The Static Stochastic Ground Holding Problem

Robert L. Hoffman *
Ryan M. Rifkin †Michael O. Ball ‡Amedeo R. Odoni §

1 Introduction

2 Model Formulation

Let $t = 1, 2, ..., T$ be a set of discrete, contiguous time intervals during which the AAR of an airport is in jeopardy. Each t could represent, say, a 15-minute time interval. Let $\overline{D}_t$ be the total number of flights scheduled to arrive during time interval t. In practice, not all flights can be given a CTA (controlled time of arrival). International flights and general aviation, for instance, are exempt from ground delay programs in the United States. Let D_t be the total number of (non-exempt) flights that will be included in the GDP. Then the number of exempt flights is given by $E_t = \overline{D}_t - D_t$ with $E_t \geq 0$.

When formulating a GDP, the decision-maker (DM) has the option to exempt flights from the program based on criteria other than those already mentioned. Typically, these criteria are related to the proximity of the origination airport to the destination airport. For instance, when formulating a GDP for an airport on the east coast of the United States, the DM may choose to exempt all long-distance flights, hence, all flights originating on the west coast. So, we assume that the set of flights that are candidates for a given GDP is partitioned into disjoint classes $e = 1, 2, ..., E$, called **tiers**. For each tier, we define a decision variable y_e, such that $y_e = 1$, if tier e is to be included in the GDP and $y_e = 0$, otherwise. Let d_t^e be the number of flights in tier e. Then we have $D_t = \sum_{e=1}^{E} d_t^e y_e$.

*University of Maryland
†Massachusetts Institute of Technology
‡University of Maryland
§Massachusetts Institute of Technology

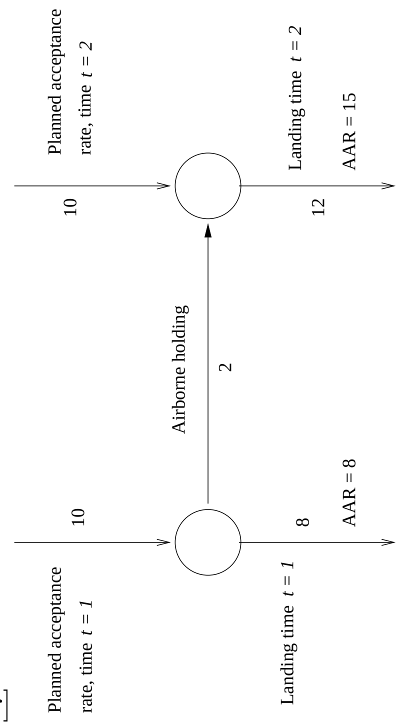

[r]

Figure 1: Airborne holding of two flights

Let A_t be the PAAR (planned arrival acceptance rate) for time period t, i.e., the number of flights that will be assigned a controlled time of arrival that falls within time period t.

Because of the stochastic nature of weather, A_t flights will not necessarily be able to land during time interval t. For instance, if $A_1 = 10$ and the airport can land only 8 flights at time $t = 1$, then 2 of the 10 flights scheduled to arrive at time $t = 1$ will be held in the air by the controllers at the destination airport. Let us assume that $A_2 = 10$ and that the capacity of the airport is, say, 15 flights for the next period so that the two airborne-held flights are able to land in the next time interval, $t = 2$. Then the airborne delay under this single AAR scenario can be depicted by the flow diagram in Figure 1.

Assuming that ground holding is cheaper than airborne holding (otherwise there is no need for a GDP), it would have been cheaper to hold two flights on the ground for one time unit and have them arrive at the terminal airspace at time $t = 2$. The single-unit ground delay of these two flights can be represented by adding a horizontal arc at the top of Figure 12.

Note that Figure 2 treats flights on an aggregate level, with no mention of which flights

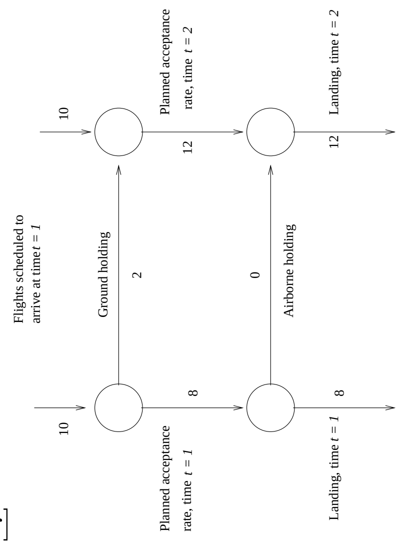

[r]

Figure 2: Ground holding of two flights

will absorb the ground delay. Recall from the introduction that the flights are sequenced by landing order prior to the formulation of a GDP. Thus, it is easily deduced that the first eight flights are allowed to land at time interval $t = 1$ while the ninth and tenth flights will absorb the ground delay and land in time interval $t = 2$.

