Coordinating multiple traffic management initiatives with integer optimization

Andrew M. Churchill, David J. Lovell
Department of Civil and Environmental Engineering &
Institute for Systems Research
University of Maryland
College Park, Maryland U.S.A.
churchil@umd.edu, lovell@umd.edu

Michael O. Ball
R.H. Smith School of Business &
Institute for Systems Research
University of Maryland
College Park, MD U.S.A.
mball@rhsmith.umd.edu

Abstract–In this paper, the problem of coordinated resource rationing is considered in the face of current operational practice. In the United States, access to congested aviation resources is typically controlled by a system of capacity rationing wherein flights are assigned to slots at specific times. This is a well-accepted and efficient system, but it is not well-equipped to handle the problem faced when a single flight is controlled by more than one rationing initiative. The question of which, if any, initiative takes precedence over the others is not easily answered. An integer optimization model is introduced in this paper to find the delay-minimizing combination of multiple slot assignments for a set of flights and rationing initiatives. Rather than approach the problem comprehensively, this model treats each rationing initiative as somewhat independent, including only a constraint to guarantee that whatever slot pairs are assigned are mutually compatible. Computational results, including a case study, are reported along with directions for continuing research.

Keywords– ground delay program, airspace flow program, resource rationing, integer programming

1. INTRODUCTION

In air transportation, demand for certain resources occasionally meets or exceeds the available capacity. This situation may be particularly acute at certain airports for arrival and departure operations, and in the airspace along certain routes. At some of these resources, congestion may occur under even nominal conditions, but at others, weather is typically the exacerbating factor.

Several systems have been developed around the world to address these demand-capacity imbalances. The underlying principle in each of these systems is that demand-capacity imbalances should be addressed before the affected flights depart, as absorbing delay on the ground is safer and less expensive than doing so in the air. As a result, flights are typically given controlled departure or arrival times corresponding to some slots they have been assigned in the rationing initiative.

In the U.S., ground delay programs (GDP) and airspace flow programs (AFP) address demand-capacity imbalances expected at airports and in the airspace, respectively. These initiatives operate independently from one another and are employed sparingly – under nominal conditions they are typically not used. In Europe, the Central Flow Management Unit (CFMU) assigns control times comprehensively to address demand-capacity imbalances throughout each flight's route. This system is employed continuously, always adjusting flight control times.

The process of assigning ground delays to alleviate these demand-capacity imbalances has been well-studied scientifically, having been first systematically outlined in [*1*]. The single initiative case, similar in principle to the GDP and AFP used in the U.S., has been formulated for deterministic [*2*] and stochastic ([*3*], [*4*]) cases, as well as with static [*2*] and dynamic ([*4*]) decision making structures. Likewise, the problem has been expanded to consider a network of airports over which individual aircraft operate multiple subsequent flights ([*5*], [*6*]). In all of these cases, however, each flight has only been affected by a single rationing initiative because only airport arrival capacity was rationed, and a single flight may only arrive once. Until the summer of 2006, airspace capacity was not explicitly rationed [*7*]. At this time, AFP was introduced, employing the same principles and software to manage disruptions as are used for GDP [*8*].

As ground holding strategies have been implemented in the U.S., there has been a strong desire to include user (e.g., airlines, private jets) priorities. This led to the development of the Collaborative Decision Making (CDM) paradigm [*9*]. This community established that the most equitable means by which capacity is rationed is the published flight schedule, although other metrics have been considered [*10*]. Through the user input this process allowed, numerous enhancements have been implemented, including exempting certain flights, facilitating slot trades, and crediting flight operators for providing timely information about flight status.

At the other end of the complexity spectrum from single initiative models used in practice are the models that assign delays comprehensively, while considering any and all resources a flight encounters along its route. Models such as this ([*11*], [*12*], [*13*], [*14*]) are more consistent with the European approach to managing demand-capacity imbalances, but face difficulties due to the tremendous data requirements and complexity resulting from their comprehensive viewpoint.

