DYNAMIC STOCHASTIC OPTIMIZATION MODEL FOR AIR TRAFFIC FLOW MANAGEMENT WITH EN ROUTE AND AIRPORT CAPACITY CONSTRAINTS

Avijit Mukherjee

Institute for Systems Research, University of Maryland, College Park, MD 20740

avijit@isr.umd.edu

Mark Hansen

Department of Civil and Environmental Engineering

University of California, Berkeley, CA 94720

mhansen@ce.berkeley.edu

Introduction

In this paper, we present a linear dynamic stochastic optimization model for Air Traffic Flow Management (ATFM) that accounts for uncertainty in both airport and en route airspace capacity. Rather than analyzing this problem in its full generality, we focus on the case in which there is a single destination airport and a small number of arrival fixes subject to blockage or reduced capacity as a result of weather. This would typically occur when there is weather in the vicinity of the airport that affects capacity of some of the standard arrival fixes of the airport, along with the acceptance rate (or arrival capacity) of the airport itself. The main decisions in the model are pre-departure delay (or ground holding) and local rerouting of inbound flights.

We perform experiments by applying the model to manage arrival flow at Dallas Fort-Worth Intl. Airport (DFW). We consider hypothetical cases when weather blocks some of the standard arrival fixes of DFW, and reduces the airport capacity. Such situations commonly arise at DFW when thunderstorms occur in the vicinity of the airport, and effectively block some of the standard routes and/or fixes delivering inbound traffic. In our model, ground delays are dynamically revised based on updated information on capacity forecasts. When weather impacts en route fixes, and the cost ratio between unit airborne and ground delay is low, some flights are rerouted through weather-free regions. When the airport is the only bottleneck, rerouting flights is not necessary. Results from our experiments show that when weather impact on en route fixes is severe and persistent, substantial benefits in terms of delay savings is achieved from dynamically rerouting flights via weather-free arrival fixes compared to cases where ground delays are dynamically readjusted but no airborne delays are allowed.

Background

The research literature includes several optimization models for ATFM. In almost all such models, the objective is to minimize total delay cost, which has two components – ground and airborne delays. Most of the optimization models in the literature are formulated as linear and/or integer programming models. Both deterministic and stochastic optimization methods have been applied to solve problems in ATFM. Stochastic approaches have been confined mostly to the ground holding problem (GHP), in which it is assumed that one or more airports are the only bottlenecks in the system. Previous research that considers airspace as well as airport congestion has focused on deterministic problems.

Bertsimas-Stock [2] addressed the ATFM problem under deterministic time-varying en route airspace and airport capacities. They proposed an Integer Programming (IP) formulation. They showed the problem is equivalent to the "Job Shop Scheduling" problem [3], which is *NP-hard*. Nevertheless, many constraints in their formulation are facets of the polyhedron defined by the set of all constraints, and therefore in many instances, even ones with large numbers of variables and constraints, integer solutions are obtained directly from the LP relaxation of their model.

The Bertsimas-Stock model is "disaggregate" in formulation because the decision variables are related to each individual flight's release (or departure) time, amount of airborne holding, and route. This makes the problem very large for realistic situations involving thousands of flights and hundreds of congestible NAS resources. The computation time can become exorbitant if LP relaxation doesn't yield integer optimal solutions. But the disaggregate formulation also has advantages. It allows flight-specific sets of route(s) and cost functions to be used as inputs. Disaggregation also allows variations of the original formulation that include route choice, allow dynamic rerouting, consider delay propagation, and capture dependence between arrival and departure capacities of airports. A limitation of the model is that it is deterministic, and thus ignores uncertainty in resource capacities.

Uncertainty in resource capacities in ATFM has been addressed mainly in context of the GHP ([1], [4], [7], [8].) In all of these models, uncertainty has been addressed by considering a set of scenarios, each corresponding to a time-varying airport capacity profile. Richetta and Odoni [8] also used the concept of a scenario tree that represents evolving information about which scenario will be realized. In Mukherjee and Hansen [4], we formulated a multi-stage stochastic integer program with recourse that allows dynamic revision of ground delays imposed on non-departed flights, based on updated information on the capacity conditions at the destination airport. Updated information is obtained through branching of the scenario tree. We specified a set of coupling constraints that impose the condition that unless a particular scenario is realized completely, the ground holding decisions cannot be made exclusively based on that scenario. In other words, the decision to release a flight must consider all scenarios that have not been eliminated from possibility at the time of release. The decision variables are related to departure times of individual flights. The disaggregate formulation allowed us to easily incorporate non-linear measures of delay, without having to reformulate the problem as a non-linear optimization model. Such measures of delay are useful to address equity issues in GHP [9].