Figure 2 is easily extended to allow for an arbitrary number of time periods, T. Moreover, arcs can been added to allow for the exemption of flights, as in Figure 2. The exempted flights flow directly to the bottom nodes, indicating that they will be accepted to the terminal airspace under any conditions. These flights are not exempt, however, from possible airborne holding. Figure 2 assumes the following notation.

E_t = number of flights scheduled to arrive at time t that are exempt from ground delay.

b_t = number of flights that can be landed at the airport (the AAR) during time period t.

G_t = number of (non-exempt) flights wishing to land at time t but not accepted to the terminal airspace at time t (these flights have been ground delayed).

W_t = number of flights held in the air from time period t to $t+1$ (these flights have been airborne delayed).

Next, we extend Figure 2 to accommodate multiple AAR scenarios. Let $s = 1, 2, ..., S$ be a set of AAR scenarios, each with an associated probability, p_s. The capacity of the airport (AAR) at time t varies with the scenario, s, hence, so does the amount of airborne holding at time t. We generalize the above notation as follows.

W_t^s = number of flights held in the air from time period t to $t+1$,under scenario s.

b_t^s = number of flights that can be landed at the airport (the AAR) during time period t, under scenario s.

p_s = the probability of scenario s occurring.

We obtain Figure 2 from Figure 2 by replicating the lower portion of Figure 2 once for each scenario. Our integer programming model will be based on this diagram.

Deterministic flow model (S=1)5in

Stochastic flow model (S=2)5in

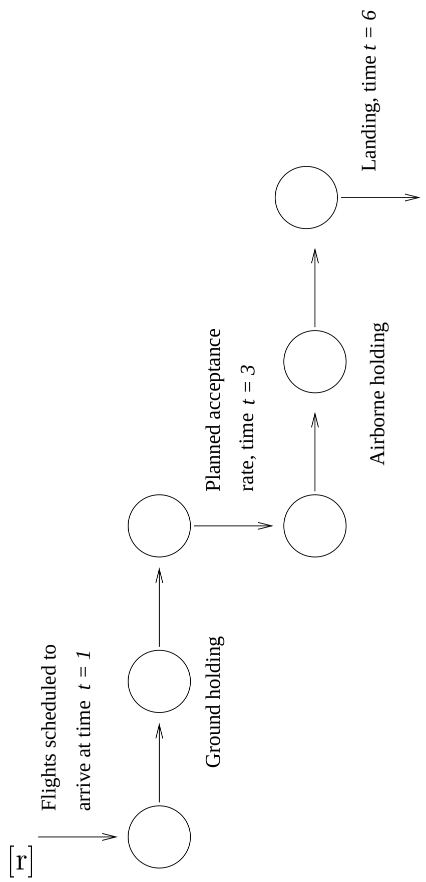

Figure 3: The path of a single flight

The hypothetical path of a single flight through the flow diagram is depicted in Figure 3. The flight enters the system at time $t = 1$, indicating that it is originally scheduled to arrive at time $t = 1$. Next, the flight absorbs two units of ground delay, then is allowed to enter the terminal airspace and absorbs two unit of airborne delay. Note that the flow pattern gives the illusion that the flight is being held on the ground just prior to arriving at the terminal space (attempting to land). In reality, it will absorb its ground delay at its origin airport at some earlier time, which is easily computed based on the estimated enroute travel time for the flight.

Let c_g and c_a be, respectively, the costs of holding a flight for one time unit on the ground and one time unit in the air. Given an outcome (scenario), s, the total delay cost of the GDP is given by $\sum_{t=1}^{T}\left(c_g G_t + c_a W_t^s\right)$, and the expected total delay cost of the GDP is given by $\sum_{s=1}^{S}\left[p_s \cdot \sum_{t=1}^{T}\left(c_g G_t + c_a W_t^s\right)\right]$.

In formulating a GDP, the DM tries to minimize the total expected delay cost subject to the constraints imposed by Figure 2. We now develop those constraints.

First, we model the flow through the upper (round) nodes of Figure 2, using the following relation:

$$G_{t-1} + D_t = A_t + G_t \qquad t = 1, \ldots, T, \ G_0 = 0 \tag{1}$$

To see that these equations are valid, recall that A_t is the number of flights that will be accepted to the terminal (airport) airspace (but not necessarily land) at time t, that D_t is the number of flights scheduled to land, and that W_{t-1} is the number of flights in the airport airspace at time $t-1$ that were unable to land. Then $W_{t-1} + D_t$ is the number of flights wishing to land at time t and (1) simply says that this number is equal to the number accepted to the airspace at time t (A_t) plus the number rejected from the airspace at time t (G_t).