An important distinction with these comprehensive models versus those described previously is that these must consider all flights simultaneously to build a plan, whereas the previous models consider only those flights explicitly affected by demand-capacity imbalances. While it is likely that such comprehensive approaches will minimize total delays, they have not been implemented in practice for a wide range of possible reasons. The single initiative models are more easily embedded into decision support system where human intervention is possible, especially to accont for dynamically changing conditions. Related to this is the challenge of integrating the global models within the CDM paradigm. Our goal with the research presented here is to develop models that lie somewhere in between so as to achieve a more global perspective while preserving the important CDM and practical decision support features.

We seek to model an application setting where some flights are affected by multiple rationing initiatives, but each of these initiatives functions, to a large degree, independently. This is a less complex problem than the comprehensive approach, since it does not require planning for flights not affected by constrained resources. Considering these rationing initiatives in the independent fashion as is done today, with an eye toward how they may better be coordinated is the objective of this paper. A numerical example of this conflicting situation is provided as a further introduction to this problem.

To address this problem area, several models are shown. First, a formulation for the single resource rationing problem is defined. Then, using that as a basis, a formulation is shown for rationing several connected resources to ensure feasible slot pair assignments. Finally, computational results on a realistic case study demonstrating both the applicability and feasibility of this coordinated formulation are shown.

2. Motivating Example

The models presented in this paper address the U.S. paradigm of rationing capacity independently at each resource. Specifically, the problem of coordinating the potentially conflicting times assigned by multiple independent rationing initiatives is examined. The impact of this problem is examined in the following example illustrated by the space-time diagram shown in .

In this example, a single flight is travelling from its origin airport to its destination and is passing through a stormy region of reduced capacity. Further, capacity at its destination is reduced. Thus, it is subject to rationing both en route and at its destination. The nominal interoperation times at each resource are two minutes, but under the degraded conditions, these increase to three minutes. Rationing is performed according to the accepted principle of using the schedule as a baseline.

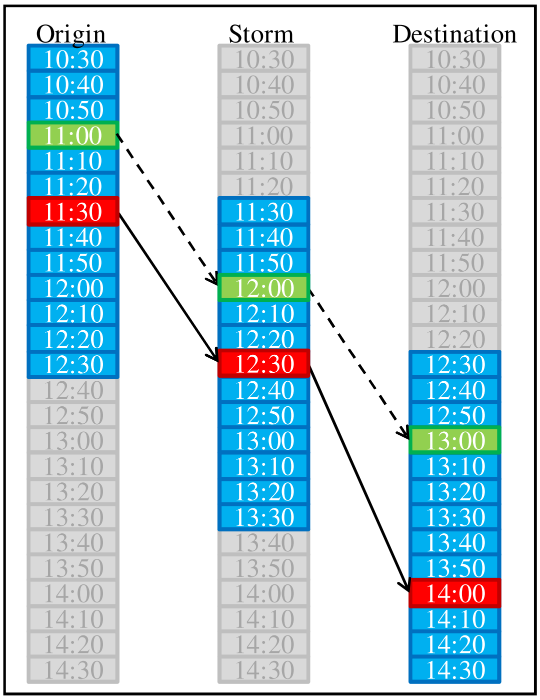

Figure 1 – Example of conflicting rationing

The total travel time is one hour, and the storm lies directly in the middle of the trip, as represented in the figure. The dashed lines represent the two hour scheduled trip departing at 11:00 that arrives at 13:00. The slot assignments after rationing are shown with the solid lines.

The first leg of the trip shown (origin-storm) does not present any problems with the assigned slots, as the flight can depart the origin at any time. However, using this slot at the storm will require a 90 minute travel time between the storm and destination. This is unacceptably more than the nominal travel time, resulting in an infeasible combination.

Clearly, however, this situation can be resolved by prioritizing one initiative over the other and exempting the flight through the secondary initiative by automatically granting it a slot other than the one it would have received. In the case of multiple flights affected by more than two different resources, the solution becomes much more difficult. The various considerations clearly yield many possible combinations of outcomes. Clearly, this is a complex combinatorial problem, and several approaches which leverage various properties of this problem are presented in the next section.

