Resource Allocation in Flow-Constrained Areas with Stochastic Termination Times

Moein Ganji[14], Alex Nguyen[13], David Lovell[14], Michael Ball[12]
University of Maryland, College Park, MD 20742
[1] Institute for Systems Research
[2] Robert H. Smith School of Business
[3] Department of Electrical and Computer Engineering
[4] Department of Civil and Environmental Engineering

Abstract
In this paper we formulate an optimization problem for the assignment of dispositions to flights whose preferred flight plans pass through a flow-constrained area. For each flight, the disposition can be either to depart as scheduled but via a secondary route that avoids the flow-constrained area, or to use the originally intended route but to depart with a controlled departure time and accompanying ground delay. We anticipate that the capacity through the flow-constrained area will increase at some future time once the weather activity clears. The model is a two-stage stochastic program that represents the time of this capacity windfall as a random variable, and determines expected costs given a second-stage decision, conditioning on that time. The goal is to minimize the expected cost over the entire distribution of possible capacity increase times.

I. INTRODUCTION

A flow-constrained area (FCA) is a region of the national airspace system (NAS) where a capacity-demand imbalance is expected due to some unexpected condition such as adverse weather, security concerns, special-use airspace, or others. Flow-constrained areas might be drawn as polygons in a two-dimensional space, although in practice they are usually represented by a single straight line, functioning as a cordon.

When an FCA has been defined, it is then often the case that an airspace flow program (AFP) is invoked by the Federal Aviation Administration (FAA). An AFP is a traffic management initiative (TMI) issued by the FAA to resolve the anticipated capacity-demand imbalance associated with the FCA. It is the goal of this paper to develop a method by which, given the aggregate data described here, specific orders for individual flights can be developed for a single FCA that a) maximize the utilization of the constrained airspace, b) prevent the capacity of the FCA from being exceeded, and c) achieve a system-wide delay minimization objective. We recognize that this model cannot be directly applied to AFP planning as it does not address issues related to the manner in which the FAA and the flight operators collaborate in reaching a final decision regarding each flight. Our goal here is to develop relevant stochastic optimization models. We intend to address issues related to collaborative decision making (CDM) in later papers.

II RELATED RESEARCH

The research most closely related to this paper has to do with airport ground holding. Much work has been done in addressing the airport ground holding problem, including the development of stochastic integer programming models, [1], [4], [5], [6], [11]. However, there is still much active research in the development of models for managing flights through congested areas of en route airspace under weather uncertainty.

In [10], the rerouting of a single aircraft to avoid multiple storms and minimize the expected delay was examined. In this model, the weather uncertainty was treated as a two-state Markov chain, with the weather being stationary in location and either existing or not existing at each phase in time. A dynamic programming approach was used to solve the routing of the aircraft through a gridded airspace, and the aircraft was allowed to hedge by taking a path towards a storm with the possibility that the storm may resolve by the time the aircraft arrived. The focus of the work was on finding the optimal geometrical flight path of the aircraft, and not on allocation of time slots through the weather area. Follow-on work expanded to modeling multiple aircraft with multiple states of weather and attempted to consider capacity and separation constraints at the storms.[9][8]

Initial steps at a concept of operations that describes the terminology, process, and technologies required to increase the effectiveness of uncertain weather information and the use of a probabilistic decision tree to model the state space of the weather scenarios was provided in [1]. Making use of this framework is a model recently proposed that uses a decision-tree approach with two-stage stochastic linear programming with recourse to apportion flows of aircraft over multiple routing options in the presence of uncertain weather [3]. In

ISBN: 978-0-615-20720-9

the model, an initial decision is made to assign flights to various paths to hedge against imperfect knowledge of weather conditions, and the decision is later revised using deterministic weather information at staging nodes on these network paths that are close enough to the weather that the upcoming weather activity is assumed known with perfect knowledge. Since this is a linear programming model, only continuous proportions of traffic flow can be obtained at an aggregate level, and not decisions on which individual flights should be sent and when they should arrive at the weather. In [7], a stochastic integer programming model is developed based on the use of scenario trees to addressed combined ground delay-rerouting strategies in response to en route weather events. While this model is conceptually more general than ours, by developing a more structured approach we hope to develop a more scalable model.