We should point out that, unless $S = 1$, the equations of the form 1 are not be based on conservation of flow. For each scenario s, there is one arc flowing out the bottom of any given upper (round) node, and each of these S arcs has flow A_t. This labeling is justified by the fact that once the planned arrival acceptance rate A_t is set, A_t flights are expected to arrive at the destination airport no matter which scenario occurs. (Each of these arcs could have been superscripted with their respective scenarios (A_t^s), in which case we would require coupling constraints of the form $A_t^s = A_t^{s'}$, for $s \neq s'$. The formulation presented here avoids the need for these constraints.) As a consequence, there could be more flow out of the one of these nodes than into it.

For convenience, we define u_t^s to be the number of flights that arrive at time t under scenario s (this variable need not appear in the final formulation). This is the flow out the bottom of Figure 2 and is limited by the arrival capacity at time t under scenario s, b_t^s. Thus, we have that

$$u_t^s \leq b_t^s \qquad s = 1, \ldots, S, \ t = 1, \ldots, T \tag{2}$$

By conservation of flow at the lower nodes, variables of the form u_t^s can be defined as follows.

$$u_t^s = W_{t-1}^s + A_t + E_t - W_t^s \qquad s = 1, \ldots, S, \ t = 1, \ldots, T, W_0^s = 0 \tag{3}$$

Substituting (3) into (2), we obtain the following set of constraints.

$$W_{t-1}^{s} + A_t + E_t - W_t^{s} \leq b_t^{s} \qquad s = 1, \ldots, S,\ t = 1, \ldots, T, W_0^{s} = 0 \tag{4}$$

Since the flow across each arc represents a number of flights, this value must be integer and non-negative. In all, we have the following integer program.

(SGH)

$$Minimize\ \sum_{t=1}^{T} \left(c_g G_t + \sum_{s=1}^{S} c_a A_t^{s} \cdot p\left(s\right) \right) \tag{5}$$

$$subject\ to$$

$$(1), (4)$$

$$G_t \in \mathcal{Z}_+, A_t \in \mathcal{Z}_+, W_{s,t} \in \mathcal{Z}_+$$

Since the model treats flights on an aggregate level, the total number of constraints is independent of the number of flights and kept quite small: $O\left(S \cdot T\right)$. The number of variables is also $O\left(S \cdot T\right)$.

For the case $S = 1$, the model is deterministic and the diagram in Figure 2 reduces to the diagram in Figure 2. The deterministic case $S = 1$ is easily solved by a greedy algorithm which lands as many flights as possible during each time period, remembering to take exempt flights into account. When $S \geq 2$, the problem is truly stochastic and the greedy algorithm is ill-defined because the phrase "land as many flights as possible" becomes ambiguous.

The model resembles a classic production-inventory model in which an inventory can be held in one of two states (see Section 4.5 of [7]such as crude oil. A quantity of materials (the flights) is purchased when it enters the diagram at a time period and can then be held in raw form (ground holding) or finished form (airborne holding). The materials are sold when they exit the diagram at the bottom. But the model presented here is distinguished in three ways. First, it does not seek production levels to be set for each time period; these levels are set in advance by the parameters of the form D_t. Second, and more crucially, the product can be held over (in two different types of inventory) at two different costs, as in a raw material/finished good production model such as crude oil/refined oil. However, the amount held over in the second form is stochastic, depending on a random variable. This prevents the model from being solved by standard techniques. (See [11] for a treatment of stochastic network routing.) Third, the model cannot be translated into a network model, even though it is depicted by a flow diagram. (The dual of the model, however, can be, as shown in Section 5.)

3 Pre-processing

In this section, we show that for each tier e and each time interval t, the variables y_e and E_t can be eliminated from the formulation of SGH by pre-processing. We start with the variables y_e.

Part of the output of model SGH would be values for the variables y_e. From an optimization standpoint, the exemption of tier e from the program can only worsen the objective function because one has lost the ability to ground delay the flights of tier e. Therefore, there will be an optimal solution to SGH satisfying $y_e = 1$ for all e, as the following proposition shows:

Proposition 1 *For every feasible solution,* $\mathbf{z}$, *to* SGH*, there is a feasible solution,* $\mathbf{z}^*$, *with the properties that* $F(\mathbf{z}^*) = F(\mathbf{z})$ *and* $y_e = 1$ *for* $e = 1, 2, ..., E$.

PROOF:*Let* $\mathbf{z}$ *be any feasible solution to* SGH *with corresponding function value* $F(\mathbf{z})$. *If* $y_e = 1$ *for every* e*, then we set* $z = z^*$ *and there is nothing to show. Otherwise, we create a new solution,* $\mathbf{z}^*$*,as follows. For each index* e *such that* $y_e = 0$ *in solution* $\mathbf{z}$*, we set* $y_e = 1$. *By definition* D_t*, this increases the value of* D_t *by the amount* d_t^e *and decreases the value of* E_t *by the same amount. This jeopardizes equality in equation (1), but equality can be maintained by adding the amount* d_t^e *to the value of* A_t*. This also preserves feasibility in constraints (4). Since the values of* W_t^s *and* G_t *remain unchanged for every* s *and every* t*, we have created a new feasible solution,* $\mathbf{z}^*$ *with the same function value as* $\mathbf{z}$*. The result follows from the fact that* y_e *does not appear in the objective function and that the index* e *was chosen arbitrarily.*$\square$

As a consequence of Proposition 1, we can make the simplifying assumption that $y_e = 1$ for $e = 1, 2, ..., E$ and discard these variables from the formulation of the problem. In practice, this means that if the values of the variables y_e are set a priori by the DM (as opposed to being output by the solution), then optimality may be sacrificed.