Nilim et al. [6] addressed weather related uncertainty in airspace condition in routing individual flights. They formulated a robust Markov Decision Process problem for dynamically rerouting an aircraft across convective weather impacted region. Decisions in their problem are aircraft speed and heading. This research is among the first to address uncertainty in en route airspace conditions. However, their methodology is more relevant to controlling individual aircraft while en route – i.e. Air Traffic Control, whereas ATFM requires tools at the planning stage of mitigating demand-capacity imbalance, which is what we propose in this paper.

In summary, while there is considerable literature on ATFM optimization, a significant gap remains. While some models consider multiple constraints in both airport and en route capacities, and others address uncertainty in airport capacity scenarios, the literature does not include models that consider both airport and en route capacity constraints while addressing uncertainty. That void is addressed in this paper.

Dynamic Stochastic Optimization Model for ATFM

In this section we present a linear stochastic optimization model for ATFM in which ground delays of non-departed flights can be revised after updated information on evolving capacity at the airport and arrival fixes becomes available. While en route, flights can be rerouted at any en route fix along their path where alternative routes are available

Scenario Tree

Uncertainty in capacities is addressed by considering multiple scenarios. Each scenario represents the joint evolution of capacity profiles of multiple NAS resources – arrivals fixes for an airport and the airport itself. A scenario tree consists of branches with branching points. Each branch corresponds to a particular scenario or group of scenarios. A branching point, or node, in the scenario tree represents a time when certain scenarios are eliminated as possibilities for a particular day. Thus, as the day progresses, more specific information about the capacity scenario that will be realized on that day becomes available. Branching points may coincide with times when capacity profiles of different scenarios diverge, or they may precede those times as a result of forecast information. (See [5] for detailed representation of a scenario tree.)

Let Θ denote the set of capacity profile scenarios, and the unconditional probability of occurrence of a scenario $\xi \in \Theta$ be given by $P\{\xi\}$. Let B be the total number of branches of the scenario tree; $B \geq |\Theta|$. Each branch corresponds to a set of scenarios. Let the η_b scenarios corresponding to branch $b \in \{1,..,B\}$ be given by the set

$\Omega_b = \{\psi_1^b,...,\psi_{\eta_b}^b\} \subset \Theta$. The time periods corresponding to start and end nodes of a branch are denoted by o_b and e_b.

Network Representation of Airspace

The airspace is represented via network $G = (N, A)$ formed by a set of nodes N and a set of directed links A interconnecting the nodes. A node can be an airport or an en route fix; and a link is a pre-specified flight trajectory between two nodes. Let $P \subset N$ denote the set of arrival fixes of the destination airport.
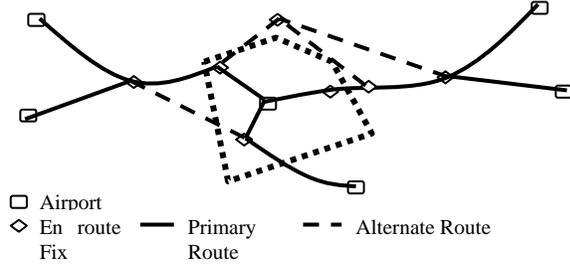

Figure 1. Network Representation of Airspace

Figure 1 illustrates the network representation of the airspace. Between each origin airport and the destination, there is a set of links that forms the primary (or shortest distance) route. Also, as shown in the figure, rerouting can occur at some pre-specified set of en route fixes using alternate links. Moreover, there can be multiple fixes along the primary path from an origin, where rerouting decisions can be made. The dotted bold line in the figure represents the boundary of a hypothetical airspace in the vicinity of an airport. The standard arrival fixes for the airport lie within this region. The capacity of some of the arrival fixes as well as the airport may be affected by adverse weather. Several scenarios, as discussed earlier, depict time-varying capacity profiles of NAS resources within the impacted region.

Model Formulation

Let Φ denote the set of flights, each of which originates from an airport denoted by org_f, $f \in \Phi$. The destination airport is denoted by a node k. The planning horizon is divided into time periods of equal intervals, and is denoted by a set $\Gamma = \{1,...,T\}$. Let dep_f and arr_f denote the scheduled departure and arrival time periods of a flight $f \in \Phi$. Let $\Lambda_f \subset A$

denote the set of links that forms the available routes for flight f. The minimum travel time of a flight on various links on its flight path is given by the parameters $\tau_{i,j}^f; (i, j) \in \Lambda_f; i, j \in N$. To reduce the problem size, we define a time window (a set of time periods) for each flight, within which it must have crossed a NAS resource, if the resource lies on the route chosen by the flight. The time window is denoted by $L_i^f; i, j \in N : (i, j) \in \Lambda_f$. Defining such a time window also imposes an upper limit to the amount of airborne holding a flight can be subject to at various locations on its flight path. Let $C_j^\xi(t), j \in P \bigcup \{k\}$, denote the scenario-specific time-varying capacity profiles of various resources under each scenario $\xi \in \Theta$. A resource can be an arrival fix $p \in P$ or the destination airport k. Let $\lambda > 1$ denote the cost ratio between one unit of airborne and ground delay for all flights.