Recently, a Ration-by-Distance (RBD) method was proposed as an alternative to the Ration-by-Schedule (RBS) method currently used for Ground Delay Programs (GDPs) that maximizes expected throughput into an airport and minimizes total delay if the GDP cancels earlier than anticipated [3]. This approach considers probabilities of scenarios of GDP cancellation times and assigns a greater proportion of delays to shorter-haul flights such that when the GDP clears and all flights are allowed to depart unrestricted, the aircraft are in such a position that the expected total delay can be minimized. While this problem was applied to GDPs, the principles of a probabilistic clearing time where there is a sudden increase in capacity and making initial decisions such that the aircraft are positioned to take the most advantage of the clearing is similar to our problem.

III. MODEL

A. *Model inputs*

Our base model inputs consist of information about the FCA, which is consistent with the information used in AFP planning.

- Location of the FCA
- Nominal (good weather) capacity of the FCA
- Reduced FCA (bad weather) capacity of the FCA
- Start time of the AFP
- Planned end time of the AFP

From a list of scheduled flights and their flight plans, we determine the set of flights whose paths cross the FCA, and who therefore would be subject to departure time and/or route controls under an AFP. We also require a set of alternate routes for each flight (see Figure 1). The alternate route for each flight should be dependent on the geometry of the FCA and the origin-destination pair it serves. These most likely would be submitted by carriers in response to an AFP; for the purposes of this paper it is assumed they are submitted exogenously, although for testing purposes it was necessary to synthesize some alternate routes.

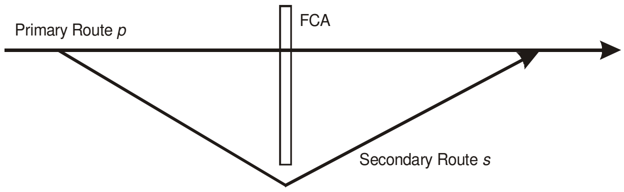

Figure 1. Primary route through the FCA and secondary route bypassing FCA.

B. *Controls*

In order not to exceed the (reduced) FCA capacity, each flight will be assigned one of two dispositions in the initial plan reacting to the FCA:

1. The flight is assigned to its primary route, with a controlled departure time that is no earlier than its scheduled departure time. Given an estimate of en route time, this is tantamount to an appointment (i.e., a slot) at the FCA boundary. Some flights might be important enough that they depart on time, the AFP notwithstanding. Other flights might be assigned some ground delay.
2. The flight is assigned to its secondary route, and is assumed to depart at its scheduled departure time.

This limited set of conceived actions imposes several important assumptions:

- We do not consider airborne holding as a metering mechanism to synchronize a flight on its primary route with its slot time at the FCA
- We assume that any necessary number of flights can be assigned to their secondary routes without exceeding any capacity constraints in other parts of the airspace. (In fact, our model can easily be extended to handle such “other” capacity constraints but in this initial version we do not include them.)
- We assume that, when the weather clears, the FCA capacity increases immediately (“in one step”), back to some higher capacity.
- The random variable is the time at which the FCA capacity increases back to a higher value. We assume that perfect knowledge of the realization of this random variable is not gained until the scenario actually occurs, and so no recourse can be taken until the scenario is realized.

C. *Scenarios and future responses*

The outputs of this model are:

ISBN: 978-0-615-20720-9

1. An initial plan that designates whether a flight is assigned to its primary route or secondary route; for those assigned to their primary route an amount of ground delay (possibly zero) is assigned.
2. A recourse action for each flight under each possible early clearance time.

We model the time at which the weather clears (i.e. FAC capacity increases) as a discrete random variable, with some exogenous distribution. For any realization of the capacity increase time, the flights in question will be in some particular configuration as specified in the initial plan. Some will have departed, either on their primary or secondary routes, some will already have completed their journeys, and some will still be at their departure airports.

Flights that were originally assigned to their primary route and that have already taken off will be assumed to continue with that plan. For any such flight, the primary route is assumed to be best, so no recourse action is necessary.

We now consider flights originally assigned to their primary route that have not yet taken off. We need not consider transferring them to their secondary routes, because if that were a good idea in the improved capacity situation, it would also have been a good idea in the initial plan. Thus, the only possible change in disposition for these flights involves potentially moving their controlled departure time, i.e. reducing their assigned ground delay. Constraints are required to define the range of times the fight can arrive at the FCA boundary based on the required en route time and the time the recourse action is taken (clearance time). We also explicitly enforce the restriction that under such situations the assigned ground delay (or equivalently arrival time at FCA) cannot be increased.