We comment that proposition 1 would not hold in an alternative model in which the variables y_e appear in the objective function with associated penalties.

Now we show how to eliminate through preprocessing the variables of the form E_t, which reflect the number of flights scheduled to arrive at the airport at time t. For each scenario s, we perform the following procedure, shown here in pseudocode:

$$
\begin{array}{l}
t \leftarrow 1 \\
C_E \leftarrow 0 \\
E^* \leftarrow E \\
while(t < T)\{ \\
\quad E^*_{t+1} \leftarrow E^*_{t+1} + max(0, E^*_t - b^s_t) \\
\quad C_E \leftarrow C_E + c_a \cdot max(0, E^*_t - b^s_t) \\
\quad b^s_t \leftarrow b^s_t - \min(b^s_t, E^*_t) \\
\quad E^*_t \leftarrow 0 \\
\quad t \leftarrow t + 1 \\
\}
\end{array}
\tag{6}
$$

Intuitively, for each time period t, we subtract the number of exempt flights from the capacity for that time period. If there are more exempt flights than capacity, we set the capacity to zero and add any remaining exempt flights to the next time period. Once we have performed the above procedure for every scenario s, we set all the E_t to 0, yielding a new model. This model has the same set of feasible solutions as our original model, and the cost of any feasible solution in the two models differs by C_E, the total costs that must be added to offset those cases where "excess" exempt flights are moved forward from one time period to the next in the transformation — these flights would have experienced airborne delay in our original model. If we are interested only in the optimal solution itself, rather than the cost of the optimal solution, we can ignore the value C_E (see Section 4).

Henceforth, we assume that in any formulation of SGH, there are no exempt flights so that $\bar{D}_t = D_t$, and that $y_e = 1$ for every e. We have the following simplification of SGH:

(SGH_2)

$$\min \sum_{t=1}^{T} \left(c_g G_t + \sum_{s=1}^{S} c_a W_t^s \cdot p(s) \right)$$

subject to

$$G_{t-1} + D_t = A_t + G_t \qquad t = 1, \ldots, T, G_0 = 0$$

$$W_{t-1}^s + A_t - W_t^s \le b_t^s \qquad s = 1, \ldots, S,\ t = 1, \ldots, T, W_0^s = 0$$

$$G_t \in \mathcal{Z}_+, A_t \in \mathcal{Z}_+, W_{s,t} \in \mathcal{Z}_+$$

4 Holding Costs

Practical implementation of our model would require that values be fixed for the ground and airborne holding costs. Prior to this work, the common interpretation of the ground and airborne holding costs were as actual costs incurred by the airlines. Any assessment of these costs would be highly complex and situation specific, and would involve complex factors such as crew costs, fuel consumption, connectivity to other flights, and passenger safety. These types of calcualtions are very hard to make for every flight in real time; additionally, it is only by aggregating flights as a flow rather than treating flights individually that our models become tractable. Richetta pioneered this approach in [12], where he proposed that flights be grouped into a small number (e.g., three) of classes and that each class would have a single ground delay cost and a single airborne delay cost representative of the whole class.

There are figures to support the implementation of this approach. For instance, the Air Transportation Association estimated that for a typical flight, every minute of ground delay costs \$20.35 while every minute of airborne delay costs \$45.85.

On the other hand, one of the fundamental tenets of collaborative decision making is that the FAA should not make economic decisions on behalf of the airlines. Since our model is designed for FAA use, it is inappropriate to view the holding costs as the financial costs incurred by the airlines. We instead suggest the following interpretation.

The primary mission of the FAA, hence, the ATCSCC (Air Traffic Control Systems Command Center), is to ensure passenger safety while simultaneously managing resources within the NAS (national air space) in an efficient, even-handed manner. These two objectives are inherently somewhat in conflict. A ground delay program which was aggressive enough to ensure that, under any reasonable circumstances, no flight was subjected to airborne holding would, on average, lead to a gross under-utilization of airport arrival capacity and the services provided by the air traffic controllers. On the other hand, running no ground delay programs whatsoever would allow any available airport arrival capacities to be maximally exploited, but this would clearly pose unacceptable risks to passenger safety. The DM must strike a balance between restricted traffic flow and efficient allocation of resources.

