Evaluating Air Traffic Flow Management in a Collaborative Decision-Making Environment

Douglas Fearing (corresponding author)

Affiliation: Technology and Operations Management, Harvard Business School,
Address: Harvard Business School, Morgan Hall 491, Soldiers Field Road, Boston, MA 02136
Phone: 617-495-6614
Email: dfearing@hbs.edu

Cynthia Barnhart

Affiliation: Department of Civil and Environmental Engineering, Massachusetts Institute of Technology
Address: Massachusetts Institute of Technology, Room 2-106, 77 Massachusetts Avenue, Cambridge, MA 02139
Phone: 617-253-3815
Email: cbarnhart@mit.edu

Effective word count: 7487 (6237 text words + 1 figure + 4 tables)

Revision submitted on August 30th, 2010.

Abstract: The Collaborative Decision-Making (CDM) framework introduced into Ground Delay Programs (GDPs) in the late 1990s is an integral component in the Federal Aviation Administration's (FAA's) Traffic Flow Management (TFM) procedures. CDM allows the FAA to act as a mediator when managing TFM programs, pushing as much decision-making as possible to individual airlines. Though this approach has been highly successful in practice, it creates a new question for the research community – how should proposed enhancements to TFM be evaluated in a CDM environment? In our paper, we develop a sequential evaluation procedure to address this question. Our procedure includes airline disruption responses and a quasi-compression operation, attempting to mimic the three-stage CDM process. To model airline disruption responses, we develop an integer optimization model that balances operational and passenger considerations to make determinations on which flights to cancel, swap, or delay. We demonstrate the value of our procedure by analyzing an optimization-based TFM approach in the CDM environment.

In this paper, we develop a sequential evaluation procedure to evaluate air traffic flow management approaches in a collaborative decision-making (CDM) environment. After allocating traffic flow management (TFM) capacities, the sequential evaluation procedure i) uses an airline disruption response model to approximate each airline's flight cancelations and swaps, ii) performs a quasi-compression operation on the combined schedule, and iii) calculates passenger delays based on the result.

In Section 1, we provide a brief review of related literature. Section 2 describes the components of the sequential evaluation procedure, including the airline disruption responses and quasi-compression. Section 3 describes how we construct the scenarios used in our analysis, including the estimation of planned aircraft routings and passenger itineraries. Last, in Section 4, we use the approach to analyze the benefits of an optimization-based TFM allocation approach in a CDM environment.

1 Literature Review

There have been many publications related to Collaborative Decision-Making (CDM). The CDM paradigm was developed through a joint effort involving the FAA, the airline community, Metron Aviation Inc., the Volpe National Transportation Systems Center, and the academic community. The CDM approach was first incorporated into the FAA's Ground Delay Program (GDP) in the late 1990s and later incorporated into Airspace Flow Programs (AFPs) for managing en route congestion (*1*). Chang, et al. provides a good description of the highly successful CDM approach (*2*). Ball, et al. provides an efficient dual network flow formulation for the static-stochastic GDP and show how this formulation can be incorporated into a CDM framework with equity considerations (*3*). Kotnyek and Richetta (*4*) extend the Richetta and Odoni (*5*) static-stochastic formulation to incorporate equity consideration. In the process, they demonstrate that the Ball et al. model can be described as a special case. Vossen and Ball extend an integer programming formulation for the GDP to approximate the CDM process (*6*). They also describe how the compression step of the CDM process can be considered a mediated bartering step between airlines. In compression, the FAA performs a sequence of 1-to-1 exchanges to fill in gaps created by cancelations and reward the airlines that canceled the flights.

The dynamic nature of CDM makes it difficult to evaluate the impact of proposed TFM enhancements, because efficiencies gained during allocation could be lost in subsequent airline responses and compression. This is especially true in a multi-resource setting where it is not clear how compression would or should be applied. There are many papers that have evaluated the impacts of proposed multi-resource TFM enhancements in terms of their impact on allocated flight delays (e.g., Bertsimas and Patterson (*7*), Andreatta, Brunetta, and Guastalla (*8*), Bertsimas, Lulli, and Odoni (*9*), and Barnhart, et al. (*10*)). We believe this paper represents the first attempt to characterize the impact of the CDM process in the context of a multi-resource optimization-based allocation.

2   Sequential Evaluation Procedure

The sequential evaluation procedure we have developed is depicted in Figure 1. The procedure takes as input an allocated schedule based on one or more TFM programs. From this schedule, each airline is allotted capacity (or slots) for TFM resources (i.e., airports or flow constrained areas (FCAs)). Capacities are assumed to be infinite for resources with no TFM program in place. To attempt to mimic airline operations, we group sub-contracted carriers with their parent carriers. The full list of the 15 airline groupings we consider is provided in Table 1. These groupings are determined based on the prevalence of corresponding multi-carrier itineraries in the Airline Origin and Destination Survey (DB1B) from the Bureau of Transportation Statistics (*11*).
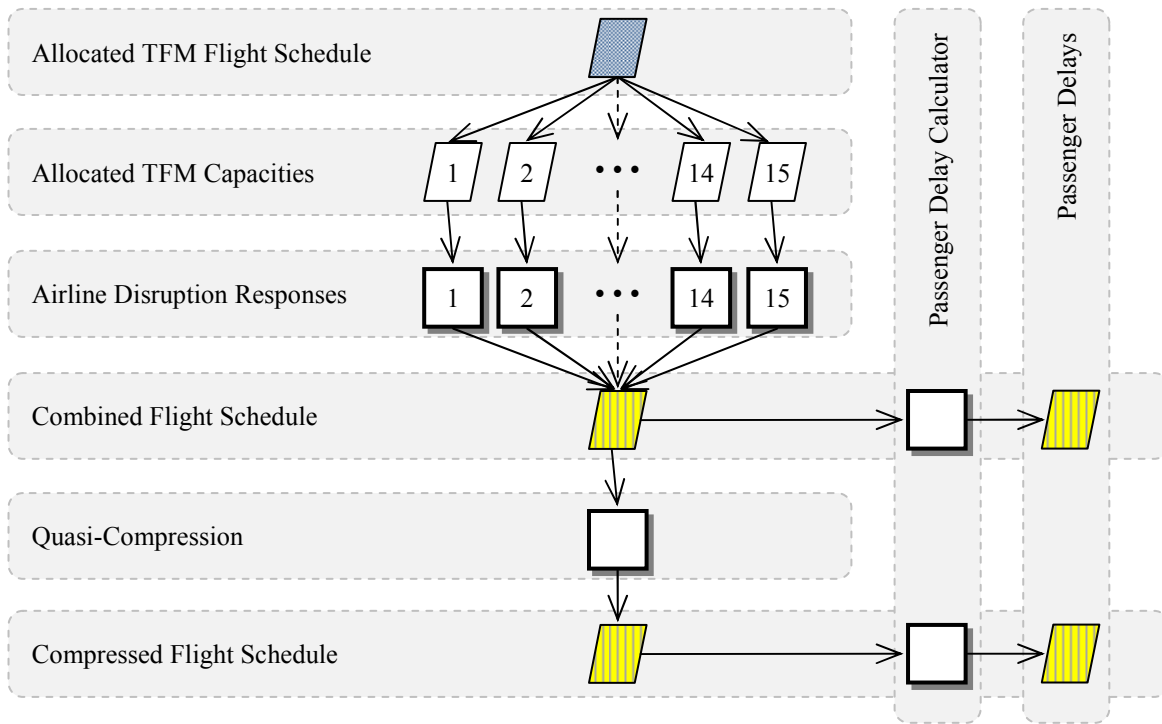
Figure 1: Process workflow for sequential evaluation procedure with input in blue and outputs in yellow