All other flights not yet considered were originally assigned to their secondary routes, with departure times as originally scheduled. These secondary routes avoid the FCA somehow. Under the FCA capacity windfall, some of those flights may now have an opportunity to use the FCA. If a flight has not yet taken off, and it is decided that it can use the FCA, the lowest cost way to do this is to re-assign it back to their primary route, with some controlled departure time no earlier than their scheduled departure time. If, on the other hand, the flight has already taken off, then the only mechanism to allow it the use of the FCA is a hybrid route that includes that portion (and perhaps more) of the secondary route already flown, plus a deviation that traverses the FCA and presumably rejoins the primary route at some point after the FCA (see Figure 2). A flight that is already en route via its secondary route may or may not prefer such a hybrid path, depending on the difference in cost (time, fuel, etc.) between doing that and continuing on its secondary route. There may be many possible hybrid routes, and perhaps only a limited set of those would be acceptable to carriers and air traffic control (ATC).

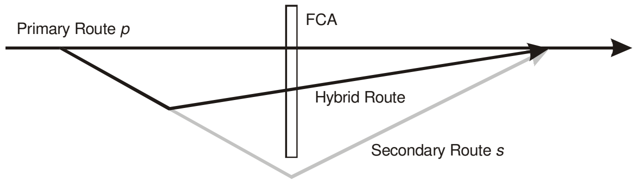

Figure 2. Reverting from secondary route back to primary route through FCA.

For each possible value of the capacity windfall time, we determine the expected locations of all affected flights at that time, and also what would be the best change in disposition, if any, for each of those flights according to a system performance metric. With this information, we can compute the conditional cost associated with flying these flights under that realization of the stochastic event. Ultimately, then, the goal of the optimization problem is to minimize the expected total cost, given these conditional costs and their probabilities.

D. Model development

We start by defining the discrete lattice on which time will be represented. We assume there is an index set $\{1,\ldots,T\}$ of size T that demarcates equally spaced time slots, each of duration Δt. Each of these represents a possible appointment time window at the FCA. The nominal capacity of the FCA should be specified in terms of the maximum number of flights permissible during one of these time windows. The number of time slots T then depends directly on Δt and the total duration of an AFP, perhaps inflated to allow for ending times later than the original estimate. The reference time $t=1$ can be chosen as the earliest scheduled departure time of all of the affected flights. The actual time indicated by the index t is then $(t-1/2)\Delta t$.

The flights affected by the FCA can be determined from the filed flight plans for that day, minus known cancellations and re-routes at the time the AFP is invoked. These flights are indexed according to the set $\{1,\ldots,F\}$. In the rest of the paper, any specific reference to a time period t and flight f assumes that $t\in\{1,2,\ldots,T\}$ and $f\in\{1,\ldots,F\}$.

1) First stage (Initial Plan)

There are two sets of assignment variables that are related to decisions about the dispositions of flights. One set represents the initial plan, which are the decisions provided by the model that will be enacted immediately once the model is run and the AFP is declared. The second set represents conditional decisions (recourse actions) based on the random variable representing the time at which the capacity windfall takes place, which we do not know at the time of the execution of this optimization problem, but that we condition for when determining the best initial plan.

 ISBN: 978-0-615-20720-9

For the initial plan, we define the following set of binary decision variables:

$$x^p_{f,t} = \begin{cases} 1, \text{ if flight } f \text{ uses its primary route and} \\ \text{has an appointment time } t \text{ at the FCA} \\ 0, \text{ otherwise} \end{cases}$$

$$x^s_f = \begin{cases} 1, \text{ if flight } f \text{ is assigned to its secondary} \\ \text{route} \\ 0, \text{ otherwise} \end{cases}$$

Every flight *f* needs to have an assigned disposition under the initial plan, thus:

$$\sum_t x^p_{f,t} + x^s_f = 1 \quad \forall f \qquad (1)$$

We require that any flight that is assigned to its primary route cannot be given an appointment slot at the FCA that is earlier than its scheduled departure time plus the expected en route time required to arrive at the FCA. If $E_f \Delta t$ represents the en route time (from its origin to the FCA) for flight *f*, and $D_f \Delta t$ is the scheduled departure time for flight *f*, then:

$$\sum_{t=1}^{D_f+E_f} x^p_{f,t} = 0 \quad \forall f \qquad (2)$$

This construction requires that en route times and scheduled departure times are represented on the same discrete lattice as the FCA appointment times.

