connect to flights b, c, d, and e.

Flight departure re-timing: In Figure 2(iii), we illustrate another strategy in which speed change decisions are complemented with flight holding decisions. In the example, it might be more cost efficient to speed up a to a lesser extent than is necessary to preserve passenger connections to b, c and d, and then to hold the departures of flights b, c, d to allow downstream connections from a.

(i) Original flight plan for flight *a*

(ii) Alternative flight plans for flight *a*

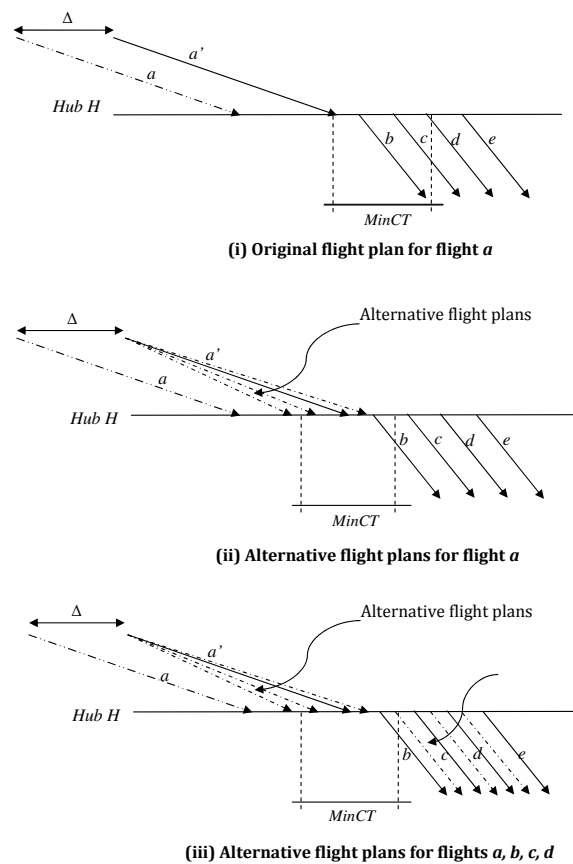(iii) Alternative flight plans for flights *a, b, c, d*

Figure 2 **Flexibility provided in disruption management by choosing alternate flight plans**

3.1. Example

We illustrate the advantage provided by optimizing flight speeds in the disrupted scenario shown in Figure 2. We evaluate, for each flight speed possible, the fuel cost and the passenger-related delay costs to the airline. To do this, we use the tools described in §3.1.1 and §3.1.2.

3.1.1.   Flight Planning Engine Flight plans used in our experiments are generated using JetPlan (Jeppesen 2010), a flight planning tool developed by Jeppesen Commercial and Military Aviation. Jeppesen's flight planning engine uses information about each flight, its planned (or current) schedule, airways, weather patterns, possible aircraft and engine configurations, and payload during the day of interest. It then generates flight plans corresponding to different speeds and travel times for each flight. The flight plan generator takes into account the fuel burn due to the payload consisting of cargo, passengers, luggage hold, and fuel weight. Included in fuel are contingency and reserve fuel.

3.1.2.   Passenger Delay Evaluation Module For each possible choice of flight plan and the schedule associated with that choice, we evaluate the impacts on passengers using an airline disruption management simulator (Davis et al. 2002) (Vaaben 2009). The purpose of this simulator is to compute the estimated true *realized* passenger delay costs of a set of delayed flights and the corresponding recovery actions on the day of operations. A passenger is defined to be *disrupted* if they cannot take their originally planned itinerary due to cancelations or misconnections. The simulator performs passenger re-accommodation for disrupted passengers by solving the passenger recovery problem with the actual cost values experienced by the airline. Summed with the delay costs experienced by passengers on delayed (but not disrupted) flights, this provides an estimate of the true passenger-related delay cost *to the airline.* These are computed using the delay cost specified in §6.3, which include passenger-delay related costs to the airline, hotel, meal reimbursements and goodwill costs.

3.1.3.   Results Table 1 shows the changes in (i) fuel costs of flight a (taken from the actual operation of a major European airline), and (ii) the corresponding realized passenger-related delay costs to the airline; by operating a at different speeds, given Δ equal to one hour. The flight speeds and fuel costs in columns 1, 2 and 3 are generated using the Flight Planning Engine. Passenger-related delay costs in column 4 are computed for the schedule specified in column 2, using the Passenger Delay Evaluation Module. Table 1 and Figure 3 summarize the fuel costs and the passenger-related delay costs to the airline, corresponding to the different flight plans (or CI values), thus depicting the trade-off between the flying time and total cost.

Compared to the flying time of 431 minutes (at CI 500), the originally planned flight time of 454 minutes (or CI 40) results in many more disrupted passengers, more extensive delays and increased need for re-accommodation. At CI 500, there is a sharp reduction in the passenger cost function as several passengers can make their planned connections, while at CI 40, these passengers misconnected. More fuel is burned with the faster speed, but not so much as to offset the improvements

Cost Index (CI)	Flight Time	Fuel burn ($)	Passenger-related delay cost ($)	Total Cost ($)
20	455	53772.80	103396.50	157179.30
40	454	53776.10	103337.10	157113.20
60	454	53777.00	103337.10	157114.10
80	453	53838.80	103337.10	157175.90
100	451	53957.80	103396.50	157354.30
300	442	55962.10	102010.00	157972.10
500	431	58401.00	42715.30	101116.40
700	427	60013.00	41308.60	101321.60
900	426	60551.60	41249.20	101800.80
1100	424	61651.30	38361.60	100012.90
1500	423	62942.70	38302.20	101244.90

Table 1   **Flight time - cost trade-offs associated with different flight plans**
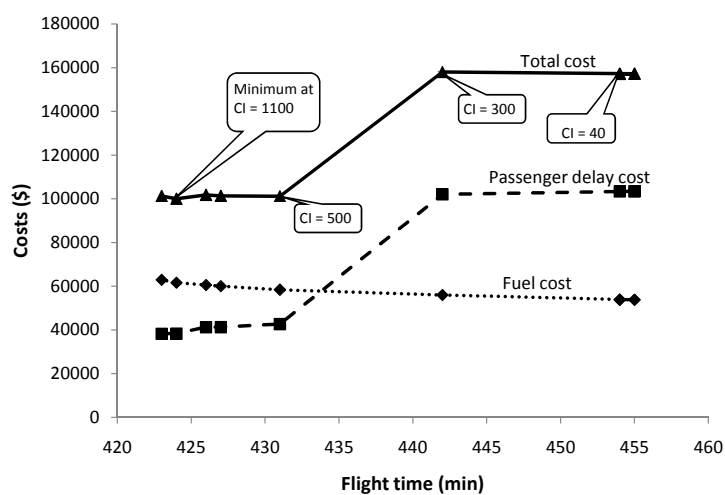

Figure 3   **Trade-off between flight time and associated costs**

in passenger-related delay costs to the airline. The sum of passenger delay and fuel burn costs is minimized at a flight time of 424 minutes (at CI 1100). However, the airline from which this example is extracted, typically operates at CI 30 and allows its dispatchers and pilots to speed up to a maximum of CI 300. It is clear from the figure that neither of these CI values truly minimizes costs. While speeding up the flight from CI 40 to CI 300 may be viewed by the pilot as 'making up time', in fact it simply burns excess fuel and increases total cost.

The example illustrates that it is possible to optimize (minimize) total costs by increasing (or decreasing) aircraft speeds relative to those planned. The appropriate speed change and the realized benefit depends on both Δ and on the degree and structure of passenger itinerary connectivity, and is not well-captured using the static rule-of-thumb approach currently used in practice. There has been growing understanding of the shortcomings of current practice, as discussed in Burrows et al. (2001) and Altus (2010), but models to overcome these limitations have not been built.

The example further illustrates that a more effective way to choose a flight speed is to optimize costs in real-time. We present in the following section an optimization-based framework that

allows us to optimize operating fuel costs and performance costs as measured by passenger service reliability, thereby minimizing total costs incurred.

Traditional disruption management practice does not capture elements of speed changes as a means to enable planned (but delayed) connections during operations. Flight planners also do not capture the network impacts of the schedule during operations as described in Altus (2010). Our work serves to illustrate that by combining these elements, improvements in total costs are achievable.

Our integrated flight planning and disruption management approach uses at its core, an optimization model that requires input information about possible flight plans in a disrupted scenario and combines them with existing mechanisms of swaps, cancelations and holding flights. Our optimization model is schematically shown in Figure 4. The flight planning engine that feeds into this module has been described earlier in §3.1.1.
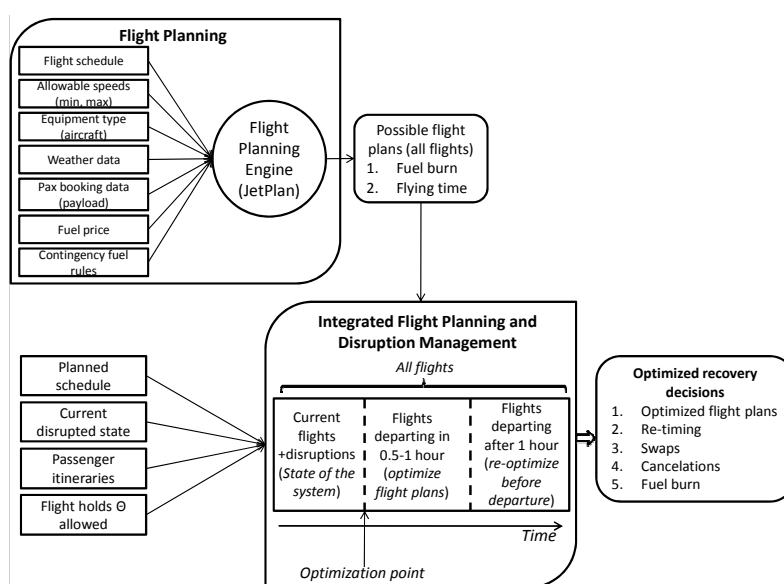
Figure 4   Optimization module for integrating flight planning and disruption management

Our experimental approach is described in Figure 5. The optimization module is integrated with a data analysis module that analyzes historical data to generate statistically significant scenarios. This serves as input into the optimization. We evaluate the solutions from the optimization using an independently built simulator, described previously in §3.1.2, to compute *true* delay and disruption costs.