To approximate airline disruption responses, we utilize a modified TFM scheduling model, which is described in Section 2.2. The output of this model is a feasible airline flight schedule, incorporating flight cancelations and maintaining aircraft flow balance. Because allocated capacities are derived from a feasible controlled schedule, the combined airline schedules form a feasible schedule relative to the initial TFM programs. The combined flight schedule and its calculated passenger delays are two of the outputs from our sequential evaluation procedure. We then perform a quasi-compression to fill in gaps created by flight cancelations, ensuring that no flight receives more delay in the compressed schedule. The details of and motivation for the quasi-compression procedure are described in Section 2.4. The compressed flight schedule and its calculated passenger delays form the final two outputs of our sequential evaluation procedure. Before describing the individual components of the sequential evaluation procedure, we define the required data and notation in Section 2.1.

Airline Codes	Airline Names
AA, MQ	American Airlines & American Eagle Airlines
DL, EV, OH	Delta Airlines, Atlantic Southeast Airlines, & Comair
CO, XE	Continental Airlines & ExpressJet Airlines
NW, 9E	Northwest Airlines & Pinnacle Airlines
AS	Alaska Airlines
AQ	Aloha Airlines
B6	JetBlue Airways
F9	Frontier Airlines
FL	AirTran Airways
HA	Hawaiin Airlines
OO	SkyWest Airlines
UA	United Airlines
US	US Airways
WN	Southwest Airlines
YV	Mesa Airlines

Table 1: Airline groupings for sequential evaluation procedure

2.1   Data and Notation

The notation used for the airline disruption response model builds on the notation defined in (*10*):

$$\mathcal{T} \quad = \quad \text{set of discrete time intervals;}$$

$$\mathcal{R} \quad = \quad \text{set of capacity-controlled resources;}$$

$$b_{rt} \quad = \quad \text{capacity of resource } r \text{ over time interval } t;$$

$$\mathcal{F} \quad = \quad \text{set of flights;}$$

$$|f| \quad = \quad \text{number of steps in controlled flight plan for flight } f;$$

$$\mathcal{I}(f) \quad = \quad \text{set of step indices in controlled flight plan for flight } f;$$

$$r(f,i) \quad = \quad \text{resource required by flight step } i \text{ for flight } f;$$

$$\alpha(f,i) \quad = \quad \text{earliest start time for flight step } i \text{ for flight } f; \text{ and}$$

$$\delta(f,i) \quad = \quad \text{processing time of flight step } i \text{ for flight } f.$$

Because each airline's disruption response is determined independently, there is no need to index these terms by airline. For example, when determining the disruption response for Southwest, $\mathcal{F}$ will refer to all of the Southwest flights, and b_{rt} will refer to the capacity allocated to Southwest for resource r and time interval t. For each flight, we extend $\mathcal{I}(f)$ to include a departure index, 1, and an arrival index, $|f|$.

In addition to the above, the airline disruption response depends on planned aircraft routings and booked passenger itineraries. We specify the aircraft type for each scheduled flight, $p(f) \in \mathcal{P}$, where $\mathcal{P}$

represents the set of aircraft types for the airline. The aircraft type is a combination of the operating carrier and International Civil Aviation Organization (ICAO) aircraft code for the flight. Each planned aircraft routing, $s \in \mathcal{S}$, is then represented as a sequence of flights an aircraft is scheduled to fly, $[f_1, \ldots, f_n]$, where $[s]_1$ represents the first flight and $[s]_{(0)}$ the last. The minimum aircraft turnaround time (i.e., the minimum time between arrival and subsequent departure) is assumed to be constant across all airports and aircraft types, and is denoted by m. To maintain aircraft flow balance, we also specify the origin and destination of each flight, $orig(f)$ and $dest(f)$ respectively, letting $\mathcal{A}$ represent the set of all airports.

Passenger itineraries are either non-stop, represented by a single flight f, or one-stop, represented by a pair of flights (f_1, f_2). Each flight $f \in \mathcal{F}$ represents a valid non-stop itinerary, and we let the set OS represent the set of booked one-stop itineraries, where at least one flight in each flight pair is in $\mathcal{F}$. We do not require both flights to be in $\mathcal{F}$, because one flight may be operated by another airline. For these cases, we need to know the planned connection time between the arrival of flight f_1 and the departure of flight f_2, $w(f_1, f_2)$, which we compare to the minimum connection time, $w_{\min}$, to determine a missed connection. We let $np(f)$ represent the number of non-stop passengers booked on flight f and $np(f_1, f_2)$ represent the number of passengers booked on the one-stop itinerary represented by the pair of flights $(f_1, f_2) \in OS$. The number of seats on flight f is represented by $ns(f)$.