No similar constraint is applied to flights assigned to their secondary routes under the initial plan, because they are not metered at any point and hence are expected to depart at their originally scheduled departure time. There is no provision in the model for a flight to depart early, despite the fact that the secondary route takes more time than the primary route (since, subject to minor variations, airlines do not allow flights to take off before their scheduled departure times).

It might be the case that for a particular flight *f*, there is a latest slot time l_f at the FCA that the carrier who owns that flight would be willing to accept. Slots later than l_f can be prevented via the following constraint:

$$\sum_{t=l_f+1}^{T} x^p_{f,t} = 0 \qquad (3)$$

For any flight for which l_f is not explicitly provided, l_f is an effective time beyond which it would never make sense not to choose the secondary route.

The initial constrained capacity (maximum number of flights) for time window t can now be defined as C^0_t and the constraint to enforce it is:

$$\sum_f x^p_{f,t} \le C^0_t \quad \forall t \qquad (4)$$

2) Second Stage (Revised Plan)

The variables and constraints defined so far represent the first stage of the stochastic program. It is assumed that these decisions will be enacted deterministically immediately after the FCA is declared. Next, we describe the second stage of the stochastic program – those variables that represent the conditional decisions we expect *would be* made if any of a number of possible capacity windfall times happens to come true in the future. We model the time slot at which this occurs as a discrete random variable with domain Ω and probability mass function

$$f_U(u) = \Pr\{U = u\} \quad \forall u \in \Omega$$

Under a capacity windfall, a flight that was originally assigned to its primary route with a controlled departure time might still be given the same general disposition, although its departure time could be moved earlier if that were beneficial to the system goal. We let

$$y^p_{f,t} \mid u = \begin{cases} 1, \text{ if at the time } U = u \text{ of the capacity windfall,} \\ \text{flight } f \text{ is assigned to its primary route with} \\ \text{appointment slot } t \text{ at the FCA} \\ 0, \text{ otherwise} \end{cases}$$

We will (shortly) introduce other variables for the other possible second stage flight dispositions, and we will require that all flights be assigned a disposition under every possible realization of the stochastic event *U*. For now, we proceed by obviating values of $y^p_{f,t} \mid u$ that would either be physically infeasible or politically imprudent. Later, structural constraints plus pressure from the objective function will lead to the best possible selection of second stage dispositions for all flights.

First, it is impossible to assign a flight to a slot that would require it to depart before its scheduled departure time:

$$y^p_{f,t} \mid u = x^p_{f,t} \quad \forall f, u, \quad \forall t \in \{1, ..., D_f + E_f\} \qquad (5)$$

This constraint works with constraint (2) to achieve the required result.

Given the timing *U* of the capacity windfall, some flights may already have taken off. If they did so via their primary route (with a controlled departure time), then their second stage disposition should match that of the first stage:

$$y^p_{f,t} \mid u = x^p_{f,t} \quad \forall f, u, \quad \forall t \in \{1, ..., u + E_f\} \qquad (6)$$

A closer look at constraint (6) reveals that it also satisfies an important requirement for flights that have not yet taken off. For any particular flight *f* and given the capacity windfall time *u*, the collection of primary stage variables $\{x^p_{f,t}\}_{t=1}^{t=u+E_f}$ will either contain one at exactly one position or it will consist entirely of zeros. In the former case, this means that the flight has already taken off, and that situation has been dealt with. In the latter case, this is indicative of the fact that these slot times are infeasible. Thus, even for flights that have not yet

 ISBN: 978-0-615-20720-9

taken off, constraints (2) and (6) insure that they will not be assigned, in the second stage, to their primary routes with slot times that they cannot achieve.