Figure 5   **Experimental approach for integrated flight planning and disruption management**
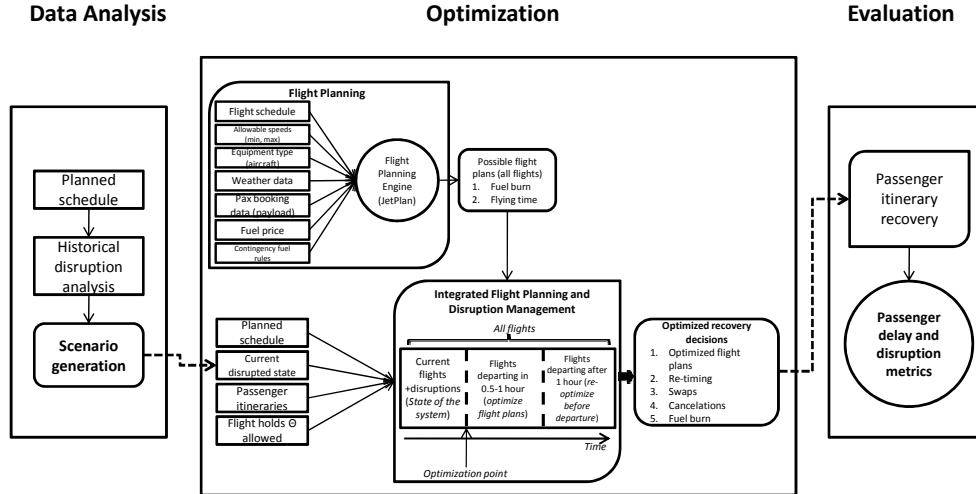
4.   Modeling Framework

In our enhanced disruption management approach incorporating flight planning, schedule and flight plan optimization is performed prior to each flight, at the time just before the flight plan is filed for the flight. This provides the ability to produce different flight planning solutions during operations; solutions designed to capture the features of aircraft and passenger connectivity for that flight given current schedules, and further network effects that propagate throughout the network. We focus on aircraft and passenger disruption management. With suitable modifications, this can be extended to include crew disruption management also.

Given a disrupted schedule, an airline defines a recovery time-window of duration T, starting from the current time on the day of operations, beyond which normal operations should be resumed. The time-window is defined to begin at least one hour prior to departure of the long-haul flights that are delayed at departure to their hub. It consists of both the arrival and departure banks of long-haul flights and extends over 48 hours, or even 72 hours in case of a very large disruption.

At the start of the time window, we assume that we have a snapshot of the airline's schedule and resource allocations at that point in time. That is, we assume knowledge of: (i) current aircraft locations, planned aircraft rotations and maintenances; (ii) currently delayed flights, and (ii) planned passenger itineraries (and therefore disrupted itineraries). We refer to this information as the *airline system state* at time t.

Knowing the airline system state, estimated times of departure for flights delayed at departure, and scheduled times for non-disrupted flights, we create time-space network representations (described in §5) for aircraft and passenger movements. We create appropriate copies (also described

in §5) representing possible arrival and departures of flights in the time-space networks, allowing for re-scheduling as well as speed changes for flights. With these networks underlying the model, we solve the enhanced disruption management and flight planning formulation(s) (presented in §5) which provides a schedule that minimizes the sum of passenger delay and fuel burn costs.

4.1. Assumptions

The following assumptions are considered when building our models: (i) A flight cannot be cleared for departure prior to its scheduled departure time; (ii) The decrease in payload (and hence the decrease in fuel burn) due to passengers mis-connecting is negligible; and (iii) If a flight plan with a significantly different arrival time at the destination airport is used, there is a landing slot available at that time.

5. Mathematical Models
5.1. Time-space networks

Our model is based on *time-space network* representations of the airline's schedule. The nodes in a time-space network are associated with both time and location, and an arc between two nodes indicates a possible movement between the two locations (or same location) and times. Given the state of the system at time t, we create time-space networks within the time-window T whose arcs are based on (i) estimated departure and arrival times of disrupted flights in the system, and (ii) scheduled departure and arrival times of non-disrupted flights, and (iii) possible re-timings and speed changes of flights. In fact, we create two different types of time-space networks: (i) an aircraft flow network for each aircraft and (ii) a passenger flow network for the passengers.

We introduce some notation useful for building our network representations. Let F be the set of flights f that have both departure and arrival within the time window T. Among these, let F' be the set of long-haul flights, for which we consider speed changes, re-timing and swaps; and $F - F'$ be the set of short-haul flights for which we consider flight re-timings alone but no speed changes (speed changes are not significant for short-haul flights). Let A be the set of aircraft a available and A' be the set of aircraft that operate long-haul flights. Let Π be the set of fleet types available, and $\pi(a)$ denote the fleet type assigned to aircraft a in the original schedule. Let $F_{\pi(a)}$ represent the set of flights of the same fleet type as aircraft a, and $F_{FML(a)}$ the set of flights of the same fleet family as aircraft a. For aircraft $a \in A'$, we allow tail swaps (swaps within $F_{\pi(a)}$) and fleet swaps (swaps within $F_{FML(a)}$). Let P be the set of passenger itineraries operated within the time window T, and n_p represent the number of passengers on itinerary $p \in P$. Let D_L and D_S represent the set of long-haul and short-haul flights, respectively, that are immediately downstream of any long-haul flight $g \in F'$ in a passenger itinerary $p \in P$.

5.1.1.   Aircraft flow networks For each aircraft $a \in A$, we create a time-space network N_a to track its movement over the flight schedule. N_a spans the length of the time window T and consists of flights that can be operated by aircraft a, that is, N_a contains $F_{\pi(a)}$ and $F_{FML(a)}$. The operations of these flights are captured by multiple arc copies for each flight, that represent possible departure and arrival times of each flight, as we will describe below. Each node in the aircraft flow network represents either a possible departure time of a flight f, or a possible arrival time of the flight *plus* the minimum turn time of aircraft a. We represent the set of nodes in N_a as N_a'. In N_a for each aircraft a, a supply $s^n = 1$ is associated with the node n where the aircraft is known to start at the beginning of the time window T, and a demand of $s^n = -1$ where it completes the last flight in the network. All changes to aircraft a's path and schedule are thus limited within the time-window.

For each flight leg f, we denote the set of flight copies (over all networks N_a) as C_f. Each flight leg copy $k \in C_f$ connects a possible departure time of flight f to a possible arrival time (corresponding to a specific flight plan) plus the minimum turn time of aircraft a. N_n^- is the set of incoming arcs to each node $n \in N_a'$ and N_n^+ is the set of outgoing arcs from each node $n \in N_a'$. A ground arc exists from each node to the next (time-wise) node at each location, and allows feasible aircraft paths to be defined. We refer to the set of ground arcs in N_a as G_a.

Because we consider disruptions to long-haul flights, we create flight copies to model recovery for these flights. For each long-haul flight, that is, for all $f \in N_a$ where $a \in A'$, we create copies of the following types. The first represents the originally scheduled (or estimated, in the case of delayed flights) time of departure and arrival. The second represents alternative departure times of the flight (compared to the original) without speed changes. For this, we generate copies of the flight every 5 minutes until a maximum departure delay of R minutes after its estimated departure time. The third represents flight plans that involve speed changes (created using the Flight Planning Engine), and therefore, represent block time changes to the original flight, but with the same departure time as that of the original flight. The fourth type represents flight plans involving speed changes for the different departure times of flights specified is a combination of the second and third types, representing both a different departure time and a different speed (block time) of operation.

Additionally, copies are created to model the holding of downstream flights in passenger itineraries. Let Θ (in minutes) be the maximum extent to which downstream flights departures are allowed to be intentionally held or delayed in order to facilitate passenger connections. That is, downstream connections are allowed to arrive at most Θ minutes late at their destination. Then, in each network N_a, we have the following cases. First, for all long-haul flights $f_1 \in D_L$ (that is, f_1 follows a long-haul flight in a passenger itinerary) we will capture possible speed changes as

well as delaying departure times for this flight. Let the maximum decrease in block time possible due to speeding up f be δ. We create possible departure nodes of flight f every 5 minutes until a maximum departure delay of $\Theta + \delta$ (if such nodes are not already created in previous steps), and corresponding flight copies representing speed changes due to alternative flight plans for each possible departure time. Thus, we ensure that the arrival delay of the downstream flight copies at the destination is no more then Θ minutes. Second, for all short-haul flights $f_2 \in D_S$, we simply make copies of the flight with its scheduled block time and departure arcs at 5 minute intervals until a maximum departure delay (and a corresponding arrival delay) of Θ. However, from the short-haul flights held for passengers, delay can be propagated further downstream to other short-haul flights due to aircraft arriving late. Therefore, for flights $f_3 \in F - F'$ downstream of f_2, we again create copies of the flight with its scheduled block time and departure arcs at 5 minute intervals until a maximum departure delay (and a corresponding arrival delay) of Θ.