In so doing, the decision maker at the ATCSCC must address airline needs and perceptions. In particular, an airline finds it aggravating when a GDP is aborted in mid-operation. This happens whenever the DM has clear evidence that the original capacity and weather forecasts were overly pessimistic. Of course, the GDP was unnecessary only in hindsight but it is hard to justify ground delay to passengers who see "blue-sky" conditions at their departure airport. For this reason, the airlines are frequently willing to risk small or moderate amounts of airborne delay and tend to favor relatively liberal ground holding policies.

Our alternative perspective, then, on the delay cost ratio (air to ground) is that of a numerical expression of the willingness to trade ground holding for air holding, in light of the twin objectives of safety and manageability on the one hand and efficient resource use on the other. The higher the cost ratio, the more conservative would be the resulting GDP. In particular, we envision that the specialist at the ATCSCC might explore several values of the cost ratio before implementing a ground delay program.

Lastly, we point out that the structure of the objective function allows the delay costs, both airborne and ground, to vary with t. This grants the decision maker the flexibility to establish time intervals of varying length. For instance, one could use periods of increasing length so that short-term planning is done on a refined level while long-range planning is more coarse.

5 Theoretical Results

In this section, we show that the integer programming problem SGH_2 can be solved to optimality in polynomial time. In particular, we will show that SGH_2 is a dual network flow problem. As immediate corollaries, the constraint matrix associated with SGH_2 is totally unimodular, and the LP relaxation of SGH_2 [edit] yields integral solutions.

We say that an m row, n column $(0, 1, -1)$ matrix M is a *network matrix* if there exists a directed tree R on $m+1$ nodes and a one-to-one mapping of the rows of our matrix onto the edges of R with the property that each column of the matrix corresponds to the characteristic vector of a path in R. This definition was introduced in [16] and is explored in [9]. The more familiar *node-arc incidence matrices*, with a single 1 and -1 in each column, are a special case of this construction. Network matrices are desireable because they are totally unimodular and give rise to integral polyhedra.

Theorem 2 *Let A^T be the transpose of the constraint matrix associated with SGH_2. A^T is a network matrix.*

PROOF SKETCH: *We will illustrate the construction of the tree R for the case of two time periods and two scenarios $(T = 2, Q = 2)$. It will be clear that this construction can be extended to an arbitrary number of time periods and scenarios. The general construction is explicitly formulated as a matrix transformation in [8], providing a more formal proof [edit].*

For a two time period, two scenario problem, with slack variables, the primal problem is:

$$\min \quad c_g \cdot G_1 + c_g \cdot G_2 + c_a \cdot W_{11} + c_a \cdot W_{12} + c_a \cdot W_{21} + c_a \cdot W_{22}$$

$$\begin{array}{cccccccccccccc}
A_1 & & +G_1 & & & & & & & & & & = & D_1 \\
 & A_2 & -G_1 & +G_2 & & & & & & & & & = & D_2 \\
-A_1 & & & & +W_{11} & & & & -S_{11} & & & & = & -M_{11} \\
 & -A_2 & & & -W_{11} & +W_{12} & & & & -S_{12} & & & = & -M_{12} \\
-A_1 & & & & & & +W_{21} & & & & -S_{21} & & = & -M_{21} \\
 & -A_2 & & & & & -W_{21} & W_{22} & & & & -S_{22} & = & -M_{22}
\end{array}$$

We write the matrix associated with this problem as:

$$
A = \begin{bmatrix}
A_1 & & G_1 & & & & & & & & & \\
& A_2 & -G_1 & G_2 & & & & & & & & \\
-A_1 & & & & W_{11} & & & & -S_{11} & & & \\
& -A_2 & & & -W_{11} & W_{12} & & & & -S_{12} & & \\
-A_1 & & & & & & W_{21} & & & & -S_{21} & \\
& -A_2 & & & & & -W_{21} & W_{22} & & & & -S_{22}
\end{bmatrix}
$$

All the entries in A should be 0, 1, or -1; the variable names have been added to mark each column with its corresponding variable [edit]. To show that A^T is a network matrix, we must find a directed tree R on 13 nodes such that every column of A^T (equivalently, every row *of A) corresponds to the characteristic vector of a path in R. Figure 4 shows the required tree. We have labeled each arc of R with the corresponding primal variable [edit]. It can be verified [edit] that every row of A [edit] corresponds to the characteristic vector of a path in R. For example, the fourth row of A [edit] corresponds to the path formed by traversing [edit] arc S_{12} backward, then arc W_{12} forward, then arcs A_2 and W_{11} backward. Figure 5 shows the graph G spanned by the tree R. The added arcs correspond to the associated dual cost coefficients (the right-hand-side of the primal) [edit]. This construction extends to an arbitrary of number of scenarios and time periods [edit].*

Corollary 3 *The constraint matrix associated with SGH_2 is totally unimodular.*

Corollary 4 *The LP relaxation of SGH_2 yields integral solutions [edit].*

Corollary 5 *The integer program SGH_2 can be solved to optimality in polynomial time.*