Looking at constraints (5) and (6), it is clear that they can be combined:

$$y^{p}_{f,t}|u = x^{p}_{f,t} \quad \forall f, u, \quad \forall t \in \{1,\ldots,\max(u, D_f) + E_f\} \tag{7}$$

On the other hand, for flights that already took off via their secondary routes (and therefore at their scheduled departure times), the only possible second stage dispositions are secondary or hybrid routes, so assignments to primary routes for these flights must be prevented:

$$\sum_t y^{p}_{f,t} \mid u \le 1 - x^{s}_{f} \quad \forall u, \forall f \ni D_f < u \tag{8}$$

In addition, we will not allow a flight whose controlled departure time is being moved in the face of a capacity windfall to be worse off than it was before this event materialized:

$$y^{p}_{f,t} \mid u \le \sum_{q \ge t} x^{p}_{f,q} + x^{s}_{f} \quad \forall u, f, t \tag{9}$$

Notice that we want to allow for the possibility that flights originally assigned to their secondary routes can revert, under the appropriate circumstances and if the optimization decides this is best, to their primary route if they have not already taken off, which is why the variable x^{s}_{f} appears in constraint (9).

For flights that were originally assigned to the secondary route, the increased capacity at the FCA might allow some of these flights to pass through the FCA and thus improve their flight path by returning to the primary route at some point after the FCA or continuing directly to the destination. For a flight that has not yet departed, one could choose to have the same structure apply, but the portions of the total flight path spent on the secondary and reverting routes would then have to have length zero. In this paper, as will be shown later, we use a different approach. We define the second-stage decision variables for this choice as follows:

$$y^{h}_{f,t} \mid u = \begin{cases} 1, & \text{if flight } f \text{ was originally assigned to its secondary route, but under capacity clearing time } u \text{ has been assigned an FCA appointment slot } t \\ 0, & \text{otherwise} \end{cases}$$

This decision can only be reached for flights that were originally assigned to their secondary routes:

$$y^{h}_{f,t} \mid u \le x^{s}_{f} \quad \forall u, f, t \tag{10}$$

However, it should be obvious that the objective function will enforce this behavior implicitly.

The flights in question will be on their secondary routes and diverting onto a hybrid route that passes through the FCA. We need to impose constraints that insure that these flights are only assigned to FCA time slots they can feasibly reach. If a flight diverts from its secondary route to its hybrid route at time t^d there will be an earliest time it can reach the FCA. Figure 3 illustrates the geometry used to compute the parameter used by our model:

$t^{d}_{f,t}$ is the time at which f must divert from its secondary route so as to use its hybrid route that arrives at the FCA at time t.

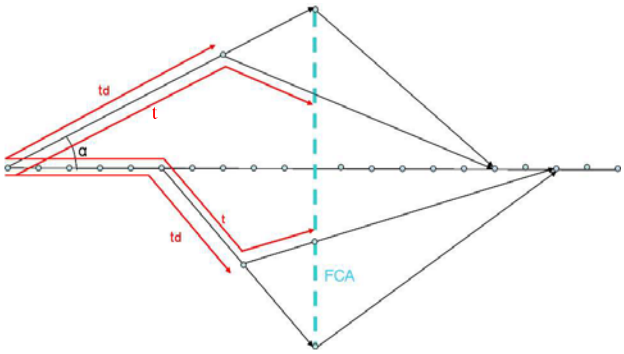

Figure 3. Possible diversion from secondary to hybrid route.

The following constraint prevents a flight from diverting to its hybrid route before the weather is actually cleared.

$$y^{h}_{f,t} \big| u = 0 \quad \forall f, u, \quad \forall t \big| t^{d}_{f,t} \le u \tag{11}$$

In addition, the time slot assignment cannot be later than the latest time for which it would be reasonable to accept an assignment at the FCA considering the geometry of its secondary route:

$$y^{h}_{f,t} \mid u = 0 \quad \forall f, u, \quad \forall t > l_f \tag{12}$$

The final option possible, if the capacity increase occurs at time u, is to leave a flight that was originally assigned to its secondary route on that route, with the same (scheduled) departure time:

$$y^{s}_{f} \mid u = \begin{cases} 1, & \text{if flight } f \text{ was originally assigned to its secondary route, and if, under AFP stop time } u \text{, that decision remains unchanged} \\ 0 & \text{otherwise} \end{cases}$$

Practically speaking, it would never make sense to assign a flight to its secondary route under the recourse if it had not also been given the same assignment in the initial plan. It might seem, therefore, that the following constraint is necessary:

$$y^{s}_{f} \mid u \le x^{s}_{f} \quad \forall u, f \tag{13}$$

However, it should be obvious that the objective function will enforce this behavior implicitly. If it was cost effective to assign a flight to its secondary route under the recourse, it would be cost effective to do so under the initial plan.