The decision variables, which are binary, are defined as follows:

$$Y_{f,i,j}^\xi(t) = \begin{cases} 1 \text{ if flight } f \text{ has crossed node } i \text{ on its flight} \\ \quad \text{path along link } (i, j) \quad \text{by the end of time} \\ \quad \text{period } t \text{ under scenario } \xi \\ \\ 0 \text{ otherwise} \end{cases}$$

$$f \in \Phi, (i, j) \in \Lambda_f, t \in L_i^f, \xi \in \Theta$$

At the origin airport $(i = org_f)$, the decision variables represent the time period by when a flight has departed under different scenarios. Any difference from the scheduled departure time implies a ground delay. Ground delays can differ across scenarios. This means that the ground delay of a flight that has not yet departed may be revised based on updated information on evolving conditions at the destination airport and arrival fixes. Dynamic rerouting occurs at fixes along the flight path that are divergence points for alternative routes to the destination airport.

The scenario-specific planned arrival time of individual flights at the destination airport (denoted by node k) can be determined from the above decision variables, and is represented by the auxiliary variable:

$$X_f^\xi(t) = \begin{cases} 1 \text{ if flight } f \text{ is planned to arrive at the airport k} \\ \quad \text{by end of time period } t \text{ under scenario } \xi \\ \\ 0 \text{ otherwise} \end{cases}$$

The variables $X_f^\xi(t)$ are related to the decision variables $Y_{f,i,j}^\xi(t)$ by the following expression:

$$X_f^\xi(t) = \sum_{p \in P:(p,k) \in \Lambda_f} Y_{f,p,k}^\xi(t - \tau_{p,k}^f)$$

Note that the planned arrivals may be different from actual arrivals (or landings) at the airport due to capacity constraints. When arrival capacity is inadequate, certain flights will face airborne holding inside the airport TRACON. Let $W^\xi(t)$ denote the number of flights that are subject to airborne queuing delay in the airport TRACON area at time period $t \in \Gamma$, under scenario $\xi \in \Theta$. The variables $W^\xi(t)$ will be greater than zero if the airport acceptance rate at time period t is below the total arrival demand during that time period. The amount of airborne holding within the TRACON can be limited by imposing additional constraints on the variables $W^\xi(t)$.

The scenario-specific ground delay of a flight can be derived from the decision variables as follows:

$$GD_f^\xi = \sum_{t \in L_{org_f}^f} \sum_{j:(org_f,j) \in \Lambda_f} (Y_{f,org_f,j}^\xi(t) - Y_{f,org_f,j}^\xi(t-1)).(t - dep_f)$$

The scenario-specific en route delay of a flight before its planned arrival at the airport is given by the following expression:

$$ED_f^\xi = \sum_{t=arr_f+1}^{T} \left(X_f^\xi(t) - X_f^\xi(t-1)\right)(t - arr_f) - GD_f^\xi$$

The objective function given below minimizes the expected total cost of delay, where airborne delay is weighted by the cost ratio λ.

$$\sum_{\xi \in \Theta} P\{\xi\}.(\sum_{f \in \Phi} GD_f^\xi + \lambda.(\sum_{f \in \Phi} ED_f^\xi + \sum_{t=1}^{T} W^\xi(t)))$$

The set of constraints $\forall f \in \Phi,\ (i,j) \in \Lambda_f$, and $t \in L_i^f, \xi \in \Theta$ are as follows:

$$Y_{f,i,j}^\xi(t) - Y_{f,i,j}^\xi(t-1) \geq 0 \qquad (1)$$

$$Y_{f,i,j}^\xi(t) - \sum_{(i',i) \in \Lambda_f} Y_{f,i',i}^\xi(t - \tau_{i',i}^f) \leq 0 \qquad (2)$$

$$\sum_{j:(i,j) \in \Lambda_f} Y_{f,i,j}^\xi(t) \leq 1 \qquad (3)$$

$$W^\xi(t-1) - W^\xi(t) + \sum_{f \in \Phi} \left(X_f^\xi(t) - X_f^\xi(t-1)\right) \leq C_k^\xi(t)$$
$$;\quad \forall t \in \Gamma, \xi \in \Theta \qquad (4)$$