3. Single Initiative Model

Because the objective of this paper is to describe a model which coordinates multiple quasi-independent rationing initiatives, it is useful to first develop the model that represents the independent rationing process.

The broad objective of such a model is to ration capacity at a single congested resource both efficiently and equitably. The formulation shown here accomplishes this by assigning flights to slots in a delay-minimizing fashion. While an integer optimization approach is employed, the results can be mimicked through algorithmic means. The model shown here is similar in principle to many others but is included as an important transition to the multiple initiative model in the subsequent section.

One feature that does make this formulation unique from previous research is the paradigm used to describe the rationed capacity. In this model, the index of the slot to which a flight is assigned is decoupled from the time associated with that slot. This adds an additional qualification (discussed later) to some of the summation terms, but in so doing, helps to reduce formulation size. This stands in contrast to the construct used by most ground holding models of constant length time intervals, each with varying capacity.

Rationing models in other research have been formulated with a uniform lattice of time periods over which discrete capacity changes occur. Typically these periods encompass more than one flight. Because airport capacities are specified as hourly rates, this approximation presents problems, primarily that the number of different hourly rates that can be represented in each time period is limited by the length of that time period. The length of these discrete periods is on the order of 5-15 minutes. For example, using a lattice with 5 minute bins enables each bin to represent hourly capacities that are multiples of 12.

In the formulations shown in the paper, capacity divisibility is enabled by using slots of unit capacity. Slots are not necessarily assigned on a uniform lattice, thus allowing for non constant interoperation times. Thus, resource capacity is specified as a list of slots, rather than a list of time periods each with associated capacity.

Because the models shown in this paper address deterministic capacity, the primary advantage of this modeling paradigm is the capacity divisibility. However, for stochastic formulations that consider specific potential outcomes, this methodology may greatly simplify formulations and solutions. For example, in cases in which capacity may only increase over time in various scenarios, this paradigm greatly simplifies solutions, as flights will stay assigned to the same slot index, but the operation time associated with that index will decrease.

A. Formulation

The model described here is an assignment formulation with limitations placed on the set of slots to which a given flight may be assigned. Several input data are required to understand this formulation.

The set F comprises the individual flights affected by the rationing initiative, each with prespecified scheduled arrival time α_f. The set S comprises the slots to which those flights will be assigned, each with associated slot beginning time τ_s. In this work, the number of slots always equals or exceeds the number of slots, as shown in (1). In this work, a feasible solution is assumed to exist – this also implies the condition shown in (1). In building a case study, this is a trivial condition to enforce, as many additional slots with small interoperation times may be created after the planned initiative end time. This simulates the reality at most airports, at which operations may be extended late into the night to accept the day's flights.

$$|S| \geq |F| \qquad (1)$$

The decision variables x_{fs} are integer valued, assuming a value of one when flight f is assigned to slot s and zero in all other cases. Decision variables are only created for combinations of f and s for which the condition in (2) is met. This helps to reduce the size of the constraint matrix by eliminating unnecessary variables. In principle, others could be considered, but would be necessarily fixed to zero.

$$\tau_s \geq \alpha_f \qquad (2)$$

The first constraint set in this formulation is shown in (3). This enforces the condition that each flight must be assigned to exactly one slot. An additional condition is imposed that the slot to which each flight is assigned must begin at or after the flights scheduled arrival, as in (2).

$$\sum_{\substack{s\in S:\\ \tau_s \geq \alpha_f}} x_{fs} = 1 \qquad \forall f \in F \qquad (3)$$

The second constraint set, (4), enforces the capacity of each slot to be at most one flight. As discussed, the construct of using single-flight slots is also somewhat unique. Other models have assumed longer slot lengths to avoid the problem of capacity divisibility, but have then assigned multiple flights to each slot.