In our experiments, we choose values of Θ to be 0, 5, 10 and 15 minutes. We impose a limit of 15 minutes on Θ so that arrival delay of a downstream flight f due to delay propagated to it from an upstream flight via passenger connections is limited to 15 minutes. This is so that the on-time performance of the system (determined by delays greater than 15 min) is not deteriorated.

Scheduled maintenance for an aircraft a, if scheduled within the time window T, is modeled by creating an artificial 'flight leg' in N_a, beginning at the start of maintenance at the maintenance station and ending at the end of the scheduled maintenance at the same station. If maintenance can be delayed, we capture it by creating copies of this arc. Maintenance arcs for aircraft a are only present in N_a and not in other networks.

Three types of costs are associated with each flight copy in the aircraft flow networks. First is the incremental fuel cost c_f^k, for each flight copy $k \in C_f$ for each flight f, compared to the 'normal' CI of operation. c_f^k can be positive or negative, and is obtained from the flight planning engine. Second are aircraft swap costs of operating the flight with an aircraft other than that planned. Let s_f^a denote the swap cost of operating flight f with aircraft a (equals zero if a is the originally planned aircraft routing). Let $\zeta_{(a,f)}^k$ be an indicator that is 1 for all copies $k \in C_f$ that belong to N_a. Then with flight copy $k \in C_f$, a swap cost of $s_f^a \zeta_{(a,f)}^k$ is associated. Third are incremental costs d_f^k of delayed departure, of \$10 per minute, associated with each minute a flight is delayed beyond its scheduled departure.

The above description of construction of N_a for each aircraft a is equivalent to creating a single aircraft flow network, with flight copies for each aircraft exactly as described above; and modeling it as a multi-commodity network flow problem with each aircraft as a commodity. However, for

ease of exposition, we have described it as an aircraft flow network for each individual aircraft. Our model differs from previous literature, where typically, each fleet has been modeled as a commodity (Thengvall et al. (2000), Andersson and Varbrand (2004), Yan and Young (1996) and Bratu and Barnhart (2006)).

5.1.2.   Passenger flow networks Similar to aircraft flow networks, we also construct a *passenger flow network* N_P that captures all passenger itineraries. Each node in the passenger flow network represents either a (scheduled or possible) departure of flight f, or an arrival of flight f for a passenger on that itinerary. N_P contains flight copies combined from aircraft flow networks N_a for all a (the minimum turn time for each aircraft, however is not included in defining the nodes in N_P), with the exception of maintenance arcs. Thus, the flight copies represent the operation of the flight with differing departure and arrival times, and by different aircraft. Connection arcs at each location connect successive flight legs, when feasible (exceeds the minimum connecting time) in a passenger itinerary. On this network, each passenger is modeled as a commodity, with its origin at the departure node of the first flight leg on their itinerary, and its destination as the latest node of the passenger's destination airport in the time-window. The destination node of passenger p is set in this manner to capture potential disruptions, and re-accommodation.

Let N_p' be the set of nodes on itinerary p and G_p be the set of ground arcs for p in the passenger flow network N_P. To model passenger re-accommodation, first, we generate candidate itineraries $R(p)$ for each passenger type p. If passenger itinerary p is not disrupted at time t, $R(p)$ consists only of the originally scheduled itinerary. If passenger itinerary p is disrupted in the scenario considered, $R(p)$ is a list of candidate itineraries or paths on the passenger flow network from the origin of p to its destination, with each starting after p's original departure, by at least the amount of disruption of p. $R(p)$ also includes a virtual itinerary to indicate re-accommodation to another airline's network, or perhaps, cancelation of the passenger trip at its origin. d_p^r represents the arrival delay of passengers originally on itinerary p who are re-accommodated on itinerary r. Passenger-related costs c_p^r denote the cost of using itinerary r to accommodate passenger p. This is based on the actual arrival time of itinerary $r \in R(p)$ at the destination, and includes delay costs and goodwill costs. Parameters Cap_f are the number of seats on flight f and parameter δ_f^r is 1 if flight f is on itinerary r and zero otherwise.

5.1.3.   Variables Let x_f^k be a binary variable that takes on value 1 if copy k of flight leg f is present in the solution and 0 otherwise, y_g be a binary variable that is 1 if ground arc g is present in the solution and 0 otherwise, and z_f be a binary variable that is is 1 if flight f is canceled in

the solution and 0 otherwise. Let ρ_p^r be the number of passengers originally on itinerary p who are re-accommodated on itinerary r (ρ_p^p equals the number of non-disrupted passengers traveling on their originally scheduled itinerary).

5.2. Aircraft Recovery and Passenger Re-accommodation Model

We propose models to minimize the sum of multiple operating costs, including incremental fuel costs, swap costs and passenger-related delay costs (for re-accommodation and recovery). The following is our formulation for combined aircraft recovery and passenger re-accommodation including flight planning opportunities.

$$\min \sum_{f \in F} \sum_{k \in C_f} (c_f^k + s_f^a \zeta_{(a,f)}^k + d_f^k) x_f^k + \sum_{p \in P} c_p^r \rho_p^r \tag{5.1}$$

$$\text{s.t.} \sum_{k \in C_f} x_f^k + z_f = 1 \qquad\qquad \forall\, f \in F \tag{5.2}$$

$$\sum_{g \in N_n^-} y_g + \sum_{(f,k) \in N_n^-} x_f^k + s^n = \sum_{g \in N_n^+} y_g + \sum_{(f,k) \in N_n^+} x_f^k \qquad\qquad \forall\, n \in N_a', \forall\, a \in A \tag{5.3}$$

$$\sum_{r \in R(p)} \rho_p^r = n_p \qquad\qquad \forall\, p \in P \tag{5.4}$$

$$\sum_{p \in P} \sum_{r \in R(p)} \delta_f^r \rho_p^r \le Cap_f(1 - z_f) \qquad\qquad \forall\, f \in F \tag{5.5}$$

$$x_f^k \in \{0,1\} \qquad\qquad \forall\, k \in C_f, \forall\, f \in F \tag{5.6}$$

$$\rho_p^r \in \mathbb{Z}^+ \qquad\qquad \forall\, r \in R(p), \forall\, p \in P \tag{5.7}$$

$$y_g \ge 0 \qquad\qquad \forall\, g \in G_a, \forall\, a \in A \tag{5.8}$$

The objective (5.1) is to minimize the sum of incremental fuel costs, swap costs, incremental delay costs and passenger delay costs. Constraints (5.2) ensure that a flight is either operated using one of the copies created, or canceled. This flight cover constraint is also applied to the copies of maintenance arcs (described in §5.1.1) with the corresponding z variable set to 0 to disallow cancelation of maintenance. Thus we ensure that compulsory maintenance is carried out, and to exactly the aircraft to which it is assigned. Constraints (5.3) require flow balance for each aircraft. Constraints (5.4) ensure that all passengers reach their destinations either on their original itinerary or an alternate one. Constraints (5.5) ensure that no passengers are assigned to a canceled flight leg, and restrict the number of passengers assigned to a flight leg to its capacity. Constraints (5.6), (5.7) and (5.8) restrict variable values to appropriate binary or integer values. The constraint that y_g is binary can be relaxed to $y_g \ge 0$ because x variables are binary.

5.3. Approximate Aircraft and Passenger Recovery Model to Trade-off Fuel Burn and Passenger Cost

Solving the aircraft and passenger recovery model with passenger re-accommodation described in (5.1) - (5.8) can require excessive time for real-time decision making. Feasible solutions obtained when the model is stopped after 5 minutes result in high operating costs. This is likely due to the large sizes of the problem caused by the aircraft-specific networks, flight copies from alternate flight plans, departure times and the corresponding copies to model aircraft swaps, and capacity constraints (5.5) that often result in fractional solutions (as observed in Barnhart et al. (2002) and Bratu and Barnhart (2006)). Thus (5.1) - (5.8) may not suitable for application when decisions have to be made in a few minutes. To address this, we introduce an alternative model that captures approximately the trade-off between fuel burn and passenger delays.