$$\sum_{f \in \Phi:(p,k) \in \Lambda_f} \left(Y_{f,p,k}^\xi(t) - Y_{f,p,k}^\xi(t-1)\right) \leq C_p^\xi(t);$$
$$t \in \Gamma, p \in P, \xi \in \Theta \qquad (5)$$

$$W^\xi(0) = 0;\quad \forall \xi \in \Theta \qquad (6)$$

$$W^\xi(T) = 0;\quad \forall \xi \in \Theta \qquad (7)$$

$$X_f^\xi(T) = 1;\quad \forall f \in \Phi, \xi \in \Theta \qquad (8)$$

$$Y_{f,i,j}^{\psi_1^b}(t) = ... = Y_{f,i,j}^{\psi_q^b}(t) = ... = Y_{f,i,j}^{\psi_{\eta_b}^b}(t); \qquad (9)$$
$$\psi_q^b \in \Omega_b : \eta_b \geq 2 \text{ and } o_b \leq t \leq e_b$$

$$Y_{f,i,j}^\xi(t) \in \{0,1\} \qquad (10)$$

$$W^\xi(t) \geq 0;\quad \forall \xi \in \Theta, t \in \Gamma \qquad (11)$$

Constraints (1) reflect the non-decreasing property of decision variables $Y_{f,i,j}^\xi(t)$. Thus if a flight f has crossed a node i and entered the link $(i,j) \in \Lambda_f$ by end of time period t_0, under scenario ξ, then the variables $Y_{f,i,j}^\xi(t)$ remain 1 for all time periods $t \geq t_0$. Constraints (2) reflect the requirement that a flight can enter a link $(i,j) \in \Lambda_f$ by time period t under any scenario ξ, only if it has flown a minimum distance corresponding to an incoming link to node i -- $\tau_{i',i}^f; (i',i) \in \Lambda_f$. Constraints (3) state the condition that only one among several available links can be chosen at a node for rerouting. Constraints (2) and (3) together ensure that each flight selects only one route (among several available) under any specific scenario.

Constraints (4) impose an upper bound, which is the airport scenario-specific airport capacity, to the number of flights that can land at the destination airport during any time period t, under each scenario ξ. Variables $W^\xi(t)$ reflect the number of flights that are planned to arrive during or before the time period t, that are subject to airborne queuing, due to deficit in the airport capacity. Constraints (5) impose the scenario-specific capacity of various arrival fixes to be the upper bound to the corresponding throughput at each of those fixes.

Constraints (6) state that the system is empty, i.e. there are no flights in airborne queuing delay at

the airport, in the beginning of the planning horizon. Constraints (7) – (8) ensure all flights land by the end of planning horizon.

Constraints (9) impose the condition that the decision variables $Y^\xi_{f,i,j}(t)$ are the same for all scenarios ξ that are represented by the evolving branch of the scenario tree at time t - i.e. $\xi \in \Omega_b : o_b \leq t \leq e_b$. Note that if the number of scenarios represented by the evolving branch is one, then the coupling constraints are no longer imposed. In other words, scenario-specific decisions can be based on a specific capacity scenario only after it has been realized. As long as the active branch of the scenario tree represents more than one scenario, the decisions must be same under all those scenarios.