If a passenger itinerary is disrupted, the estimated delay associated with re-accommodation depends on the itinerary and the type of disruption. For non-stop itineraries, the only type of disruption we consider is a flight cancelation. Thus, we specify the re-accommodation delay associated with a flight cancelation for non-stop passengers booked on flight f as $rd(f)$. For one-stop itineraries, the possible types of disruptions are first flight cancelation, second flight cancelation, or missed connection; we specify the re-accommodation delays as $rd^1(f_1, f_2)$, $rd^2(f_1, f_2)$, and $rd^x(f_1, f_2)$ respectively.

2.2 Ration-by-Schedule Exponential Penalty Model

In this section, we describe the deterministic, multi-resource Ration-by-Schedule Exponential Penalty (RBS-EP) formulation that provides the foundation for the airline disruption response model. Further details regarding the development of the RBS-EP model can be found in (**10**).

2.2.1 Decision Variables

$$y_{fit} = \begin{cases} 1, & \text{if flight plan step } i \text{ for flight } f \text{ has started by time } t; \text{ and} \\ 0, & \text{otherwise.} \end{cases} \tag{1}$$

2.2.2 Constraints

$$y_{fit} \leq y_{fi(t+1)} \qquad\qquad \forall f \in \mathcal{F}, \forall i \in \mathcal{I}(f), \forall t \in \{0 \ldots T - 2\}; \quad (2)$$

$$y_{fi(T-1)} = 1 \qquad\qquad \forall f \in \mathcal{F}, \forall i \in \mathcal{I}(f); \quad (3)$$

$$y_{fi(\alpha(f,i)-1)} = 0 \qquad\qquad \forall f \in \mathcal{F}, \forall i \in \mathcal{I}(\mathrm{F}) \text{ s.t. } \alpha(f,i) > 0; \quad (4)$$

$$y_{f(i+1)t} = y_{fi(t-\alpha(f,i+1)+\alpha(f,i))} \quad \forall f \in \mathcal{F}, \forall i \in \mathcal{I}(f) \setminus \{|f|\}; \quad (5)$$

$$\sum_{\{(f,i): r(f,i)=\hat{r}\}} \left(y_{fit} - y_{fi(t-\delta(f,i))} \right) \leq b_{\hat{r}t} \qquad\qquad \forall \hat{r} \in \mathcal{R}, \forall t \in \mathcal{T}. \quad (6)$$

2.2.3 Objective Function

In the RBS-EP model, the cost of each additional interval of delay grows exponentially beyond a threshold determined by independent first-scheduled, first-served allocations for each TFM-controlled resource. To express the objective function, we let $d^{RBS}(f, i)$ be the delay that would be assigned to step i of flight f based on independent allocations. The following two equations define the cost, c_{ft}, associated with the last step of flight f being scheduled at time t.

$$d^{RBS}_{MAX}(f) = \max_{i \in \mathcal{I}(f)} \{d^{RBS}(f, i)\}; \text{ and} \qquad (7)$$

$$c_{ft} = \min\{t - \alpha(f, |f|), d^{RBS}_{MAX}(f)\} + \sum_{\epsilon=1}^{t-\alpha(f,|f|)-d^{RBS}_{MAX}(f)} \lambda^{\epsilon}. \qquad (8)$$

In equation (8), the base of the exponent, λ, controls the trade-off between fairness and delay. Using this notation, the RBS-EP objective function is:

$$\min \sum_{f \in \mathcal{F}} \left[\sum_{t=\alpha(f,|f|)}^{T-1} c_{ft} \left(y_{f|f|t} - y_{f|f|(t-1)} \right) \right]. \qquad (9)$$

2.3 Airline Disruption Response Model

In the airline disruption response model, we cancel and reschedule flights using an integer optimization model extended from the RBS-EP formulation. The decisions on flight schedules and cancelations are based on a multi-part objective function that approximates both operational and passenger considerations. Additionally, using planned aircraft routings, the model ensures that flow balance is maintained for each airline, airport, and aircraft type.

2.3.1 Decision Variables

For each one-stop itinerary $(f_1, f_2) \in \mathcal{OS}$ where both flights are in $\mathcal{F}$, we introduce two binary variables to represent whether the itinerary is disrupted due to the second flight being canceled or due to a missed

connection. Additionally, we introduce a linear decision variable to measure the delay when no disruptions occur.

$$z_{f_1 f_2} = \begin{cases} 1, & \text{if itinerary } (f_1, f_2) \text{ is disrupted due to flight } f_2 \text{ being canceled; and} \\ 0, & \text{otherwise.} \end{cases} \tag{10}$$

$$x_{f_1 f_2} = \begin{cases} 1, & \text{if itinerary } (f_1, f_2) \text{ is disrupted due to a missed connection; and} \\ 0, & \text{otherwise.} \end{cases} \tag{11}$$

$$d_{f_1 f_2} \geq 0 \tag{12}$$

2.3.2   Constraints

The first change to the RBS-EP model is allowing cancelations by replacing constraints (3) with the following to ensure that if one step of a flight is canceled, then so are all others.

$$y_{fi(T-1)} = y_{f(i+1)(T-1)} \quad \forall f \in \mathcal{F}, \forall i \in \mathcal{I}(f) \setminus \{|f|\}. \tag{13}$$

Additionally, we add constraints to ensure that for each airport, $a \in \mathcal{A}$, and aircraft type, $p \in \mathcal{P}$, flow balance is maintained. To simplify exposition, in equation (14), we define $\Delta(a, p, t)$ as the difference between the number of flights of aircraft type p that arrive at airport a by time $(t - m)$ and the number of flights of aircraft type p that depart from airport a by time t.

$$\Delta(a, p, t) \triangleq \sum_{\{f \in \mathcal{F}: dest(f)=a, p(f)=p\}} y_{f1(t-m)} - \sum_{\{f \in \mathcal{F}: orig(f)=a, p(f)=p\}} y_{f|f|t}. \tag{14}$$

Note that if airport a has no aircraft of type p available at the beginning of the model time period, $\Delta(a, p, t)$ equals the number of these aircraft available for departure. Because the model time period ends at time T, we use $\Delta(a, p, T + m)$ to specify the difference between arrivals and departures at the end of the model time period.