In addition to the notation in (5.1) - (5.8), let $IT(p)$ be the set of flight legs in itinerary p, $IT(p,l)$ the lth flight leg in itinerary p. Let n_f be the number of booked passengers whose itineraries terminate with flight leg f; δ_f^p equal 1 if itinerary p terminates with flight leg f, and 0 otherwise. We let $MC(p,f,k)$ denote the set of flight leg copies f' (the flight to which f connects in itinerary p) in the passenger flow network N_P, to which there is insufficient time to connect from copy k of flight leg f. Let λ_p be a binary variable that is 1 if itinerary p is disrupted and 0 otherwise, and let $\tilde{c}_p$ be the approximate cost of disruption per passenger on itinerary p. $\tilde{c}_p$ is an *approximate* costs of re-accommodation for each disrupted itinerary $p \in P$, because we assume that if passenger itinerary p is disrupted, the passengers on itinerary p are re-accommodated on the next available itinerary to the destination in the next flight bank. Based on this assumption, we compute the per passenger estimated arrival delay cost to the airline for passengers on itinerary p. c_p estimates the costs incurred *by the airline* due to passenger delays, including recovery costs and goodwill costs corresponding to the arrival delay. Setting a cost per itinerary p also allows the capture of non-linearity in costs, where higher delays incur disproportionately higher costs compared to smaller delays. Our modified aircraft recovery model with passenger disruptions is as follows:

$$\min \sum_{f \in F} \sum_{k \in C_f} (c_f^k + s_f^a \zeta_{(a,f)}^k + d_f^k) x_f^k + \sum_{f \in F} c_f z_f + \sum_{p \in P} \tilde{c}_p n_p \lambda_p \tag{5.9}$$

$$\text{s.t.} \sum_{k \in C_f} x_f^k + z_f = 1 \qquad\qquad \forall f \in F \tag{5.10}$$

$$\sum_{g \in N_n^-} y_g + \sum_{(f,k) \in N_n^-} x_f^k + s^n = \sum_{g \in N_n^+} y_g + \sum_{(f,k) \in N_n^+} x_f^k \qquad\qquad \forall n \in N_a', \forall a \in A \tag{5.11}$$

$$x_{IT(p,l)}^k + \sum_{m \in MC(p,IT(p,l),k)} x_{IT(p,l+1)}^m - \lambda_p \le 1 \qquad\qquad \forall k \in C_{IT(p,l)},$$

$$\forall\, l \in 1, .., |IT(p)| - 1, \forall\, p \in P \quad (5.12)$$

$$\lambda_p \geq z_f \qquad\qquad \forall\, f \in IT(p), \forall\, p \in P \quad (5.13)$$

$$x_f^k \in \{0,1\} \qquad\qquad \forall\, k \in C_f, \forall\, f \in F \quad (5.14)$$

$$z_f \in \{0,1\} \qquad\qquad \forall\, f \in F \quad (5.15)$$

$$\lambda_p \in \{0,1\} \qquad\qquad \forall\, p \in P \quad (5.16)$$

$$y_g \geq 0 \qquad\qquad \forall\, g \in G_a, \forall\, a \in A \quad (5.17)$$

The objective function (5.9) sums up the the incremental fuel costs of flights, swap costs, the incremental costs of flight delays, costs of flight cancelations and the costs to the airline of passenger itinerary disruptions. Constraints (5.12) ensure that itineraries with insufficient connection time are classified as disrupted. Because the value of c_p is greater than zero, this constraint ensures that λ_p is 1 if and only if both terms in (5.12) are 1, that is, if passengers on itinerary p cannot connect from one leg on their itinerary to the following leg on their itinerary. Constraints (5.13) similarly ensure that if a flight leg is canceled, all itineraries containing the flight are classified as disrupted. Constraints (5.16) can be relaxed to $0 \leq \lambda_p \leq 1$ because x and z variables are binary. In all other cases, λ variables will be zero because of the positive cost associated with them in the objective. Constraints (5.10), (5.11), (5.14), (5.15) and (5.17) ensure flight cover (or scheduled maintenance as described in §5.2), aircraft balance, and binary values of variables x, z and y respectively, as discussed for (5.1) - (5.8).

6. Experimental Setup
6.1. Network Structure and Experiment Design

In this section, we demonstrate the potential impact of disruption management enhanced with flight planning, using data obtained from a major European airline (specified in §3.1.3). The airline operates a hub-and-spoke network with about 250 flights per day serving about 60 cities daily and multiple continents. (This does not include feeder airline flights.) The airline operates a banked schedule at its hub. About 243 flights, or 93% of the flights operated by the airline are into or out of the hub. 10% of the flights (approximately 30 arrivals and departures per day) operated are long-haul, and present significant opportunities for speed changes. The remaining 90% of flights are medium-haul and short-haul. Aircraft rotations on this network are typically designed as cycles originating from and ending at the hub, with each cycle consisting of 2 to 4 flights. This is particularly true of short- and medium-haul flights that operate within Europe, which are operated as short cycles around the hub. Long-haul flight operations comprise more than 30% of the flying hours of the airline per day. About 40% of the passengers have at least one long-haul flight on

their itinerary. Because these itineraries bring in more revenue than itineraries with only short-haul flights, we estimate that about 50% of the revenue is associated with passenger itineraries containing long-haul flights.

In our experiments, we focus on disruptions of long-haul flights inbound to the hub. The first reason for this focus is a significant percentage of passengers connect at the hub from international locations into Europe and vice versa, and therefore the hub presents the best opportunity to affect passenger connectivity. A second reason is flight planning opportunities, in particular, speed changes, are significant for long-haul flights.

Our models are implemented in C++ and use Xpress. Computational experiments are conducted on a server using a 64 bit Intel Xeon E5440 2.83 GHz processor with 4 cores and 16 GB RAM. Because our models are designed for real-time application, we limit the solve time to 2 minutes and evaluate the best solution found.

6.2.   Historical Delay Analysis and Scenario Generation

We describe here our scenario generation process, depicted in the Data Analysis module of Figure 5. We conduct an analysis of delays of long-haul flights that are inbound to the hub and generate distributions of these delays. Our historical delay analysis is conducted for data available for the months of June and July 2008. Unfortunately, passenger information for this period is not available. We have passenger data only for a period of two weeks in November 2008 (for which we do not have flight delay data). We replicate each instance of disruption for each day for which the same schedule occurs and passenger data is available. Thus, each delay scenario may be solved multiple times, once for each day a similar schedule re-occurs for the two weeks in November 2008 for which passenger data is available. We test and evaluate a total of 60 scenarios.

6.3.   Parameter assumptions

We assume the following values for the parameters in the model:

- Passenger-related delay costs to the airline = \$1.09/passenger per minute, for 2008. This number is the airline's estimate of *its own cost* incurred for passenger delays, including recovery, re-accommodation and goodwill cost.

- Fuel cost = \$3.65/gal or \$0.478/lb. This estimate is a result of the airline AOCC's reported costs (including taxes) of €700 - €800 per metric ton of fuel in February 2010. We assume an average cost of €750 per metric ton. This value is converted to €0.34/lb or \$0.43/lb in February 2010. (with density 0.82 kg/litre, 3.6 gal/litre, €1 = \$1.27 European Central Bank (2010) for Nov 2008). Further, guided by the IATA fuel price development charts (International

Air Transport Association 2010), a ratio of 0.903 for costs in November 2008 to February 2010 is applied, to convert the price to \$0.478/lb in November 2008.

- T = approximately 1.5 days or the end of the propagation boundary, encompassing at least two successive arrival banks at the hub to allow for aircraft swaps.

- Normal CI = 30; rule-of-thumb maximum CI specified by the airline = 300

- c_p = Cost per disrupted passenger in model (5.9) - (5.17) = \$457.8. This cost is calculated assuming that disrupted passengers are re-accommodated in the next bank, with an average re-accommodation time of 7 hours, by calculating the average time to the next connecting flight for different passenger itineraries.

- swap cost s_k^f = \$500 per tail swap, \$1000 per fleet swap

- flight cancelation cost = \$20,000

6.4. Models being compared

For the evaluation module used in Figure 5, we compute passenger delay and disruption metrics for solutions to the following optimization models.

6.4.1. Baseline for comparison: Sequential recovery We use as a basis for comparison the solution to our model allowing aircraft recovery with possible departure delays but no flight changes and passenger disruptions not accounted for. We accomplish this by solving formulation (5.9) - (5.17), with $c_k^f = 0 \, \forall \, k \in K, \forall \, f \in F$ and $c_p = 0 \, \forall \, p \in P$. The objective in this model is to minimize the sum of flight departure delay costs, swap costs and cancelation costs.

6.4.2. Airline rule-of-thumb with flight planning In this case, we generate solutions allowing flights to be delayed intentionally and allowing flight speeds at the normal flight speed or at the maximum speed specified by the rule-of-thumb. To generate these solutions, we solve the formulation (5.9) - (5.17) without passenger connection constraints (5.12), (5.13) and (5.16); with only two arc copies per flight representing the normal flight speed and the maximum speed specified by the rule-of-thumb, and with $c_p = 0$. The objective in this model is to minimize the sum of fuel costs, flight departure delay costs and cancelation costs.

6.4.3. Aircraft-centric (sequential) recovery with speed changes, with flight planning. In this case, we generate solutions allowing flight holding and flight speed changes, but ignore the resulting passenger disruption costs. To generate these solutions, we solve formulation (5.9) - (5.17) without passenger connection constraints (5.12), (5.13) and (5.16) and with $c_p = 0$. The objective in this model is to minimize the sum of fuel costs, flight departure delay costs and cancelation costs.

6.4.4. Passenger-centric disruption management approach without flight planning
In this case, the problem is to find the best solutions to (5.9) - (5.17) but to disallow flight arcs corresponding to speed changes. We let $\Theta = 0, 5, 10$, and 15 minutes and solve the corresponding formulation (5.9) - (5.17)

6.4.5. Enhanced disruption management with flight planning We solve the complete model (5.9) - (5.17), as described in §5, with $\Theta = 0, 5, 10$, and 15 minutes.