Experimental Results

Problem Setup

We apply the above model to manage flow of aircraft inbound to Dallas Fort Worth Intl. Airport (DFW) under uncertainty in adverse weather occurrence time and location. Flights arriving at DFW usually fly over one of the four standard arrival fixes – BYP, CQY, JEN, and UKW. Under fair weather conditions, each of the arrival fixes serve air traffic originating from a set of airports, whose primary route to DFW passes over it. We consider a hypothetical convective weather event that affects the two eastern arrival fixes – BYP and CQY – along with the acceptance rates at DFW itself. A shaded region in Figure 2 shows the location of the convective weather.
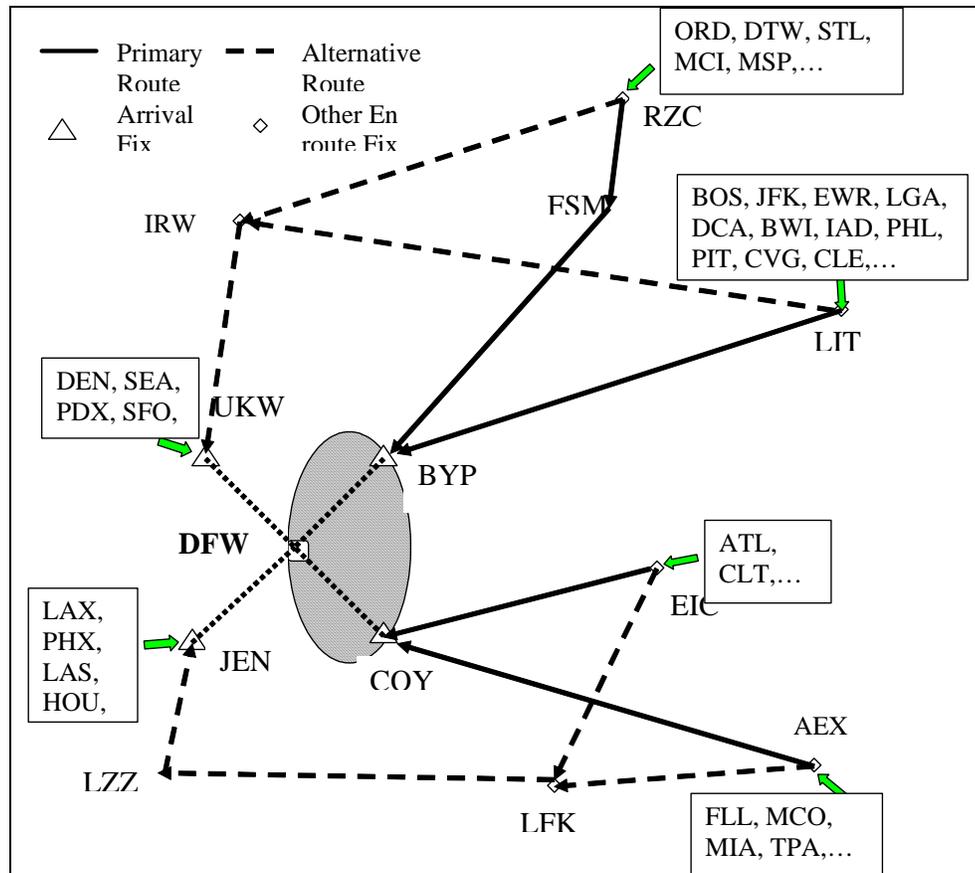

Figure 2. Standard Arrival Routes and Rerouting Options at DFW

We consider 351 flights that are scheduled to arrive at DFW by 12:15 PM on July 14, 2003. The planning horizon of 12 hours and 15 minutes is divided into 49 quarter-hour time periods, starting from midnight – 12:15 AM, and ending at noon –

12:15 PM. Thus $\Gamma = \{1,..,49\}$ is the set of time periods. The minimum flight time on any link is based on the Great Circle Distance between two points in airspace, and an assumed average speed of

an aircraft. The flight time is expressed in 15-minute units and rounded to the nearest integer.

In this experiment, we assume that the presence of convective weather blocks the affected arrival fixes so that the capacity goes to zero. The maximum throughput of the arrival fixes in fair weather conditions is assumed to be 15 aircraft/qhr (based on 5 miles-in-trail horizontal separation between aircraft and an assumed speed of 300 nautical miles per hour). For DFW itself, we assume that the 15-minute arrival capacity reduces to 15 arrivals/qhr under convective weather, and in fair weather conditions it goes up to 35 arrivals/qhr.

Capacity scenarios and the scenario tree are derived as follows. Convective weather activity at DFW is assumed to be present from the beginning of the day (beginning of planning horizon) until the end of a time period t_{dfw}, after which the weather clears off. Similar weather activity occurs at the arrival fixes – BYP and CQY – where the clearance time period is t_{fix}. Both t_{dfw} and t_{fix} can take any one of five possible clearance times – 8:30AM, 9:00AM, 9:30AM, 10:00AM, and 10:30AM. Note that t_{dfw} and t_{fix} can be different. The scenario tree representing capacity changes at DFW is shown in Figure 3. The set of scenarios $s_i^{dfw}, i = 1,..,5$, represent the time-varying airport arrival capacity, under different weather clearance times. A similar scenario tree can be constructed for the capacity changes at the two arrival fixes; and the set of scenarios is given by $s_i^{fix}, i = 1,..,5$.
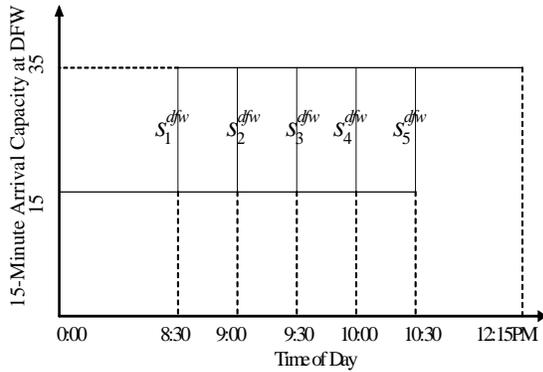

Figure 3. Scenarios of Airport Acceptance Rates at DFW

Combining the possible clearance times for the airport and the arrival fixes, there are a maximum of 25 total number of scenarios denoted by the set