ISBN: 978-0-615-20720-9

Constraints 10 and 13 can be combined into a single constraint:

$$y^h_{f,t} \mid u + y^s_f \mid u \le x^s_f \quad \forall u, f, t \tag{14}$$

It would be possible, given the constraints developed so far, to assign a flight to a hybrid route that essentially reverts to the primary route immediately. In other words, this would be an assignment that is tantamount to taking off on the primary route at the scheduled departure time, which is a more logical way to interpret this outcome. Therefore we introduce the following constraint to enforce this behavior:

$$y^h_{f,D_f+E_f} \mid u = 0 \quad \forall f, u \tag{15}$$

For each time scenario *u*, every flight *f* must be assigned to one of these dispositions. Furthermore, if the disposition involves being scheduled into a slot appointment at the FCA, no more than one slot can be assigned to a given flight. Given that the decision variables are required to be binary, the following constraint addresses both of these concerns:

$$\sum_t y^p_{f,t} \mid u + \sum_t y^h_{f,t} \mid u + y^s_f \mid u = 1 \quad \forall u, f \tag{16}$$

For any value $U=u$, there will be a new capacity profile $C^u(t)$ that agrees with $C^0(t)$ up to time $t=u$, but represents an increase in capacity beyond that point. For example, if $C^0(t)$ had been a constant vector, then $C^u(t)$ could be a step function that makes a jump at time $t=u$. On the other hand, if $C^0(t)$ had been a periodic 0-1 function, then $C^u(t)$ might just have an increased duty cycle after time $t=u$. Figure 4 shows examples of both of these extremes. A wide variety of profiles for $C^u(t)$ are possible; the only real requirements are that it agree with $C^0(t)$ prior to time $t=u$, and that after that time, it supports a higher rate of flow than that was possible under the initial plan. The capacity constraint under the scenario $U=u$ can now be written as:

$$\sum_f y^p_{f,t} \mid u + \sum_f y^h_{f,t} \mid u \le C^u_t \quad \forall u, t \tag{17}$$

3) Objective Function

Since our model involves the specification of decisions that are conditioned random events, the objective function will be an expected value. To emphasize the paradigm of creating a plan (our initial plan) together with contingency plans (our recourse actions), we represent the objective function as the sum of the deterministic cost of the initial plan minus the expected savings from recourse actions.

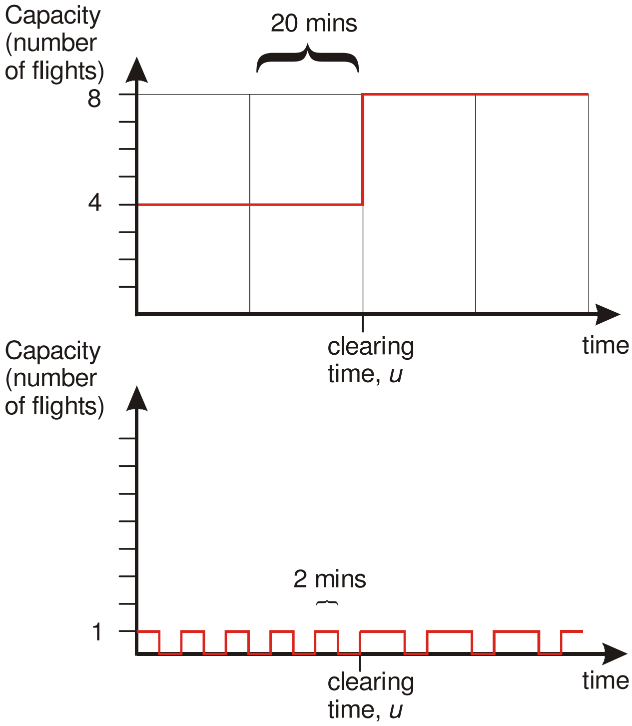

Figure 4. Capacity functions with high (top) and low (bottom) throughput before and after clearing.