6.5. Evaluation - passenger recovery and delay estimation
Each solution generation model described in §6.4 produces a schedule which we evaluate for passenger delays in the evaluation module depicted in Figure 5 and detailed in §3.1.2. Note that the simulated costs obtained using this approach are different from the objective function values for the models in §6.4, because the models approximate passenger disruption costs based on misconnections, whereas the simulator estimates the *actual* passenger costs to the airline based on actual re-accommodation. Details of our experiments are presented in §7.

7. Results and discussion
7.1. Case study 1
To illustrate the tradeoffs occurring in these problems, we first consider a simple case: one flight f is delayed by Δ minutes into the hub, while all other flights operate as scheduled. We consider different levels of delay Δ, on each of 12 days of operation of f, in November 2008. Flight f is representative of the other flights in the network in that the trends and trade-offs observed with this flight are also seen in the case of other flights. In this case, we vary Δ from 10 minutes to 60 minutes in intervals of 10 minutes.

Figure 6 shows the change in fuel burn and passenger cost curves for different levels of Δ, for selected representative days of operation. The horizontal axis represents the arrival delay of flight f and the vertical axis represents costs incurred. For each value of Δ, the fuel cost curve can be plotted to reflect speed changes in f, resulting in different arrival delays and corresponding fuel burn. Fuel cost curves are marked by Δ values from 10 to 60 in the upper portion of the figure. As the value of Δ increases, the fuel cost curve itself does not change shape, but shifts to the right to reflect increased arrival delay.

In the case when downstream flights are not held for passengers ($\Theta = 0$), the passenger-related airline costs incurred for different levels of flight arrival delay are shown, for instances across five days of data. (These are indicated in the lower part of Figure 6.) As arrival delay increases, passenger delay increases and more passenger misconnects occur. The delay cost curve incurs a
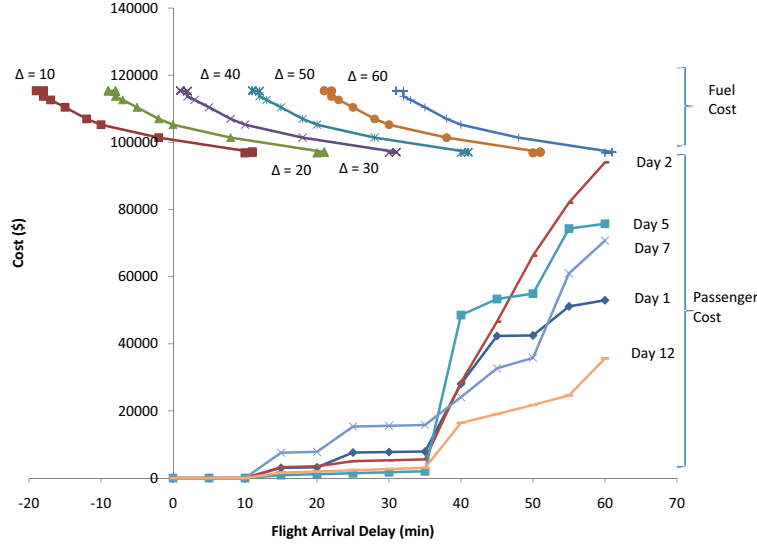
Figure 6 **Trade-offs between fuel burn and passenger delay costs over multiple days**

'jump' when a set of passengers misconnect and require recovery and re-accommodation. The total cost curve that is a sum of fuel costs and passenger costs changes dynamically with changes in the delay Δ of flight f. This we illustrate using Figure 7, which demonstrates the changes in the total cost curve for different Δ, for one day of operations (represented by Day 2 in Figure 6).

Figure 7 serves to illustrate that total costs can differ dramatically with changes in Δ. During operations, the propagation impacts of Δ can be adjusted when the departure times of downstream flights are allowed to be altered (or are altered in the course of the day, due to plans not operating exactly as planned) so that passengers can make connections. Holding downstream passenger connections opens up the possibility of the upstream flight speeding up to a smaller extent and burning less fuel, but incurring fewer misconnections. The network interactions now become more interesting, as we have the flexibility of changing speeds and departure times of inbound delayed flights as well as the outbound flight departure times.

We now describe the phenomena that occur when flight speeds and departure times are simultaneously modified to mitigate the effects of disruptions. We do so by solving the model (5.9) - (5.17), with different values of Θ. $\Theta = 0$ results in the phenomenon so far discussed and described in Figures 6 and 7. Now we present the fuel burn and passenger-related airline costs (costs estimated via simulation) when (5.9) - (5.17) is solved for $\Theta = 0$, 10 and 15, for the specific flight f and each value of Δ; and compare them to our baseline results. The results presented in Table 2 are over a 12-day period for which data is available for this flight. The costs presented are the percent savings related to the baseline disruption management model described in 6.4.1. We present the fuel burn, passenger (pax) misconnections and associated savings in costs experienced
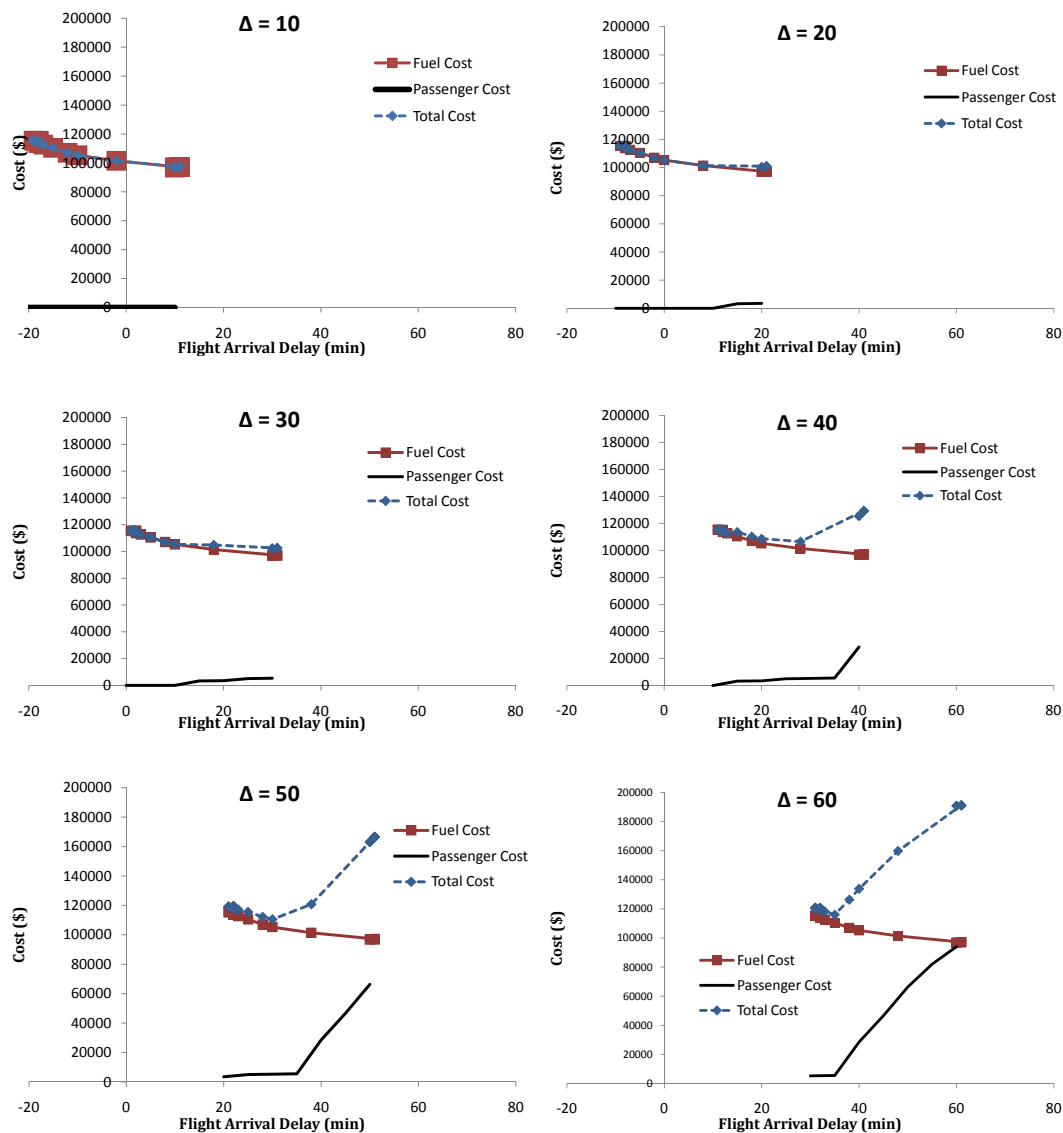
Figure 7   **Changing optimal trade-off point between fuel and passenger cost with departure delay** Δ

by the airline, for results from five different strategies of disruption management: (1) Column 1: The baseline disruption management strategy described in Section 6.4.1, which does not allow for speed changes; (2) Column 2: A disruption management strategy that combines the baseline disruption management strategy with the airline's rule-of-thumb speed up strategy specified in §6.4.2; (3) Column 3: Our enhanced disruption management strategy that combines flight planning with disruption management using (5.9) - (5.17), as described in §6.4.5, with Θ set to 0; (4) Column 4: Our enhanced disruption management strategy that combines flight planning with disruption management using (5.9) - (5.17), as described in §6.4.5, with Θ set to 10 minutes; (5) Column 5: Our enhanced disruption management strategy that combines flight planning with disruption management using (5.9) - (5.17), as described in §6.4.5, with Θ set to 15 minutes.