$\Theta = \{\xi_{11}, \xi_{12},..\xi_{ij},..\xi_{55}\}$, each of which represents the joint evolution of capacity changes at DFW and the two arrival fixes (BYP and CQY) corresponding to the weather clearance times. Scenario ξ_{ij} denotes the capacity scenario in which t_{dfw} takes the i^{th} clearance time among the five possible (ordered in increasing time of day), and t_{fix} takes the j^{th}; i.e. $\xi_{ij} = (s_i^{dfw}, s_j^{fix}), i, j \in \{1,..,5\}$. Table 1 relates each scenario with the weather clearance times. For example, scenario ξ_{14} represents the case when weather at DFW clears at 8:30 AM and the arrival fixes clear at 10:00 AM; similarly, scenario ξ_{22} corresponds to the clearance time of 9:00AM at both airport and the arrival fixes.

Table 1. Scenarios Corresponding to Different Combination of Weather Clearance Times

Weather Clearance Time at DFW	Weather Clearance Time at the Arrival Fixes (BYP and CQY)				
	8:30 AM	9:00 AM	9:30 AM	10:00 AM	10:30 AM
8:30 AM	ξ_{11}	ξ_{12}	ξ_{13}	ξ_{14}	ξ_{15}
9:00 AM	ξ_{21}	ξ_{22}	ξ_{23}	ξ_{24}	ξ_{25}
9:30 AM	ξ_{31}	ξ_{32}	ξ_{33}	ξ_{34}	ξ_{35}
10:00 AM	ξ_{41}	ξ_{42}	ξ_{43}	ξ_{44}	ξ_{45}
10:30 AM	ξ_{51}	ξ_{52}	ξ_{53}	ξ_{54}	ξ_{55}

To analyze the sensitivity of the results, we consider a set of alternative cases in which various model inputs are varied, and compare the results with the baseline case. The cases are defined as follows.

Baseline Definition (Case 1): We assume that all capacity scenarios are equally likely, i.e. $P\{\xi_{ij}\} = 0.04$, where $\xi_{ij} \in \Theta$. The cost ratio λ is set to 3, which is fairly low. Therefore in this case, we expect to see airborne holding and extra airborne time due to longer flight paths.

Case 2: High Cost Ratio. In this case we set the cost ratio λ to 25. The unconditional probabilities of capacity scenarios remain unchanged. Intuitively, we expect to see higher ground delays than in the baseline case, and little or no expected airborne delays.

Case 3: High cost ratio along with high probability of early weather clearance. The cost ratio λ is set to 25, and the probability mass function of capacity scenarios is give as follows: $P\{\xi_{11}\} = 0.4, P\{\xi_{ij}\} = 0.025$ for $\xi_{ij} \in \Theta \setminus \xi_{11}$. We

expect to see higher benefits from dynamic revision of ground delays in this case.

Case 4: Value of early information. In this case our goal is to analyze the delay cost reduction if the information on capacity changes are obtained 30 minutes earlier than the actual changes in weather conditions. Therefore in effect, branches of the scenario tree unveil two time periods earlier than that in baseline scenario tree. All other parameters and input data are kept the same as in the baseline case.

Case 5: Convective weather affecting DFW only. In this case we assume that the convective weather impacts arrival capacity of the airport (DFW) only, and not the arrival fixes. The weather clearance time at DFW -- t_{dfw} -- can take five possible values as in the baseline case. Assuming all scenarios to be equally likely, the probability of occurrence of each scenario becomes 0.2.

Results

In all cases, integer solutions were obtained from LP relaxation of the problem. The computing times were in order of 10 – 12 minutes, using CPLEX version 7.1 run on a personal computer with 2.2 GHz processor speed.

Figure 4 summarizes the results from applying the dynamic stochastic model to each of the five cases described above. The low cost ratio (λ =3) in cases 1, 4, and 5, results in airborne delays. This is because flights are allowed to depart with relatively lower ground delays in anticipation of weather clearance. If adverse weather persists later, flights are either rerouted via secondary route or subject to airborne holding. Since the airport is the only bottleneck in case 5, rerouting is not necessary to avoid en route capacity constraints, hence there is no extra flight time in this case.