The dual problem can be recast in a more familiar form as a min-cost flow problem by negating the objective function and multiplying the dual LP in equality form by N, where N is the node-arc incidence matrix associated with R with a single row deleted. Figure 6 shows the resulting flow problem for the previous example ($T = 2, Q = 2$). The labeled arcs can accomodate either positive or negative flow, as they are associated with equality constraints in the original primal problem; the flow must be nonnegative on the unlabeled arcs, which are associated with dual slack variables. The unlabeled arcs are zero-cost arcs and the unlabeled nodes are neither sources nor a sinks [edit]. Unfortunately, Figure 6 gives

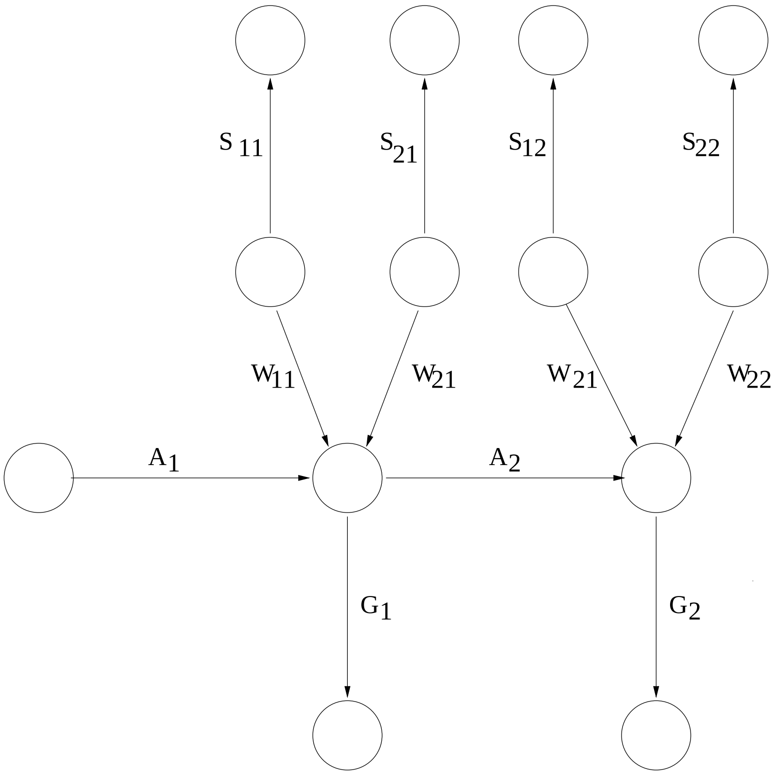

Figure 4: The dual tree R

little or no insight into the original problem because the matrix transformation destroys the relations between the arcs in the mincost flow problem and the original primal variables.

We should point out that, unlike the dual problem, the primal problem cannot be cast as a network flow problem. This can be verified by applying the network recognition algorithms provided in either [6] or [9]. Also, [8] shows directly that the flow through the middle nodes in Figure 2 cannot be described by conservation of flow equations.

6 Computational Results and Implementation

We constructed three realistic instances of SGH_2 using the following guidelines. Each test case was solved for two levels of granularity of time periods, for a total of six test cases. For each instance, we constructed three AAR scenarios. In test case 1, the AAR's were randomly generated numbers in the range 30 - 60. In test case 2, each AAR scenario drops from 60 flights per hour to 40 flights per hour for several hours, then returns to 60 flights per hour.