	Baseline recovery	Rule-of-thumb speed up (to CI 300)	Enhanced recovery: don't hold connecting flights	Enhanced recovery: hold connecting flights up to 10 min	Enhanced recovery: hold connecting flights up to 15 min
$\Delta = 10$ min					
Fuel savings per operated LH flight %	-	-4.69	0.05	0.05	0.05
pax disruption savings %	0 disruptions	N/A	N/A	N/A	N/A
delayed pax cost savings %	0 disruptions	N/A	N/A	N/A	N/A
Total cost savings %	-	-4.69	0.05	0.05	0.05
$\Delta = 20$ min					
Fuel savings per operated LH flight %	-	-4.69	-2.72	-2.72	0.05
pax disruption savings %	-	99.72	93.45	96.58	99.72
delayed pax cost savings %	-	93.66	80.30	82.80	93.66
Total cost savings %	-	-0.22	1.05	1.17	4.30
$\Delta = 30$ min					
Fuel savings per operated LH flight %	-	-4.69	-3.40	-2.73	-2.73
pax disruption savings %	-	42.17	71.23	93.73	96.58
delayed pax cost savings %	-	39.32	57.94	68.64	70.61
Total cost savings %	-	-2.54	-0.40	0.75	0.85
$\Delta = 40$ min					
Fuel savings per operated LH flight %	-	-4.69	-5.95	-4.33	-4.33
pax disruption savings %	-	56.45	70.35	94.29	94.42
delayed pax cost savings %	-	73.20	72.57	85.73	85.49
Total cost savings %	-	7.53	6.37	9.80	9.76
$\Delta = 60$ min					
Fuel savings per operated LH flight %	-	-4.69	-6.24	-8.66	-6.24
pax disruption savings %	-	13.72	70.70	71.35	72.40
delayed pax cost savings %	-	22.54	84.56	85.45	64.43
Total cost savings %	-	3.87	22.31	20.94	15.98

Table 2   **Single flight delay: simulated percentage savings for different recovery strategies, averaged over 12 days of operation**

From our analysis, we present the following findings:

1. The rule-of-thumb that is sometimes adopted by dispatchers, of speeding up to the allowable speed (Column 2 solutions), results in improved passenger costs compared to the baseline recovery model, as it improves the on-time performance of the flight. However, for medium levels of disruption, such as 20-30 minutes, it results in increased fuel consumption even in cases where it may not be required. This rule-of-thumb-based policy may be able to recover passengers for $\Delta = 10$ and $\Delta = 20$, but may fall short for larger disruptions. In comparison with our optimization-based models, the rule-of-thumb speed-up policy almost always results in higher costs.

2. Our enhanced recovery approach compared to the baseline and rule-of-thumb-practices generally reduces total costs and passenger-related delay costs for the airline significantly (except for some scenarios as explained below).

   (a) For all levels of disruption, recovery models enhanced using flight planning will hold constant or decrease total passenger-related delay costs compared to the baseline approach. For the example in Table 2, the total cost improvements ranged from 0 to 15%.

   (b) Depending on the itineraries of passengers, both flight speed changes as well as holding of

downstream flights might be necessary to reduce the number of misconnected itineraries. For example, in the case of a 20-minute initial disruption, simply allowing speed changes without holding downstream flights is sufficient to recapture 93% of the disrupted passengers back onto their original itineraries compared to the baseline case. In the case of 30- and 40-minute delays, however, only about 70% of passenger misconnects can be prevented by allowing speed changes without holding flights. By also allowing downstream flight departures to be delayed by 10 minutes, the misconnects are decreased by about 95%.

(c) Because the objective function of the model (5.9) - (5.17) estimates costs approximately, and hence differently, from those calculated in the simulation, discrepancies may be (rarely) observed in cases for which passenger delay costs calculated in the simulation are not well approximated by the passenger disruption costs in the objective function of (5.9) - (5.17). For example, note the negative value of realized total cost for $\Delta = 30$ in column 3 of Table 2. The simulated passenger delay cost savings are not as high as indicated by the objective function value when (5.9) - (5.17) is solved, resulting in increased costs instead of savings.

3. Low levels of disruption:

(a) For very low levels of disruption (for example, 10 minutes), enough slack is present in the system to absorb the disruption, and flight planning mechanisms such as speed increases and holding downstream flights are not required. Instead, we might be able to slow down the flight without incurring disruptions. However, for even fairly low levels of disruption such as 20 minutes, the interaction between speed changes and passenger delays can come into play. In the case of the flight demonstrated in Table 2, some passenger connections have a small amount of slack for which delays of 20 minutes cannot be absorbed, resulting in misconnections if the flight is not sped up. In the 20-minutes of delay case, however, almost all disruptions can be absorbed by speeding up the flight, and/or delaying downstream flights to the appropriate extent.

(b) For low levels of disruption, fuel burn costs dominate and drive the trade-off between fuel burn and passenger delay costs, as seen in the cases of 10 - 20 minutes of delay. In these cases, because fewer passengers are impacted, the balance in the optimization model tilts in the favor of decreasing fuel costs. The decision in such cases is to slow down the flight because passengers are not disrupted by the slow down. Occurrences of these levels of delays provide an opportunity to save fuel in comparison to the baseline recovery approach.

4. Intermediate levels of disruption:

(a) For intermediate levels of delay, such as 20 - 40 minutes, holding downstream flights to wait for connecting passengers can have significant benefits. In Table 2, the number of passenger misconnections decreases significantly from Column 3 to Columns 4 and 5. With the decrease in the number of misconnections, there is a corresponding decrease in passenger-related costs, for $\Delta = 20$ and $\Delta = 30$.

5. High levels of disruption:

(a) For higher levels of initial disruption (more than 30 minutes in the case shown above), passenger delay costs dominate the trade-off between fuel burn and delay costs. This is because many more downstream flight connections are impacted by a large initial disruption. To reduce the number of passenger disruptions, the optimal least total cost decision is to speed up the long-haul flight. If allowed, downstream flights are also held in order to facilitate passenger connections.

(b) We also observe from Table 2 that for departure delay levels less than 40 minutes, the number of passenger misconnections and the corresponding passenger costs significantly decrease when downstream flights are held compared to the case when only flight speeds are changed. For higher levels of delay, as shown for $\Delta = 60$, the decrease in the number of misconnections and in the passenger costs is less significant when flights are held compared to when only speed changes are allowed. (In fact, passenger delay costs increase for $\Theta = 15$.) The increase in passenger-related costs connected with holding downstream flights begins to exceed the decrease in passenger-related costs associated with re-accommodation and recovery of disrupted passengers. Thus the benefits of holding downstream flights decrease as the level of the initial disruption Δ increases to large values, because many more flights are held, and delay times for passengers on the downstream flights are increased.

7.2. Case study 2

In this section, we present a more comprehensive set of experiments and results, derived from the airline specified in §6. We consider 60 scenarios derived from the airline's historical data to test our models. In each scenario, typically, 12-14 long-haul flights are inbound to (and outbound from) the hub in each bank, among which 1-6 flights may be delayed by varying degrees from their origin into the hub. Based on the degrees of delay experienced, we categorize the scenarios, as described in Table 3. In the time-window T for which recovery is performed, there are 26-40 long-haul flights depending on the length of T. Our model formulation sizes are of the order of 200,000 rows and 200,000 columns.

If all inbound delays to the hub in the scenario are less than 20 minutes, we refer to it as a small-delay scenario; if the longest delay is greater than 20 but less than 50 minutes, we refer to it as a medium-delay scenario; if the longest delay among all flights is greater than 50 but less than 120 minutes, we refer to it as a large-delay scenario; and if there exists a delay in the scenario greater than 120 minutes we refer to it as a very large-delay scenario. Table 3 also shows the frequency of occurrences of these disruptions in June-July 2008.

Disruption type	Number	Frequency
small delays	6	0.18
medium delays	12	0.28
large delays	14	0.25
very large delays	28	0.29

Table 3   Disruption scenarios

In this section, we compare a larger set of disruption management strategies, as described in §6.4: 1. Column 1: Baseline disruption management without speed changes and not explicitly capturing passenger disruptions from §6.4.1; 2. Column 2: Aircraft-centric recovery that includes speed changes enabled by flight planning but does not explicitly capture passenger disruptions, described in §6.4.3; 3. Columns 3-6: Passenger-centric (pax-centric) disruption management approach without flight planning, with downstream departure holds set to $\Theta = 0$, 5, 10 and 15 minutes in columns 3-6 respectively. The models are described in §6.4.4; and 4. Columns 7-10: Our enhanced disruption management strategy that combines flight planning with disruption management ((5.9) - (5.17)) with downstream departure delays set to $\Theta = 0$, 5, 10 and 15 minutes respectively. We do not include the airline rule-of-thumb as it is a more costly option, as illustrated in §7.1.

As specified earlier, we impose a maximum solve time of 2 minutes for the models, and evaluate the best solution obtained thus far. MIP gaps (the ratio between the objective of the best integer solution obtained and the best linear program) equaled zero for the baseline disruption management, aircraft-centric recovery with flight planning, and passenger-centric recovery without flight planning models (except for the $\Theta = 15$ case for which small MIP gaps occurred). Our enhanced disruption management models with flight planning resulted in MIP gaps up to 50% when the disrupted scenario contained a very large number of highly delayed flights, due to the larger number of flight copies required.