For each aircraft routing, s, we assume there is an aircraft of type $p([s]_1)$ available for the first flight in the sequence at its scheduled departure time, $\alpha([s]_1, 1)$, and origin airport, $orig([s]_1)$. Additionally, for each aircraft routing, we require an aircraft of type $p([s]_{(0)})$ at the destination airport of the last flight, $dest([s]_{(0)})$, at the end of the model time period. Thus, we define $\underline{n}(a, p, t)$ as the cumulative number of aircraft of type p available at airport a by time t and $\overline{n}(a, p)$ as the number of aircraft of type p required at airport a at the end of the model time period.

$$\underline{n}(a, p, t) \triangleq |\{s \in \mathcal{S}: orig([s]_1) = a, p([s]_1) = p, \alpha([s]_1, 1) \leq t\}|; \text{ and} \tag{15}$$

$$\overline{n}(a, p) \triangleq |\{s \in \mathcal{S}: dest([s]_{(0)}) = a, p([s]_0) = p\}|. \tag{16}$$

Using these definitions, we add the following constraints to the model ensure flow balance.

$$\Delta(a,p,t) + \underline{n}(a,p,t) \quad \geq \quad 0 \quad \forall a \in \mathcal{A}, \forall p \in \mathcal{P}, \forall t \in \mathcal{T}. \tag{17}$$

$$\Delta(a,p,T+m) + \underline{n}(a,p,T) \quad = \quad \overline{n}(a,p) \quad \forall a \in \mathcal{A}, \forall p \in \mathcal{P}. \tag{18}$$

In addition to forcing delay propagation through the airline network, constraints (17) and (18) guarantee that flight cancelations occur in cycles, and thus that no deadheading (flying empty) is required to reposition aircraft for future operations.

Next, we add constraints to ensure that each of the decision variables defined in Section 2.3.1 takes the appropriate value. The first constraints ensure that we consider a one-stop itinerary to be disrupted by a second flight cancelation only if the first flight is not canceled.

$$z_{f_1 f_2} + y_{f_2|f_2|(T-1)} \geq y_{f_1|f_1|(T-1)} \quad \forall (f_1, f_2) \in \mathcal{OS}: f_1, f_2 \in \mathcal{F}. \tag{19}$$

In our minimization objective, the cost coefficient for $z_{f_1 f_2}$ will be strictly positive, which ensures that, in an optimal solution, $z_{f_1 f_2}$ will equal 1 if and only if $y_{f_1|f_1|(T-1)} = 1$ and $y_{f_2|f_2|(T-1)} = 0$, as desired.

The next constraints ensure that we consider a one-stop itinerary to be disrupted by a missed connection only if both flights are scheduled and the time between the arrival of the first flight and departure of the second is less than the minimum connection time, $w_{\min}$.

$$y_{f_2 1 t} \leq y_{f_1|f_1|(t-w_{\min})} + 1 - y_{f_1|f_1|(T-1)} + x_{f_1 f_2} \quad \forall (f_1, f_2) \in \mathcal{OS}: f_1, f_2 \in \mathcal{F}, \forall t \in \mathcal{T}. \tag{20}$$

That is, constraints (20) ensure that flight f_2 cannot depart by time t unless 1) flight f_1 arrives by time $(t - w_{\min})$, 2) flight f_1 is canceled, or 3) $x_{f_1 f_2}$ is set to 1. As with $z_{f_1 f_2}$, the cost coefficient for $x_{f_1 f_2}$ will be strictly positive, ensuring that, in an optimal solution, $x_{f_1 f_2}$ will equal 1 if and only if flights f_1 and f_2 are scheduled and the connection time between the two is less than $w_{\min}$.

Last, we add constraints to ensure that the flight delay associated with a one-stop itinerary is positive only if the itinerary is not disrupted. To specify this constraint, we first define expressions to calculate whether an individual flight is canceled, and if not, the associated delay. We can calculate the cancelation expression as:

$$cx(f) = (1 - y_{f|f|(T-1)}). \tag{21}$$

We consider the delay for flight f to be 0 if the flight is canceled, thus we calculate the delay as:

$$d(f) = \big(T - \alpha(f, |f|)\big) y_{f|f|(T-1)} - \sum_{t=\alpha(f,|f|)}^{T-1} y_{f|f|t}. \tag{22}$$

Note that $y_{f|f|t} = 0$ for all t when flight f is canceled, therefore $d(f) = 0$ in this case, as desired. When flight f is scheduled, $y_{f|f|T-1}$ is guaranteed to equal 1 based on constraints (2). In this case, $d(f)$ counts the number of intervals that $y_{f|f|t} = 0$ between $\alpha(f, |f|)$ and T, which equals the delay for flight f. Using these definitions, we define the flight delay constraints for one-stop itineraries where both flights are in $\mathcal{F}$.

$$d_{f_1 f_2} \geq d(f_2) - T\big(cx(f_1) + x_{f_1 f_2}\big) \quad \forall (f_1, f_2) \in \mathcal{OS}: f_1, f_2 \in \mathcal{F}. \tag{23}$$

The second term in the right-hand side of constraints (23) allows the flight delay decision variable, $d_{f_1 f_2}$, to equal 0 if either flight f_1 is canceled or there is a missed connection. Thus, as with constraints (19) and (20), because the costs associated with flight delays are strictly positive, $d_{f_1 f_2}$ will equal $d(f_2)$ if and only if flight f_1 is scheduled and there is no missed connection. Note that in the case flight f_2 is canceled, $d(f_2)$ will equal 0 by definition.

2.3.3   Objective Function

The objective function that drives scheduling and cancelation decisions is composed of three parts: i) small fixed benefits for flight cancelations, ii) superlinear operating costs associated with flight delays, and iii) estimated passenger delay costs. For example, in deciding which flights to cancel, there would be no operating costs, so the fixed cancelation benefit must be balanced against the expected passenger delays associated with re-accommodating all passengers booked on the flight. We measure each of the objective function components relative to expected passenger delay minutes, scaling the cancelation benefit by the number of seats on the aircraft and the operating costs by the number of seats times a factor $0 < \zeta < 1$.