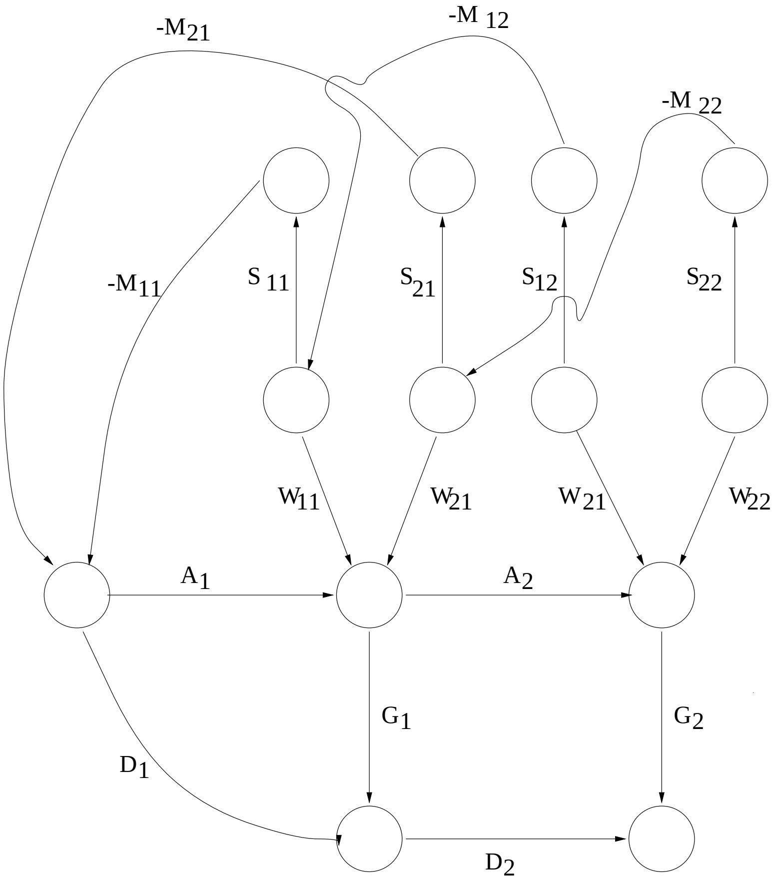

Figure 5: The dual tree R, with columns of A^T

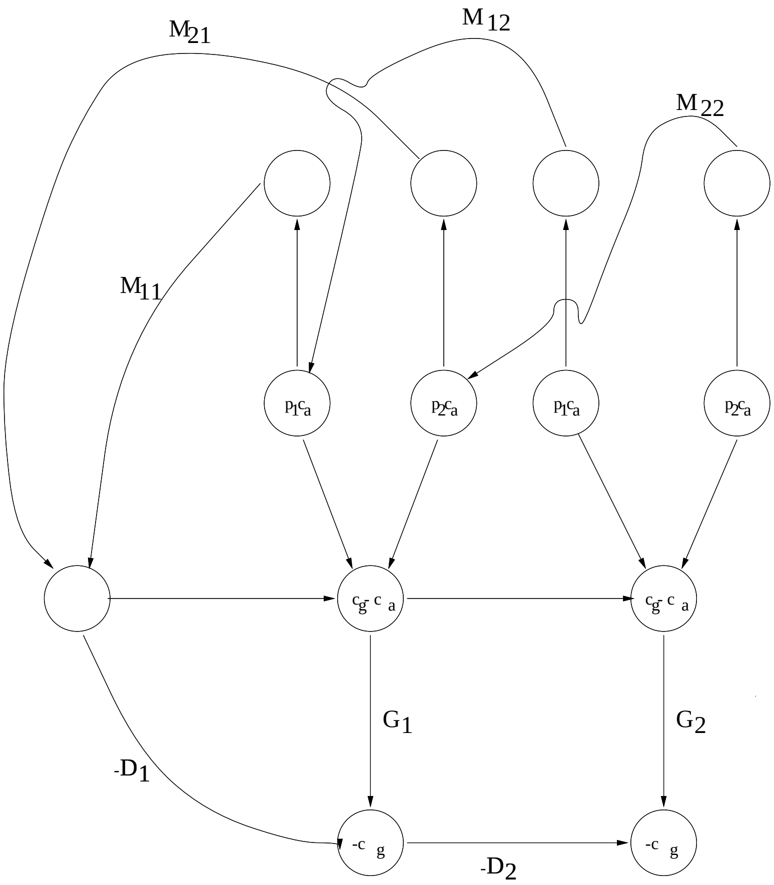

Figure 6: The transformed min-cost flow problem

The times were varied at which each AAR scenario drops and recovers. In test case 3, the first scenario is designed to be a fair weather scenario, the second is a bad weather scenario and the third is randomly generated.

In each instance, the costs used were 2.0 for ground holding and 5.0 for airborne holding. Each instance was solved using CPLEX 3.0 with the default options on a SPARC STATION 10.

The results in Table 1 show that the problem is highly tractable. The integer solution was obtained in zero nodes of the mixed integer program algorithm of CPLEX, meaning that the integer solution was obtained directly from the linear program relaxation. This is to be expected, given Corollary 5 to Theorem 2. The largest number of iterations of the simplex procedure was 289 and the longest run time was barely more than half a second. Note that the run time is almost linear in the coarseness (length of) the time periods.

We would expect that an implementation that directly exploited the dual network structure of the problem would result in substantially faster runtimes for larger problems. Because the LP model is solved so quickly on problem sizes currently considered realistic, this is unnecessary. Our goal is to deliver a model to the FAA which can run "at the push of a button" using only modest computing power; the current implementation achieves these goals.

Ideally, we would hope that a small increase in the cost ratio would lead to small changes in the optimal solution. However, under certain circumstances, the model exhibits behavior in which the optimal solution is constant over a fairly wide range of cost ratios, then, when the cost ratio reaches some critical value, the character of the optimal solution changes completely. An extreme example of this effect can be seen when capacity is constant within each scenario. In this case, the optimal solution will always mimic one of the scenarios precisely. For a study of this and several other examples illustrating the behavior of optimal solutions to SGH_2, see either [8] or [14]. We note that this effect should be considered pathological — for a reasonable number of substantively different scenarios, the effect will not be observed.