Table 4 summarizes our results over the 60 scenarios. We report the *simulated* values of the passenger disruption savings, delayed passenger costs to the airline, fuel savings per operated long-haul (LH) flight in T, total cost savings, short-haul (SH) flights intentionally held, number of flights

		Baseline recovery	Speed change	Pax Centric $\Theta = 0$	Pax Centric $\Theta = 5$	Pax Centric $\Theta = 10$	Pax Centric $\Theta = 15$	Enhanced Recovery $\Theta = 0$	Enhanced Recovery $\Theta = 5$	Enhanced Recovery $\Theta = 10$	Enhanced Recovery $\Theta = 15$
Pax disruption savings %	average	–	6.06	20.29	38.64	55.64	60.79	65.91	76.53	77.96	79.14
Delayed pax cost savings to airline %	average	–	-15.40	26.43	41.09	53.58	57.90	64.26	77.77	78.52	79.56
Fuel savings per operated LH flight %	average	–	-0.082	-0.158	-0.163	-0.163	-0.167	-0.254	-0.252	-0.251	-0.249
Total cost savings to airline %	average	–	4.95	1.29	2.26	2.48	2.86	9.18	9.02	9.01	9.00
SH flights held	average	–	0.00	0.00	4.00	6.20	7.23	0.00	3.52	4.57	5.45
Nr. canceled	average	1.60	1.23	0.57	0.53	0.53	0.50	0.10	0.10	0.10	0.10
Nr. swaps	average	3.75	3.77	4.00	3.80	3.80	3.83	3.47	3.40	3.33	3.33
OTP	average	0.88	0.92	0.90	0.90	0.90	0.90	0.94	0.94	0.94	0.94
MIP gap %	average	0.00	0.00	0.00	0.00	0.00	0.33	7.13	7.21	7.12	7.87

Table 4 Simulated savings per day for different recovery strategies and holding policies Θ, over all scenarios

canceled, number of swaps, on-time performance (OTP) for the long-haul fleets, and MIP gap for different models of interest.

In Table 4 we show that relative to the (traditional) sequential recovery model, the enhanced recovery models perform better than: (i) the model that uses speed changes alone (column 2); and (ii) passenger-centric recovery models (columns 3-6) without flight speed changes but with different levels of flight holds Θ. The enhanced recovery models essentially capture the flexibility exhibited by both classes of models and thus exhibit superior performance.

These results show an improvement of 66-79% in the number of passengers disrupted, compared to a sequential recovery approach. Passenger delay costs (incurred by the airline) can be decreased by 64-79%. This comes at a cost of additional fuel burn costs of about 0.25% per operated flight, and results in an overall cost savings of about 9% for the airline. The enhanced recovery models with flight planning also decrease significantly the number of flight cancelations (from 1.6 to 0.1 per scenario) by allowing an increased number of swaps, thereby increasing flexibility in recovery. Also, the possibility of speed-ups and the decreased numbers of cancelations help improve the on-time performance of the airline from 0.88 to 0.94 for the long-haul fleet. However, we also observe a high degree of variability in the savings from our enhanced recovery models. This follows from our discussion in §7.1 concluding that different delay scenarios and different levels of passenger connectivity result in very different trade-offs in fuel and delay costs (Figures 6 and 7). Thus, to gain further insight, we present Tables 5 - 8 detailing the performances of the different models for different levels of disruption.

From these tables, we conclude the following:

		Baseline recovery	Speed change	Pax Centric $\Theta = 0$	Pax Centric $\Theta = 5$	Pax Centric $\Theta = 10$	Pax Centric $\Theta = 15$	Enhanced Recovery $\Theta = 0$	Enhanced Recovery $\Theta = 5$	Enhanced Recovery $\Theta = 10$	Enhanced Recovery $\Theta = 15$
Pax disruption savings %	average	–	0.00	0.00	48.01	98.05	98.05	82.18	98.79	98.79	98.79
Delayed pax cost savings to airline %	average	–	-108.70	23.01	39.94	58.87	62.31	57.84	67.69	67.69	68.53
Fuel savings per operated LH flight %	average	–	-0.007	0.000	0.000	0.000	0.000	-0.007	-0.007	-0.007	-0.007
Total cost savings to airline %	average	–	0.03	0.00	0.47	1.36	1.36	1.64	1.67	1.66	1.72
SH flights held	average	0.00	0.00	0.00	3.20	4.20	4.20	0.00	1.40	1.40	1.40
Nr. canceled	average	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Nr. swaps	average	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
OTP	average	0.93	1.00	0.93	0.93	0.93	0.93	1.00	1.00	1.00	1.00
MIP gap %	average	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Table 5 Small delay scenarios: Simulated savings per day for different recovery strategies and different holding policies Θ

		Baseline recovery	Speed change	Pax Centric $\Theta = 0$	Pax Centric $\Theta = 5$	Pax Centric $\Theta = 10$	Pax Centric $\Theta = 15$	Enhanced Recovery $\Theta = 0$	Enhanced Recovery $\Theta = 5$	Enhanced Recovery $\Theta = 10$	Enhanced Recovery $\Theta = 15$
Pax disruption savings %	average	–	0.00	0.00	35.01	70.38	77.07	61.40	86.76	86.76	86.76
Delayed pax cost savings to airline %	average	–	-60.91	0.00	32.07	57.80	68.28	57.15	85.69	85.69	85.69
Fuel savings per operated LH flight %	average	–	-0.013	0.000	0.000	0.000	0.000	-0.014	-0.013	-0.013	-0.013
Total cost savings to airline %	average	–	0.04	0.00	-0.01	-0.08	-0.11	0.35	0.10	-0.08	-0.08
SH flights held	average	0.00	0.00	0.00	2.83	4.58	4.92	0.00	2.58	2.58	2.58
Nr. Canceled	average	0.08	0.08	0.08	0.08	0.08	0.08	0.08	0.08	0.08	0.08
Nr. Swaps	average	5.83	4.83	5.83	5.83	5.83	5.83	4.83	4.83	4.83	4.83
OTP	average	0.95	1.00	0.95	0.95	0.95	0.95	1.00	1.00	1.00	1.00
MIP gap %	average	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Table 6 Medium delay scenarios: Simulated savings per day for different recovery strategies and different holding policies Θ

Small disruptions: Our enhanced recovery models can recover a significant fraction (82-98%) of disrupted passengers. For small disruptions these are few in number, so most can be recovered using speed changes and by holding downstream flight departures. A very small increase in fuel burn costs of 0.007% is observed due to speed increases (and speed reductions, where appropriate). The result is an overall cost savings of 1.64 - 1.72 % for the small delay scenarios.

Medium disruptions: A significant number of passengers (61-86%) are prevented from missing their connections. The fraction is not as large, however, as in the case of small-disruptions. Compared to the small-disruption case, a larger number of passengers are negatively impacted by the

		Baseline recovery	Speed change	Pax Centric $\Theta = 0$	Pax Centric $\Theta = 5$	Pax Centric $\Theta = 10$	Pax Centric $\Theta = 15$	Enhanced Recovery $\Theta = 0$	Enhanced Recovery $\Theta = 5$	Enhanced Recovery $\Theta = 10$	Enhanced Recovery $\Theta = 15$
Pax disruption savings %	average	–	0	5.53	23.78	45.53	55.84	57.41	72.03	76.74	81.15
Delayed pax cost savings to airline %	average	–	0.6	6.58	21.17	37.03	43.89	46.72	66.85	69.8	74.17
Fuel savings per operated LH flight %	average	–	-0.008	-0.011	-0.011	-0.011	-0.011	-0.032	-0.032	-0.033	-0.028
Total cost savings to airline %	average	–	0	0	0.11	0.56	0.44	1.47	1.24	1.09	1.15
SH flights held	average	0.00	0.00	0.00	5.57	10.29	12.00	0.00	5.21	6.00	7.79
Nr. Canceled	average	0.14	0.14	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Nr. Swaps	average	3.75	3.79	3.86	3.57	3.57	3.57	3.57	3.57	3.57	3.57
OTP	average	0.90	0.93	0.91	0.91	0.91	0.91	0.94	0.94	0.94	0.93
MIP gap %	average	0.00	0.00	0.00	0.00	0.00	0.81	3.41	2.36	2.03	1.72

Table 7 **Large delay scenarios: Simulated savings per day for different recovery strategies and different holding policies** Θ

		Baseline recovery	Speed change	Pax Centric $\Theta = 0$	Pax Centric $\Theta = 5$	Pax Centric $\Theta = 10$	Pax Centric $\Theta = 15$	Enhanced Recovery $\Theta = 0$	Enhanced Recovery $\Theta = 5$	Enhanced Recovery $\Theta = 10$	Enhanced Recovery $\Theta = 15$
Pax disruption savings %	average	–	12.53	39.31	45.69	47.11	50.01	69.08	70.63	71.31	71.64
Delayed pax cost savings to airline %	average	–	11.79	51.52	56.30	56.89	58.16	77.18	78.07	78.21	78.24
Fuel savings per operated LH flight %	average	–	-0.159	-0.322	-0.332	-0.332	-0.341	-0.504	-0.500	-0.497	-0.496
Total cost savings to airline %	average	–	10.21	2.66	4.65	4.90	5.76	18.11	18.01	18.15	18.10
SH flights held	average	0.00	0.00	0.00	3.86	5.24	6.41	0.00	3.45	5.24	6.21
Nr. Canceled	average	3.21	2.45	1.14	1.07	1.07	1.00	0.17	0.17	0.17	0.17
Nr. Swaps	average	3.38	3.97	4.00	3.72	3.72	3.79	3.45	3.31	3.17	3.17
OTP	average	0.83	0.87	0.87	0.87	0.87	0.88	0.91	0.91	0.91	0.91
MIP gap %	average	0.00	0.00	0.00	0.00	0.00	0.29	13.11	13.79	13.75	15.46

Table 8 **Very large delay scenarios: Simulated savings per day for different recovery strategies and different holding policies** Θ

medium-disruptions, and fewer affected connections can be re-connected using flight speed changes and departure holds. As expected, higher speed ups are required compared to the small-disruption cases, thus burning more fuel (0.014 %) on average. The higher speed up per LH flight, combined with the fewer passengers saved from misconnecting cause the total cost savings *for the airline* to be small, as to be almost negligible in these cases. However, significant benefits are observed in costs experienced by passengers as a significant number are saved from misconnecting.