In all five cases, ground delays of flights are revised as updated information becomes available through scenario tree branching. The model adopts a "wait-and-see" policy in assigning ground delays to many flights. If the departure time of a flight is close to a branching point in the scenario tree, it may be delayed until the branching occurs. If the branching reveals improvement in capacity conditions, the flight may be released with little or no further delay. Otherwise, its ground delay may be extended. Such policy reduces the risk of airborne delays.
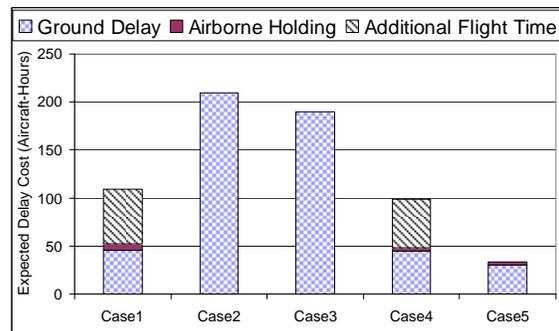

Figure 4. Expected Delay Costs (Aircraft-hours)[1]

In cases 2 and 3, due to the high cost ratio, ground delays are severe and there are no airborne delays. In both cases, flights are released only when there is no chance of their airborne delays. In other words at any instant ground delays are imposed on flights based on the capacity profile of worst scenario among all that belong to the active branch of scenario tree at that instant. Ground delays are revised to lower amounts only if favorable conditions are realized. Therefore the high cost ratio effectively imposes an additional constraint of zero airborne delay, which makes the total delay costs in case 2 and 3 relatively higher – 91% and 71% respectively -- than case 1. The expected ground delay is lower in case 3 due to high probability of early weather clearance.

The delays in case 4 are lower than case 1 because the model can exploit the early information on capacity changes. The reduction of 10% in total delay cost in this case compared to the baseline reveals the value of 30 minutes advanced forecast of capacity. This points to value of dynamic models in assessing the value improved weather forecasting capability.

Tables 2 and 3 provide further information on the performance of the dynamic stochastic model in the baseline case. The values presented in Table 2 are the ratios, expressed in percentages, of scenario-specific ground delays obtained from the dynamic stochastic model to those under perfect information, i.e. when the realized scenario is known from the beginning of the planning horizon. Thus, the ratios reveal the loss due to imperfect information. Table 3 provides similar ratios for airborne delays.

---

[1] Costs measured in units of ground delay, with airborne delay and extra flight time multiplied by the delay cost ratio λ.

Table 2. Scenario-Specific Ground Delay Ratio, Expressed in Percentage, between Stochastic Model and that under Perfect Information in the Baseline Case

Weather Clearance Time at DFW	Weather Clearance Time at the Arrival Fixes (BYP and CQY)				
	8:30 AM	9:00 AM	9:30 AM	10:00 AM	10:30 AM
8:30 AM	700	178	135	167	526
9:00 AM	265	197	150	211	346
9:30 AM	128	150	128	125	185
10:00 AM	95	107	118	94	101
10:30 AM	72	80	85	89	91

Reduced capacity at the airport necessitates ground delays, while that at the en route fixes can be managed by rerouting flights. Therefore, for a given weather clearance time at the fixes, higher ground delays are necessary for later weather clearance at DFW. Although the dynamic stochastic model assigns higher ground delays in response to later clearance times at DFW, it is unable to match the requirements as produced under perfect information (see Table 2). For late clearance at the fixes, rerouting is necessary even under perfect information. By virtue of rerouting, the dynamic stochastic model is able to reduce the gap between airborne delays it produces and that under perfect information (see Table 3).

Table 3. Scenario-Specific Airborne Delay Ratio, Expressed in Percentage, between Stochastic Model and that under Perfect Information in the Baseline Case

Weather Clearance Time at DFW	Weather Clearance Time at the Arrival Fixes (BYP and CQY)				
	8:30 AM	9:00 AM	9:30 AM	10:00 AM	10:30 AM
8:30 AM	336	288	160	115	98
9:00 AM	308	265	151	108	97
9:30 AM	400	246	168	125	104
10:00 AM	501	243	157	155	121
10:30 AM	607	260	185	134	129

When the destination airport (DFW) is the main bottleneck, such as under scenario ξ_{51} in which weather at DFW clears at 10:30AM while that at en route fixes clears at 8:30AM, the perfect information solution absorbs virtually all delay through ground holding flights, since rerouting confers little benefit. The ground delay assigned by the stochastic model under this scenario is 28% less than that necessary under perfect information, while the airborne delay

due to imperfect information is significantly higher – 507%. The difference is reduced as en route congestion becomes increasingly severe, resulting in a perfect information solution with more rerouting. For example under the scenario ξ_{55}, airborne delay is only about 29% higher compared to that under perfect information.