The benefit for flight cancelation is defined by a single parameter, ν. Thus, we specify the cancelation component of the objective function, $g^{\mathrm{canc}}(y)$, as follows.

$$g^{\mathrm{canc}}(y) = \nu \sum_{f \in \mathcal{F}} [ns(f) \cdot cx(f)]. \tag{24}$$

In practice, the operational costs and constraints associated with flight delays are defined by many factors, including crew availability and work restrictions, aircraft maintenance schedules, etc. Because we do not have access to this information, we instead use a superlinear delay cost as an approximation. That is, we assume that as flight delays increase, it becomes increasingly difficult to operate the flight. The form of

the delay cost function we use is similar to the exponential penalty cost function utilized by the RBS-EP model. For the airline disruption response, we utilize a fixed delay threshold, $\psi \in \mathbb{N}^+$. Above the threshold, the delay costs are piecewise linear, with the slope of each subsequent piece increasing by a factor of λ. Unlike the RBS-EP model, we specify a parameter, $\rho \in \mathbb{N}^+$, which defines the length of each of the pieces. Using this notation, the un-scaled delay costs associated with flight f arriving at time t, c_{ft}, are defined as follows.

$$c_{ft} = \min\{t - \alpha(f, |f|), \psi\} + \sum_{\epsilon=1}^{t-\alpha(f,|f|)-\psi} \lambda^{\text{ceil}(\epsilon/\rho)}. \tag{25}$$

In equation (25), the function ceil($\cdot$) determines the integer ceiling of its argument. It is instructive to compare the form of equation (25) to equation (8). Using these definitions, the operating cost component of the objective function, $g^{\text{ops}}(y)$, is specified as:

$$g^{\text{ops}}(y) = \zeta \sum_{f \in \mathcal{F}} \left[ns(f) \cdot \sum_{t=\alpha(f,|f|)}^{T-1} c_{ft}\left(y_{f|f|t} - y_{f|f|(t-1)}\right) \right]. \tag{26}$$

If flight f is canceled, the corresponding flight delay cost term in equation (26) will equal 0, because $\left(y_{f|f|t} - y_{f|f|(t-1)}\right) = 0$ for all t.

The last and most complex component of the objective function represents the estimated passenger delay costs. To calculate these costs, we split passenger itineraries into four categories: i) non-stop itineraries, ii) one-stop itineraries where $f_1 \notin \mathcal{F}$, iii) one-stop itineraries where $f_2 \notin \mathcal{F}$, and iv) all other one-stop itineraries. The estimated passenger delay cost, $g^{\text{pdel}}(y)$, is the sum of these four categories: $g_1^{\text{pdel}}(y)$, $g_2^{\text{pdel}}(y)$, $g_3^{\text{pdel}}(y)$, and $g_4^{\text{pdel}}(y)$ respectively.

For non-stop passengers, the delays are due to either flight delays or flight cancelation, thus:

$$g_1^{\text{pdel}}(y) = \sum_{f \in \mathcal{F}} np(f)[d(f) + rd(f)cx(f)]. \tag{27}$$

For one-stop itineraries where $f_1 \notin \mathcal{F}$ we assume that flight f_1 arrives on-time, thus all of the delay costs are associated with flight f_2, and:

$$g_2^{\text{pdel}}(y) = \sum_{\{(f_1,f_2) \in \mathcal{OS}: f_1 \notin \mathcal{F}\}} np(f_1, f_2)[d(f_2) + rd^2(f_1, f_2)cx(f_2)]. \tag{28}$$

For one-stop itineraries where $f_2 \notin \mathcal{F}$, we assume that flight f_2 departs on-time, thus the connection is missed flight f_1 arrives too late to satisfy the minimum connection time. The last arrival interval for flight f_1 to avoid a missed connection is $\bar{t}(f_1, f_2) = \alpha(f_1, |f_1|) + w(f_1, f_2) - w_{\min}$, therefore:

$$
g_3^{\text{pdel}}(y) = \sum_{\{(f_1, f_2) \in \mathcal{OS} : f_2 \notin \mathcal{F}\}} np(f_1, f_2) \Big[rd^1(f_1, f_2) cx(f_1) \\
+ rd^{\text{x}}(f_1, f_2) \Big(1 - y_{f_1 | f_1 | \bar{t}(f_1, f_2)} - cx(f_1) \Big) \Big].
\tag{29}
$$

For the second term in equation (29), we accumulate the re-accommodation delay $rd^{\text{x}}(f_1, f_2)$ if flight f_1 has not arrived by time $\bar{t}(f_1, f_2)$ and it is not canceled.

For one-stop itineraries where both flights are under the airline's control, disruptions can occur due to cancelations of either flight or due to a missed connection. Note that if the itinerary is not disrupted, passengers experience flight delays based only on the second flight in the itinerary, thus:

$$
g_4^{\text{pdel}}(y) = \sum_{\{(f_1, f_2) \in \mathcal{OS} : f_1 f_2 \in \mathcal{F}\}} np(f_1, f_2) \big[rd^1(f_1, f_2) cx(f_1) + rd^2(f_1, f_2) z_{f_1 f_2} \\
+ rd^{\text{x}}(f_1, f_2) x_{f_1 f_2} + d_{f_1 f_2} \big].
\tag{30}
$$

In equation (30), we use $z_{f_1 f_2}$ instead of $cx(f_2)$ to eliminate double-counting in the case that flights f_1 and f_2 are both canceled. Using $d_{f_1 f_2}$ in place of $d(f_2)$ eliminates double-counting associated with delays on the second flight in the itinerary.

The overall objective function equals the sum of the three cost components: $g^{\text{canc}}(y)$, $g^{\text{ops}}(y)$, and $g^{\text{pdel}}(y)$. Outside of the estimated re-accommodation delays, which we specific as inputs, the objective function is parameterized by ν, ζ, ψ, λ, and ρ.

2.4   Quasi-Compression

When evaluating different allocation techniques, it is important to consider something like compression to ensure that the allocations that lead to a larger number of cancelations do not look worse due simply to the absence of schedule compression. The challenge is that because of the nature of the 1-to-1 exchanges performed by the FAA, compression is defined only in the single-resource TFM setting. Thus, to approximate the impact of compression in the multi-resource setting, we have developed an optimization-based quasi-compression procedure. The resulting schedule provides an approximate lower-bound on the resulting flight and passenger delays.