Test Case	Numb flights	Numb periods	Min per period	Opt Func Value	Time (sec)	Simplex Iteration	Node B&B
1a	624	12	60	642.00	0.07	51	0
1b	624	48	15	3329.00	0.30	267	0
2a	624	12	60	1088.00	0.07	59	0
2b	624	48	15	4209.00	0.37	289	0
3a	624	12	60	485.00	0.05	55	0
3b	624	48	15	1745.50	0.20	223	0

Table 1: SGH model performance

7 Conclusion

The stochastic ground holding model presented in this paper finds the optimal trade-off between airborne-holding and ground holding in the formulation of a ground delay program. By treating flights on an aggregate level, as a flow, the integer programming model avoids the need to consider the assignment of every flight to each time period in the planning horizon. Additionally, the fact that the model is dual network allows us to relax the integrality constraints, as we are guaranteed integer optimal solutions directly from the LP relaxation. The combination of these two facts result in a highly tractable model which requires only modest computing power. Additionally, the model is designed to be an integral component of the collaborative decision-making process, and can be easily integrated into existing tools.

Our model assumes the production of multiple arrival acceptance rate scenarios with associated probabilities. The generation of the scenarios and their probabilities represents a unique research challenge. This requires fusion of weather forecasts with optimal runway configurations using a combination of statistical techniques and analysis of the impact of meteorological conditions on incoming aircraft. Moreover, the interaction of incoming aircraft with departing aircraft could lead to many interesting problems in optimization.

References

[1] "Aircraft Operating Costs by Stage of Flight and Associated Air Traffic Control Delay Costs Based on a Typical Flight," 1993, *Air Transportation Association*, Washington, 1995.

[2] Ahuja, R., T. Magnanti, and J. Orlin. 1993, *Network Flow: Theory, Algorithms and Applications,* Prentice Hall, Englewood Cliffs, New Jersey.

[3] Andreatta, G., A. Odoni, and O. Richetta, O. 1993, "Models for the Ground-Holding Problem," in *Large-Scale Computation and Information Processing in Air Traffic Control*, L. Bianco and A.R. Odoni, eds., Springer-Verlag, Berlin, 125-168.

[4] Bertsimas, D. and S. Stock. 1996, "Air Traffic Flow Management Problem with Enroute Capacities," to appear in *Operations Research.*

[5] Birge, J. and F. Louveaux. 1997, *Introduction to Stochastic Programming*, Springer-Verlag, New York, New York.

[6] Bixby, R.E. and W.H.Cunningham. 1980, "Converting Linear Programs to Network Flow Problems," *Math. of Operations Research*, August, 483-508.

[7] Glover, F., D. Klingman, and N. Phillips. 1992, *Network Models in Optimization and Their Applications in Practice*, John Wiley and Sons, Inc., New York, New York.

[8] Hoffman, R. 1997, "Integer Programming Models for Ground-Holding in Air Traffic Flow Management," Ph.D. Dissertation, University of Maryland.

[9] Nemhauser, G.L. and L.A.Wolsey. 1988, *Integer and Combinatorial Optimization*, John Wiley and Sons, Inc., New york, New York.

[10] Odoni, A.R. 1987, "The Flow Management Problem in Air Traffic Control," In *Flow Control of Congested Networks*, A.R. Odoni, L. Bianco, and G. Szego, eds., 269-288, Springer-Verlag, Berlin.

[11] Powell, Warren B., P. Jaillet and A. Odoni. 1995, "Stochastic and Dynamic Networks and Routing," *Handbook in Operations Research and Management Science*, **8**, M. O. Ball, et. al., editors, Elsevier Science, Berlin.

[12] Richetta, O. and A.R. Odoni 1993. "Solving Optimally the Static Ground-Holding Policy Problem in Air Traffic Control," *Transportation Science,* **27**, 228-238.

[13] Richetta, O. 1995, "Optimal Algorithms and a Remarkably Efficient Heuristic for the Ground-Holding Problem in Air Traffic Control," *Operations Research*, **43**, 758-770.

[14] Rifkin, R. 1998, "The Static Stochastic Ground Holding Problem," Masters Thesis, Massachusetts Institute of Technology.

[15] Terrab, M. and A.R. Odoni. 1993, "Strategic Flow Management for Air Traffic Control," *Operations Research*, **41**, 138-152.

[16] Tutte, W. T. 1965, "Lectures on Matroids", *Journal of Research of the National Bureau of Standards B*, **69B**, 1-47.

[17] Vranas, P., D. Bertsimas, and A. R. Odoni. 1994, "The Multi-Airport Ground-Holding Problem in Air Traffic Control," *Operations Research*, **42**, 249-261.

[18] Vranas, P., D. Bertsimas, and A. R. Odoni. 1994, "Dynamic Ground-Holding Policies for a Network of Airports," *Transportation Science,* **28**, 275-291.