Another revealing scenario in the baseline case is ξ_{15}, in which conditions at the airport improve early while adverse weather at the arrival fixes persists. Perfect information in this scenario leads to rerouting many flights via alternative routes, in order to avoid excessive ground delays. Therefore the total ground delay under perfect information about scenario ξ_{15} has a lower value compared to the other scenarios, while there are higher airborne delays due to extra flight time on longer routes. The "wait-and-see" policy for assigning ground delays in the dynamic stochastic model backfires if scenario ξ_{15} occurs. In this case, the excessive ground delays assigned by the model in anticipation of scenarios ξ_{12} and ξ_{13} cannot be recovered if ξ_{15} occurs. Thus total ground delay is about 426% more compared to what would happen if we knew that at the beginning of the day that scenario ξ_{15} would be realized.

Conclusions

In this paper, we present a stochastic integer programming model for managing inbound air traffic flow of an airport. We assume that convective weather, which is present in the vicinity of the airport, affects airport acceptance rates (or arrival capacity) and maximum throughput of some of the standard arrival fixes of the airport. The model allows decisions on rerouting and ground delay of flights to be revised based on updated information.

A set of capacity scenarios, each representing the joint evolution of time-varying capacities of multiple NAS resources (arrival fixes, and airport), is considered. A scenario tree, in which branching occurs when new information eliminates certain scenarios, provides the information on evolution of the system with time. The probability of occurrence of each scenario is assumed to be known at the beginning of planning horizon, and is provided as input to the optimization model.

When the capacity reduction at en route fixes is severe and persistent, and cost ratio between unit airborne and ground delays is low, rerouting flights

via weather free regions yields substantial benefits by alleviating high ground delays. The model achieves substantial reduction in ground delay releasing flights with less ground delays and dynamically rerouting them if faced with capacity restrictions at en route fixes. Rerouting is not necessary if airport is only bottleneck in the system. Revising ground delays in response to updated information on capacity scenarios is the only control knob in such cases along with those in which airborne holding is prohibited. The model adopts a "wait-and-see" policy in assigning ground delays to flights.

While there is benefit from planning cleverly with poor information, there remains great value from making the information better. Improved information can be utilized by the model to revise ATFM decisions on ground delay and rerouting of flights. Value of advanced information is revealed in the experimental results presented earlier.

By virtue of disaggregate formulation (i.e. individual flight level) the model can handle a much wider range of objective functions taking into account equity issues and non-linearity in the delay cost function. In all experimental cases integer optimal solutions were obtained directly from the LP relaxation. There is a need for further investigation of the general conditions under which this will occur.

Scalability of the model is a further issue. Computational complexity may become a major burden in extending it to deal with a region that encompasses multiple airports and en route fixes. Sector capacities constraints, along with airport and en route fixes, must be added. In such cases, formulating aggregate models may have larger advantages in computing time over the disaggregate model. Research on developing aggregate models with certain limitations on dynamic revision of decisions but substantial gain in computation times is currently under progress. Analysis of computational complexity of the stochastic optimization models would reveal whether further research is necessary to develop heuristics that can yield reasonably approximate solutions in practical computing time. Additional research is also required to develop scenarios and build scenario trees for the joint evolution of capacity profiles of multiple NAS resources required by these models. In building these trees, we face challenges analogous to the curse of dimensionality faced by stochastic dynamic programmers.

Finally, it will be necessary to adapt the stochastic optimization models into the Collaborative Decision Making framework is they are to be implemented in the air traffic management system currently used in the United States.

Acknowledgements

References

[1] Ball, M., R. Hoffman, A. Odoni, and R. Rifkin, 2003, A *Stochastic Integer Program with Dual Network Structure and Its Application to the Ground-Holding Problem*, Operations Research, Vol. 51.

[2] Bertsimas, D., and S. Patterson, 1998, *The Air Traffic Flow Management with En route Capacities*, Operations Research, Vol. 46.

[3] Garey, M. and D. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, 2000, New York, W.H. Freeman and Company.

[4] Mukherjee, A. and M. Hansen, *Dynamic Stochastic Model for a Single Airport Ground Holding Problem*, 2003, University of California, Berkeley, Paper Submitted for Publication.

[5] Mukherjee, A., *Stochastic Optimization Models for Air Traffic Flow Management*, 2004, Ph.D. Dissertation, University of California, Berkeley.

[6] Nilim, A., L. Ghaoui, and V. Duong, *Robust Dynamic Routing of Aircraft under Uncertainty*, 2002, IEE 21st Digital Avionics Systems Conference.

[7] Richetta, O., and A. Odoni, *Solving Optimally the Static Ground Holding Policy Problem in Air Traffic Control*, 1993, Transportation Science, Vol. 27.

[8] Richetta, O., and A. Odoni, *Dynamic Solution to the Ground Holding Problem in Air Traffic Control*, 1994, Transportation Research Vol. *28A*.

[9] Vossen, T., *Fair Allocation Methods in Air Traffic Management*, 2002, Ph.D. Dissertation, University of Maryland, College Park.