The goal of our quasi-compression model is to compress the multi-resource TFM schedule, ensuring that each flight receives no more delay than in the combined schedule and that flow balance is maintained for all airlines and aircraft types. To attempt to distribute delay improvements fairly between flights, we extend the RBS-EP allocation model, determining the task ordering, $\mathcal{J}(r)$, based on the scheduled arrival times in the combined schedule. We denote the set of canceled flights as $\mathcal{CX}$, and these flights are excluded from the task orderings. For each flight $f \notin \mathcal{CX}$, we denote its arrival time in the combined schedule as $t^*(f)$. To ensure that cancelations are maintained and that flight delays are not extended, we add the following constraints to the quasi-compression model:

$$y_{f|f|(T-1)} \quad = \quad 0 \quad \forall f \in \mathcal{CX}, \text{and} \tag{31}$$

$$y_{f|f|t^*(f)} \quad = \quad 1 \quad \forall f \in \mathcal{F} \setminus \mathcal{CX}. \tag{32}$$

Additionally, for each of the 15 airline groups, we add constraints (17) and (18) to ensure that flow balance is maintained in the compressed model.

2.5   Passenger Delay Calculation

Given either the combined or compressed flight schedule, we can estimate passenger delays using the passenger delay calculator from (*12*). One challenge is that because the allocated flight schedule corresponds to only a portion of a single day of operations, most disrupted passengers will be unable to be re-accommodated to their final destinations. For these passengers, instead of defaulting to the maximum re-accommodation delay as described in the referenced paper, we use the disruption cost regression model described in (*13*) to estimate the delays associated with each disrupted itinerary. If the estimate from the disruption delay model exceeds the maximum re-accommodation delay, only then do we default the passenger to the maximum re-accommodation delay.

3   Scenario Construction

The scenarios we develop in this chapter are similar to the ones developed in (*10*). The most significant divergence is that we limit the flight schedules to the 20 carriers represented in the Airline Service Quality Performance (ASQP) 2007 data set. In Section 3.1, we describe the process for generating the required data for each of these flights, and in Section 3.2 we describe how we compute the TFM program capacities. In addition to flight schedules and TFM program capacities, our sequential evaluation procedure requires planned aircraft routings and booked passenger itineraries. Thus, in Sections 3.3 and 3.4, we describe how we generate each of these data sets respectively.

3.1   Flight Schedules

To build the scenario flight schedule, we start with flight schedule data from Flight Schedule Monitor (FSM) (*14*) that corresponds to a single day of relatively clear weather operations –April 23rd, 2007. We restrict the schedule to the FSM flights that are operated by an ASQP carrier on origin-destination segments that exist in the T-100 Domestic Segment (T-100) database. This reduces the number of flights we consider from over 38,000 to just over 16,000.

3.2   TFM Program Capacities

To develop our capacity reduction scenarios, we start with 10 days of historical disruption scenarios. As in (*10*), we create two capacity reduction levels for the conflicting GDP and AFP programs, one representing historical capacities and one representing hypothetical capacities as congestion en-route increases. The specified resource capacities for each of the 20 capacity reduction scenarios (10 historical disruption days x 2 capacity reduction levels) correspond to the 38,000 flights. When we restrict our attention to the smaller set of 16,000 flights, we need to adjust the capacities accordingly. To do so, we perform a full schedule allocation and then calculate the scenario capacities based on the utilization of the 16,000 flights. For any time period during which a TFM program is not in effect, we consider resource capacities to be infinite.

3.3   Aircraft Routings

To describe a valid aircraft routing, a flight sequence s must satisfy the following requirements.

1. The flight sequence must be connected, i.e., $orig([s]_{k+1}) = dest([s]_k)$ for all k,
2. All flights in the sequence must be of the same aircraft type, i.e., $p([s]_1) = \cdots = p([s]_{(0)})$, and
3. The flight sequence must be consistent with the minimum aircraft turnaround time m, i.e., $\alpha([s]_{k+1}, 1) \geq \alpha([s]_k, |[s]_k|) + m$ for all k.

To determine the planned aircraft routings, we utilize the aircraft tail numbers. If we had accurate tail numbers for each flight, we could estimate the planned aircraft routings by tracing tail numbers over the course of the day. Unfortunately, the FSM flights do not include tail numbers, although we can determine tail numbers for most flights by joining against the ASQP database. The tail number data in ASQP is sometimes inaccurate, so a sorted sequence of flights sharing the same tail number is occasionally invalid according to the requirements described above. When we encounter inconsistencies, we split the invalid sequence up into its longest valid sub-sequences. Last, to minimize the number of aircraft routings, we greedily extend valid flight sequences (treating flights without tail numbers as unit-length sequences) by appending the next flight sequence that maintains the validity of the original sequence. Using this

procedure, we end up with 4002 aircraft routings for the 16,100 flights over the 12 hour FSM flight window (including all flights that arrive or depart in the window).

3.4   Passenger Itineraries

In (*12*), the authors developed a methodology for estimating historical passenger itineraries from aggregated demand data. One challenge is that the itineraries developed are based on the flights in ASQP, which do not match the FSM flights exactly. An additional complication is that the FSM flight schedules are based on actual flight times on the day of operations, whereas the ASQP-based itineraries are based on planned flight times. To address these problems, we perform the following three steps based on the procedures outlined in the referenced paper.

1.   Generate potential itineraries based on the FSM flights;
2.   Estimate *carrier-route* demands for the day by aggregating the passengers allocated to the ASQP-based itineraries over the same time period; and
3.   Allocate *carrier-route* passengers to the matching FSM itineraries using the discrete choice allocation model and estimated parameters.

This re-allocation process ensures that the passenger itineraries in our airline disruption response model are feasible in the planned schedule (e.g., one-stop itineraries must satisfy the minimum connection time). In total, we allocate 1,151,718 passengers to the 16,100 flights. Of the passengers, 36% are allocated to one-stop itineraries.

To estimate the expected re-accommodation delay associated with disruption of each of these itineraries, we use the disruption cost regression model developed in (*13*). To ensure that the expected delay associated with re-accommodation is strictly positive and reasonably bounded, we limit the range of the model so that estimated disruption costs are always between 2 hours and 24 hours (e.g., if the model predicts less than 2 hours, we treat the disruption cost as 2 hours). For the airline disruption response model, we estimate that a passenger connection is missed if there is less than 15 minutes of connection time (i.e., $w_{\min} = 3$).