Therefore the objective function can thus be represented as:

$$Min\left[C(X) - \sum_u P_u S(Y_u)\right] \tag{18}$$

Or more precisely:

$$Min \quad Z = z^1 + z^2 - \sum_u P_u\left(z^3_u + z^4_u\right) \tag{19}$$

Where,

$$z^1 = \sum_f \sum_t c^p_{f,t} x^p_{f,t} \tag{20}$$

$$z^2 = \sum_f c^s_f x^s_f \tag{21}$$

$$z^3_u = z^1 - \sum_f \sum_t c^p_{f,t} y^p_{f,t} \Big| u + \sum_f \sum_t c^s_f s^p_{f,t} \Big| u \tag{22}$$

$$z^4_u = \sum_f \sum_t sv^h_{f,t} y^h_{f,t} \Big| u \tag{23}$$

where

$c^p_{f,t}$ is the cost of assigning flight f to its primary route so that it arrives at the FCA at time *t*.
c^s_f is the cost of assigning flight f to its secondary route.
$sv^h_{f,t}$ is the savings incurred if flight *f* starts out on its secondary route but reverts to a hybrid route that arrives at the FCA at time *t*.
$s^p_{f,t}$ is a dummy binary variable that works as an indicator. It takes value of one when a flight initially assigned to its

ISBN: 978-0-615-20720-9

secondary route is assigned back to its primary route under revised plan.
So;

$$s^{p}_{f,t} = Min\left(x^{s}_{f}, y^{p}_{f,t}\right) \qquad (24)$$

IV. COMPUTATIONAL EXPERIMENT

We conducted a computational experiment to give some preliminary evidence as to the computational feasibility of the model and its impact on decision making. We now describe the problem data. Flights, their routes and alternate routes were generated artificially based on the airspace geometry given in Figure 5. There were three types of flights:

Short haul: length - 60 min: origin-to-FCA – 30 min, FCA-to-destination – 30 min; reroute angle – arctan(30/30).
Medium haul: length – 180 min: origin-to-FCA – 90 min, FCA-to-destination – 90 min; reroute angle – arctan(30/90).
Long haul: length – 300 min: origin-to-FCA – 150 min, FCA-to-destination – 150 min; reroute angle – arctan(30/150).

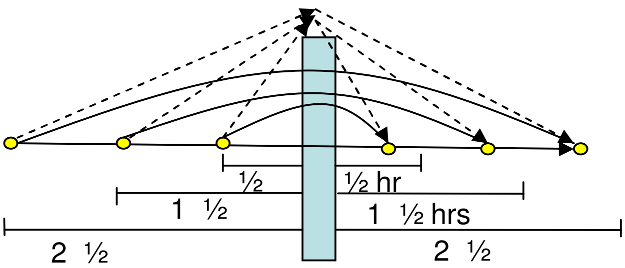

Figure 5: Airspace Geometry for Flight Generation

There were F=200 flights with one flight departing every 1 minute and departures alternating among the three flight types. First flight departed at 2:00 PM. There were T=200 time slots; each slot had a width of Δt =2 minutes. Initially, the FCA had restricted capacity of 1 flight per every three time slots (10 flights per hour). In all cases, the FCA cleared by 7:00 PM so that capacity rose to 4 flights per time slot (120 flights per hour). There were four possible early clearance times: 3:00 PM, 4:20 PM, 5:30 PM, each occurring with probability 0.25. In event of early clearance, slot capacity rose from 1/3 to 2 flights for slots between the clearance time and 7:00 PM.

Three cases were run:
All Options: This was the complete model as defined in the paper.
Reroute but No Recourse: In this case, the reroute option did not include the possibility of recourse, i.e. the $y_{f,t}^{h}$ variables are all fixed at zero. This corresponds to a decision making scenario where the possibility of rerouting after departure is not taken into account.
No Reroute: In the case, no rerouting is allowed. This corresponds to a decision making scenario under which the problem is solved only using ground delays.

Table 1 gives the results of our experiments. The costs of various solution components as well as the total expected cost are given. Note that the "All Options" scenario produced substantial cost savings over the other cases (particularly the No Reroute case). The fact that the initial plan costs (cost of the assigned ground delay $C(x_p)$ and cost of complete reroutes $C(x_s)$) changed significantly among the cases shows that taking the various recourse options into account can substantially alter the initial plan.

Also note that running times are given. A 2.8 GHz Intel® Pentium® based computer was used with 1.99 GB of RAM. The IP solver used was XPress MP® vers 2007B.