Large and very-large disruptions: A phenomenon different from those in small- and medium-disruption cases comes into play. Large disruptions propagate downstream in the long-haul schedule, affecting many flights and passengers, and requiring flight speed-ups for a large number of

flights (0.033% for large disruptions and 0.5% for very large disruptions). Flight speed changes in enhanced recovery enable a larger number of swap possibilities in the network as schedule recovery is performed, and the added flexibility also decreases the number of cancelations that might be needed in a traditional recovery model. Due to this added flexibility, a significant percentage of passengers (57-81% for large delay and 69-71% in very-large delay scenarios) can be prevented from missing their connections. The result is a total cost savings of 1.1-1.4% for large-delay and 18% for very large-delay scenarios.

Due to the high amount of delay propagation that occurs in the large and very large disruption scenarios, propagation to several flights has to be modeled, and a longer time-window T may have to be used, along with more flight copies, as needed. Because of the two minute limit on solution time, the larger-size integer program solutions are not optimal, with MIP gaps shown in Tables 7 and 8. Nonetheless, the solutions to our enhanced recovery models achieve significant improvements relative to the sequential recovery, speed change and passenger-centric models.

From tables 4, 5, 6, 7, 8, we see that flight speed changes alone with no passenger-centric costs do not add much value. Once passenger centric-costs are added into the objective, the models add much more value, as seen in the columns 'Enhanced Recovery $\Theta = 0$'. Moreover, enhanced recovery models perform better than those with passenger-centric costs and objectives for all values of Θ. The improvement is greatest for small values of Θ(such as 0 and 5), that is, when downstream passenger connections are not allowed to be held or are held only briefly. This occurs because the enhanced recovery model allows speed changes, while the passenger-centric models do not. Our enhanced recovery models also add progressively more value relative to passenger-centric recovery as the size of the initial disruption increases and more flights are impacted.

Neither speed changes alone (column 4 in Tables 4, 5, 6, 7, 8), nor passenger-centric recovery with flight holds alone (columns 5 through 8 in the tables) can generate savings as great as those of the enhanced recovery models (columns 9-12 in the tables) in which these mechanisms are combined.

To put these results in context, we weight the average savings from Tables 5 - 8 by the frequency of occurrences of such delays described in Table 3. This results in a total cost savings to the airline of about 5.7 - 5.9%, decrease in passenger disruptions of 66 - 83% and increase in fuel burn over the long-haul fleet of 0.152-0.155%. The decrease in passenger-related cost savings to the airline is about 60 - 73%. Though the cost savings to the airline are highest for very-large-delay cases and are of the order of 1-2% for other types of scenarios, in all types of scenarios, there are significant savings in passenger misconnections and disruptions. Savings in delay minutes experienced by passengers (computed using the evaluation module) are of the order of 474,823 - 485,254 minutes

per disrupted scenario compared to baseline recovery, resulting in decreased passenger-delay costs of \$17.5 - 17.9 M over the June-July period in 2008 (using a passenger-value-of-time equal to \$37.56 per hour (Air Transport Association 2009)).

7.3. Summary

Our enhanced recovery approach integrating flight planning into disruption management allows the capture of speed changes of flights and captures the interaction between fuel burn and delay costs. We show significant synergy between flight speed changes and existing mechanisms of disruption management, such as flight holds. Compared to the current state-of-the-practice at airlines, our enhanced recovery approach provides a more accurate way of dynamically quantifying the tradeoff between time-related costs and fuel-burn related costs. The rule-of-thumb policy in practice is almost always costlier than our enhanced recovery approach. Our enhanced recovery models reduce total costs and passenger-related delay costs for the airline, compared to existing approaches in practice and in the literature. For very low disruption levels, fuel costs dominate over passenger-related delay costs to the airline, and as the extent of disruption increases, passenger-related delay costs dominate. The ability to change flight speeds also contributes to higher cost savings for the airline as the extent of the initial disruption increases. Flight speed changes in combination with flight holds mitigate delay propagation effects by providing more aircraft swap opportunities and decreasing the number of required cancelations.

References

Air Transport Association. 2009. `http://www.airlines.org/Economics/DataAnalysis/Pages/CostofDelays.aspx`.

Altus, S. 2007. Flight planning: The forgotten field in airline operations. `http://www.agifors.org/studygrp/opsctl/2007/`, Presented at AGIFORS Airline Operations.

Altus, S. 2010. Dynamic cost index management in flight planning and flight following. `http://www.agifors.org/studygrp/opsctl/2010/program.html`, Presented at AGIFORS Airline Operations.

Andersson, T., P. Varbrand. 2004. The flight perturbation problem. *Transportation Planning and Technology* **27**(2) 91–117.

Associated Press. May 1, 2008. `http://www.msnbc.msn.com/id/24410809/`.

Associated Press. May 2, 2008. `http://www.usatoday.com/travel/flights/2008-05-02-slow-fuel_N.htm`.

Barnhart, C. 2009. Irregular operations: Schedule recovery and robustness. P. Belobaba, A. Odoni, C. Barnhart, eds., *The Global Airline Industry*, chap. 9. John Wiley and Sons, United Kingdom.

Barnhart, C., M. Ball, G. Nemhauser, A. Odoni. 2006. Air transportation: Irregular operations and control. G. Laporte, C. Barnhart, eds., *Transportation Handbook II*.

Barnhart, C., T. Kniker, Lohatepanont M. 2002. Itinerary-based airline fleet assignment. *Transportation Science* **36** 199–217.

Bratu, S., C. Barnhart. 2006. Flight operations recovery: New approaches considering passenger recovery. *Journal of Scheduling* **9**(3) 279–298. doi:http://dx.doi.org/10.1007/s10951-006-6781-0.

Burrows, G., C. A. Brown, T. W. Thom, J. M. C. King, J. Frearson. 2001. Real-time cost management of aircraft operations. *Management Accounting Research* **12** 281–298.

Clausen, J., A. Larsen, J. Larsen, N. J. Rezanova. 2010. Disruption management in the airline industry concepts, models and methods. *Computers and Operations Research* **37** 809–821.

Davis, N., J. Calusens, S. Tiourine. 2002. Descartes project report. (Private correspondence).

Dienst, D. 2010. Airline disruption management - exact methods for the aircraft recovery problem. Masters thesis, DTU Technical University of Denmark.

European Central Bank. 2010. Exchange Rates. `http://www.ecb.int/stats/exchange/eurofxref/html/eurofxref-graph-usd.en.html`.

International Air Transport Association. 2010. Jet Fuel Price Development. `http://www.iata.org/whatwedo/economics/fuel_monitor/Pages/price_development.aspx`.

Jeppesen. 2010. Jeppesen technology services JetPlan engine user manual. (Private correspondence).

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

Thengvall, B. G., G. Yu, J. F. Bard. 2001. Multiple fleet aircraft schedule recovery following hub closures. *Transportation Research Part A: Policy and Practice* **35**(4) 289 – 308. doi: DOI:10.1016/S0965-8564(99)00059-2. URL `http://www.sciencedirect.com/science/article/B6VG7-423RH77-2/2/a322cc8ad025ea017916916190f78138`.

Thengvall, Benjamin G., Gang Yu, Jonathan F. Bard. 2000. Balancing user preferences for aircraft schedule recovery during irregular operations. *IIE Transactions* **32**(3) 181–193.

Vaaben, B. 2009. Disruption management passenger re-accommodation strategies. `http://www.agifors.org/studygrp/opsctl/2009/`, Presented at AGIFORS Airline Operations.

Yan, S., A. Young. 1996. A decision support framework for multi-fleet routing and multi-stop flight scheduling. *Transportation Research-A* **30**(5) 379–398.

Yu, G., X. Qi. 2004. *Disruption Management: Framework, Models And Applications*. World Scientific Publishing Co., Singapore.

Airline schedule recovery involves making decisions during operations to minimize additional operating costs while getting back on schedule as quickly as possible. The mechanisms used include aircraft swaps, flight cancelations, crew swaps, reserve crews and passenger re-bookings. In this context, we introduce another mechanism, namely flight planning, which enables flight speed changes. Flight planning is the process of determining flight plan(s) specifying the route of a flight, its speed and its associated fuel burn. Our key idea in integrating flight planning and disruption management is to adjust the speeds of flights during operations, trading off flying time and fuel burn, and combining with existing mechanisms such as flight holds; all with the goal of striking the right balance of fuel costs and passenger-related delay costs incurred by the airline.