4   Results

As described in Section 2.3.3, the objective function for the airline disruption response model is described by the following five parameters: ν, ζ, ψ, λ, and ρ:

$\nu < 0$ =   cost benefit per seat associated with flight cancelations;

$0 < \zeta < 1$ =   seating capacity scaling factor applied for operating costs;

$$\psi \quad = \quad \text{linear operating cost threshold;}$$

$$\lambda \quad = \quad \text{exponential growth rate for piecewise linear operating costs; and}$$

$$\rho \quad = \quad \text{length of each segment in the superlinear portion of the operating cost curve.}$$

In our testing, we use $\nu = -6$, corresponding to negative 30 minutes of passenger delays per seat (based on 5 minute time intervals). Assuming average passenger delays of 6 hours for flight cancelations (a reasonable, low-end estimate), this parameter value suggests that all else being equal an airline would like to cancel a flight that is less than $1/12^{th}$ full. In the airline disruption model, all else is not equal, because single flight cancelations are impossible due to the flow balance constraints. Nonetheless, this seems like a subjectively reasonable threshold.

For the superlinear operating cost function, we consider two sets of parameter values, ($\zeta = 0.5$, $\psi = 6$, $\lambda = 1.5$, and $\rho = 3$) and ($\zeta = 0.5$, $\psi = 12$, $\lambda = 1.25$, and $\rho = 3$). The first suggests that flight operating costs grow linearly for up to 30 minutes of delay, before growing by a rate that increases exponentially by a factor of 1.5 every 15 minutes. The second suggests that flight operating costs grow linearly for up to an hour of delay, before growing by a rate that increases exponentially by a factor of 1.25 every 15 minutes. In each, the flight delay costs are scaled by 0.5 suggesting that baseline delay for an empty seat is half as costly as for a passenger. When referencing the two operating cost curves in our results, we refer to the first curve as *aggressive* and the second as *conservative*.

To understand the sensitivity of the model to these parameters, we compare the total costs (i.e., cancelation benefit plus operating costs plus passenger delays), based on a flight with an average 81.2% load factor (passengers divided by seats) and an average 36.4% of connecting passengers (split evenly between first leg and second leg passengers). Each of these values represents the empirical averages based on the planned passenger itineraries in our test scenarios. Under the aggressive parameter set, the airline is indifferent towards canceling this hypothetical flight based on just less than 2 hours of flight delay. For the conservative parameter set, the total cost exceeds the threshold of cancelation indifference at just less than 3 hours of flight delay.

Between the 10 days of historical disruptions scenarios, the 2 scenario types (historical and hypothetical), and the two operating cost parameterizations, we consider 40 sequential evaluation scenarios in total. We report the results for each scenario date averaged across the two parameter sets and the two scenario types; the conclusions drawn from the averaged results are the same as those drawn from each of the four individual segments.

4.1   Benefits of Optimization

In (**10**), the authors suggested that there is a significant benefit to coordinating conflicting TFM programs (i.e., GDPs and AFPs) using an optimization-based approach. The analysis looked at the equity and efficiency of competing allocation schemes including two multi-resource RBS scheduling approaches used in practice, referred to as precedence RBS and exemption RBS. One limitation with that analysis is that the authors compared only allocated flight delays. In practice, due to flight cancelations and swaps, realized flight delays can differ significantly from those allocated. It is possible to conceive situations where one approach allocates fewer delays, but does so in a way that systematically leads to a worse result. In this section, using our sequential evaluation procedure, we demonstrate that the benefits suggested in the paper, though slightly diminished, appear to be fairly robust even in the dynamic CDM environment. We perform this analysis by comparing results from the exemption RBS allocation to the Ration-by-Schedule Exponential Penalty (RBS-EP) model at each stage in the evaluation procedure: allocated, combined, and compressed. We use a constant $\lambda = 1.25$ as the base for the exponential penalty in the RBS-EP model, allowing up to 60 minutes of delay beyond the maximum RBS delay threshold. In the result tables, we refer to this as the RBS-EP model.

In Table 2, we provide statistics regarding the allocated flight schedules. The percentage of unfair delay is based on the time-order deviation metric developed in (**10**). Of the 20 individual scenarios tested, the RBS-EP model is at least as fair according to time-order deviation in 19, with the 5/9/2007 hypothetical scenario being the lone exception. In aggregate, the RBS-EP is significantly fairer, allocating 2.4% of unfair delay compared to 8.4% for exemption RBS, while reducing total flight delays by 11.6%.

| Scenario Date | Exemption RBS | | RBS-EP | |
	Avg Flt Delay	% Unfair Delay	% Delay Reduction	% Unfair Delay
5/2/2007	10.4	14.8%	14.6%	3.4%
5/9/2007	7.4	1.4%	7.9%	1.7%
6/19/2007	25.6	7.7%	13.8%	2.8%
6/27/2007	28.4	4.2%	9.3%	1.8%
6/28/2007	33.1	10.7%	16.2%	2.8%
7/5/2007	12.3	8.3%	6.0%	0.8%
7/16/2007	5.7	8.8%	10.8%	5.0%
7/18/2007	23.5	15.8%	8.6%	2.8%
7/27/2007	8.5	3.0%	11.6%	2.6%
9/27/2007	11.4	6.9%	7.5%	1.3%
Summary	**19.0**	**8.4%**	**11.6%**	**2.4%**

Table 2: Comparison of allocated flight schedules

To compare the combined airline schedules, we consider flight cancelations, flight delays, and estimated passenger delays. From this stage on, we ignore fairness, because airlines are free to reschedule flight according to their own internal objectives. To better understand how the initial allocation might impact these internal objectives, we list the sum of the objective values across each of the 15 airline disruption response models. In Table 3, we perform these comparisons averaged across both parameter sets and scenario types.