	All Options			Reroute but No Recourse			No Reroute		
air/ground delay cost	4			4			4		
Objective Function	2379.4			2622.8			6808		
C(xp)	187			455			11381		
C(xs)	764.7			722.0			0		
U:Clearance time	3:00 PM	4:20 PM	5:40 PM	3:00 PM	4:20 PM	5:40 PM	3:00 PM	4:20 PM	5:40 PM
Probability of clearance	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25
SV(yp)	61	114	0	345	364	110	10117	6247	1928
SV(yh)+SV(xs-->yp)	595.1	207.8	19.7	446.6	68.7	0.0	0.0	0.0	0.0
GD(u)	126	73	187	110	91	345	1264	5134	9453
AD(u)	169.6	556.9	745.0	275.4	653.3	722.0	0.0	0.0	0.0
C(u)	804.4	2300.6	3166.9	1211.6	2704.0	3232.8	1264.0	5134.0	9453.0
E[cost]	2379.4			2622.8			6808.0		
E[GD]	143.3			250.3			6808.0		
E[AD]	559.0			593.1			0.0		
n(xp=1)	65			73			200		
n(yp=1)	148	77	65	150	84	73	200	200	200
n(xs=1)	135			127			0		
n(yh=1)	41	55	10	0	0	0	0	0	0
n(ys=1)	11	68	125	50	116	127	0	0	0
n(xs-->yp)	83	12	0	77	11	0	0	0	0
Running Time (sec)		283			226		428		

Table 1: Computational Results

V. CONCLUSIONS AND FUTURE WORK

In this paper, we have defined the basics of a stochastic optimization model for simultaneously making ground delay and reroute decisions in response to en route airspace congestion. We have also given the results of an initial computational experiment. Future steps should include more computational experiments and model refinements aimed at improving the computational performance of the integer program and at exploring the changes in airspace planning the model provides. We anticipate the need to provide many refinements and extensions to this model to better address practical problem solving. Further, another vital direction is the development of strategies necessary to embed this model within CDM processes necessary for the delivery of practical air traffic flow management solutions.

VI. ACKNOWLEDGEMENTS

The authors gratefully acknowledge the support of the National Aeronautics and Space Administration Airspace Systems Program under ARMD NRA: NNH06ZNH001.

 ISBN: 978-0-615-20720-9

References

.

[1] M.O. Ball, R. Hoffman, A. Odoni and R. Rifkin, "A Stochastic Integer Program with Dual Network Structure and its Application to the Ground-Holding Problem", *Operations Research*, 51,167-171, 2003

[2] G. Davidson and J. Krozel, "Strategic Traffic Flow Management Concept of Operations", Aviation Technology, Integration and Operations (ATIO), 2004.

[3] R. Hoffman, J. Krozel, G. Davidson and D. Kierstead, "Probabilistic Scenario-Based Event Planning for Traffic Flow Management". AIAA Guidance, Navigation, and Control Conference, August 2007.

[4] R. Hoffman, M. O. Ball and A. Mukherjee, "Ration-by-Distance With Equity Guarantees: A New Approach to Ground Delay Program Planning and Control", Proceedings of the USA/Europe Air Traffic Management R&D Seminar, ATM 2007.

[5] B. Kotnyek and O. Richetta, "Equitable Models for the Stochastic Ground-Holding Problem Under Collaborative Decision Making", *Transportation Science*, 40, 133-146, 2006.

[6] A. Mukherjee and M. Hansen, "A dynamic stochastic model for the single airport ground holding problem", Transportation Science, 41 444-456, 2007.

[7] A. Mukherjee and M. Hansen, "A dynamic rerouting model for the air traffic flow managment", Transportation Research - Part B, in press.

[8] A. Nilim and L. El Ghaoui, "Algorithms for Air Traffic Flow Management Under Stochastic Environments," IEEE American Control Conference, 2004.

[9] A. Nilim, L. El Ghaoui and V. Duong, "Multi-Aircraft Routing and Traffic Flow Management under Uncertainty", Proc. ATM 2003, Budapest, Hungary, 2003.

[10] A. Nilim, L. El-Ghaoui, V. Duong and Hansen, M., "Trajectory-Based Air Traffic Management (TB-ATM) Under Weather Uncertainty," Proceedings of the 4rd USA/Europe Air Traffic Management R&D Seminar, Dec. 2001.

[11] O. Richetta and A. R. Odoni, Solving optimally the static ground-holding problem in air traffic control. *Transportation Science* **27**, 228-238, 1993.

ISBN: 978-0-615-20720-9