| Scenario Date | Exemption RBS | | | | RBS-EP (1.25, 60) [% Reductions] | | | |
	# Cx Flights	Avg Flt Delay	Tot Pax Delay	Objective Value	# Cx Flights	Avg Flt Delay	Tot Pax Delay	Objective Value
5/2/2007	21.0	8.7	697,061	226,059	9.6%	10.8%	8.6%	16.6%
5/9/2007	4.0	7.9	511,804	158,626	0.0%	7.2%	8.3%	7.6%
6/19/2007	101.8	21.2	6,106,069	2,212,208	27.7%	7.0%	12.6%	11.9%
6/27/2007	101.5	23.0	5,574,536	1,879,670	5.1%	9.4%	9.7%	10.4%
6/28/2007	111.3	24.4	5,263,559	1,868,648	17.8%	8.5%	9.3%	10.0%
7/5/2007	11.5	12.0	1,267,157	408,008	10.8%	5.4%	11.9%	12.1%
7/16/2007	6.5	5.7	404,223	119,429	6.3%	9.0%	6.8%	9.1%
7/18/2007	89.0	17.6	3,943,034	1,255,284	6.1%	3.0%	10.6%	10.2%
7/27/2007	12.5	8.4	1,492,187	482,756	9.5%	11.0%	14.2%	14.6%
9/27/2007	28.0	9.6	1,741,369	594,579	6.1%	5.1%	6.4%	8.4%
Summary	**487.0**	**15.5**	**27,000,999**	**9,205,267**	**13.5%**	**7.4%**	**10.5%**	**10.9%**

Table 3: Comparison of combined schedules

In the table, we see that the RBS-EP model exhibits across the board improvements as compared to the exemption RBS scheduling approach. The improvement in total objective values equals 10.9%. If we treat this as a proxy for total costs, the percentage improvement is slightly smaller than the delay reduction indicated in Table 2 (11.6%), though still substantial. Though fewer flights are canceled using the conservative cost curve (2.1% of TFM-impacted flights), the cancelation rate using the aggressive cost curve (3.3% of TFM-impacted flights) is also well within historical norms. For instance, the average cancelation rate across all ASQP flights in 2007 was 2.4%, which includes many flights not impacted by TFM programs.

One thing to note is that the RBS-EP model balances relatively smaller improvements in flight delays against larger improvements in flight cancelations. The difference in cancelation rates between the exemption RBS and RBS-EP allocations suggests the importance of considering schedule compression, because each cancelation creates an opportunity for further delay reduction. This intuition is borne out by the results in Table 4, which show that after quasi-compression, average flight delays have decreased by a relatively larger amount based on the exemption RBS allocation (i.e., the percentage reduction when comparing average flight delays for exemption RBS and RBS-EP has decreased).

Scenario Date	Exemption RBS		RBS-EP (1.25, 60) [% Reductions]	
	Avg Flt Delay	Tot Pax Delay	Avg Flt Delay	Tot Pax Delay
5/2/2007	8.1	671,695	5.7%	5.4%
5/9/2007	7.8	517,493	7.0%	8.9%
6/19/2007	19.2	5,764,674	0.8%	8.5%
6/27/2007	20.7	5,413,536	5.9%	6.9%
6/28/2007	21.0	5,007,257	-0.2%	4.2%
7/5/2007	11.8	1,274,602	5.5%	11.7%
7/16/2007	5.7	404,089	9.1%	6.7%
7/18/2007	16.1	3,853,374	3.6%	9.6%
7/27/2007	8.2	1,490,476	10.1%	14.0%
9/27/2007	9.4	1,762,680	4.2%	5.9%
Summary	**14.2**	**26,159,875**	**3.6%**	**7.6%**

Table 4: Comparison of compressed schedules

For the combined schedules, the improvement in passenger delays tracks closely the improvement in objective value (see Table 3), which is not surprising because passenger delays make up a significant component of the airline disruption response objective function. Thus, when evaluating the compressed schedule, we consider passenger delay improvement as a proxy for overall cost reduction. For the

compressed schedules, the average passenger delay improvement is 7.6%. This values compare to the 11.6% flight delay improvement calculated based on the initial allocation. This suggests that a rather significant portion – perhaps as much as 30% – of the improvement reflected by a single-stage analysis is illusory. Note that it is unlikely that the result is inflated by a full 30%, because the compression procedure utilized in practice is more restrictive than the quasi-compression we consider here. Nonetheless, the magnitude of the difference reinforces the importance of evaluating TFM allocation schemes in a CDM-like environment.

Acknowledgments

References

1. Brennan, M. Airspace flow programs - a fast path to deployment. *Journal of Air Traffic Control* **2007,** *49* (1), 51-55.

2. Chang, K.; Howard, K.; Oiesen, R.; Shishler, L.; Tanino, M.; Wambsganns, M. Enhancements to the FAA ground-delay program under collaborative decision-making. *Transportation Science* **2001,** *34,* 57-76.

3. Ball, M.; Hoffman, R.; Odoni, A.; Rifkin, R. A stochastic integer program with dual network structure and its application to the ground-holding problem. *Operations Research* **2003,** *51,* 167-171.

4. Kotnyek, B.; Richetta, O. Equitable models for the stochastic ground-holding problem under collaborative decision-making. *Transportation Science* **2006,** *40,* 133-146.

5. Richetta, O.; Odoni, A. Solving optimally the static ground-holding policy problem in air traffic control. *Transportation Science* **1993,** *27,* 228-238.

6. Vossen, T.; Ball, M. Optimization and mediated bartering models for ground delay programs. *Naval Research Logistics* **2005,** *53* (1), 75-90.

7. Bertsimas, D.; Patterson, S. The air traffic flow management problem with enroute capacities. *Operations Research* **1998,** *46* (3), 406-422.

8. Andreatta, G.; Brunetta, L.; Guastalla, G. From ground holding to free flight: an exact approach. *Transportation Science* **2000,** *34,* 394-401.

9. Bertsimas, D.; Lulli, G.; Odoni, A. The air traffic flow management problem: an integer optimization approach. *Integer Programming and Combinatorial Optimization*, 2008; pp 34-46.

10. Barnhart, C.; Bertsimas, D.; Caramanis, C.; Fearing, D. Equitable and efficient coordination in traffic flow management. *MIT working paper* **2010**.

11. Bureau of Transportation Statistics. TranStats. http://transtats.bts.gov/TableInfo.asp?Table_ID=236&DB_Short_Name=On-Time&Info_Only=0 (accessed September 3, 2009).

12. Barnhart, C.; Fearing, D.; Vaze, V. Modeling passenger travel and delays in the National Air Transportation System. *MIT working paper* **2010**.

13. Fearing, D. *The case for coordination: equity, efficiency, and passenger impacts in air traffic flow management;* MIT Ph.D. Thesis, 2010.

14. Metron Aviation. Flight Schedule Monitor. http://www.metronaviation.com/fsm.php (accessed May 2009).