$$\sum_{f\in F} x_{fs} \leq 1 \qquad \forall s \in S \qquad (4)$$

The objective of this optimization problem is specified by the function shown in (5). This function minimizes the total sum of ground delays assigned to all flights. The superlinear function of delay length is used to favor the assignment of two short delays over a single long one. This principle contributes to equity between different flight operators because flights that are similar a priori are assigned similar delays.

$$\min z = \sum_{f\in F} \sum_{\substack{s\in S:\\ \tau_s \geq \alpha_f}} \left(\tau_s - \alpha_f\right)^{1+\varepsilon} x_{fs} \qquad (5)$$

B. Problem size

One measure of formulation strength and computational tractability is the size of the constraint matrix. The theoretical maximum/minimum numbers of constraints and variables is shown in . Several realistic numerical problem sizes are shown as well.

TABLE 1 – SINGLE INITIATIVE PROBLEM SIZE

	Constraints	Variables	
		Best case	*Worst case*
Nominal	$\lvert F\rvert+\lvert S\rvert$	$\frac{\lvert F\rvert}{2}\left(\lvert F\rvert+1\right)$	$\lvert F\rvert\lvert S\rvert$
Small $\lvert F\rvert=10$, $\lvert S\rvert=10$	20	55	100
Typical $\lvert F\rvert=100$, $\lvert S\rvert=150$	250	5050	15000

4. MULTIPLE INITIATIVE MODEL

Although the first model shown in this paper has some interesting properties, the reality is that it addresses a fairly well-solved problem. In particular, its solution is also attainable through the use of the ration-by-schedule algorithm, given several assumptions. Of greater interest, however, is the case in which multiple rationing initiatives assigning conflicting slot times to flights. This problem is addressed in this section.

The essentials of this integer programming formulation are similar to those shown in the previous section, with the exception of the addition of a single constraint set. This is, of course, the stated objective of this work – to develop a coordinated rationing method consistent with current practice. The added constraints enforce the logical condition that the slot times assigned to flight that use multiple initiatives are compatible. In some ways, this model is an extension and simplification of that proposed in [*15*] in that only those regions under adverse conditions are expressly controlled.

Thus, little additional input data are required. The set I comprises the initiatives that are to be rationed. As a result of this addition, each initiative has its own independent slot set that is indexed with i as S^i. In addition, the slot times and scheduled flight arrival times are each now indexed by i as well. If the number of initiatives is one, then this formulation simply reduces to that shown previously. Thus, this model may be seen as a generalization of the previous.

Several input data are required to account for the multiple initiatives for each flight. The set V_f is defined as all initiatives visited by flight f. For a flight to be included in this model, $\lvert V_f\rvert$ must be greater than zero. The value N_f^i is the initiative visited by flight f after initiative i. This is used to maintain the ordering of initiatives.

A. Formulation

The decision variables x_{fs}^i are integer valued, assuming a value of one when flight f is assigned to slot s in initiative i and zero in all other cases. Decision variables are only created for combinations of f, s, and i for which the condition in (6) is met. Similar to condition (2), this helps to reduce the size of the constraint matrix by eliminating unnecessary variables.

$$\tau_s^i \geq \alpha_f^i \tag{6}$$

The first constraint set in the linked formulation is shown in (7). It is similar to that shown in (3), with the added dimension of each initiative. This enforces the condition that each flight must be assigned to exactly one slot in each rationing initiative.

$$\sum_{\substack{s\in S^i:\\ \tau_s^i\geq\alpha_f^i}} x_{fs}^i = 1 \qquad \forall f\in F, i\in I : i\in V_f \tag{7}$$

Likewise, constraint set (8) is similar to (4) in that it enforces the condition that each slot in each initiative may have at most one flight assigned to use it.

$$\sum_{\substack{f\in F:\\ i\in V_f}} x_{fs}^i \leq 1 \qquad \forall i\in I, s\in S^i \tag{8}$$

The constraint set that links together these multiple initiatives is shown in (9). This constraint set works by defining feasible slot combinations in each pair of initiatives for each flight. The range R_{fs}^{ij} defines the times that flight f could feasibly arrive at initiative j by using slot s in initiative i. For each flight f, some maximum delay m_f between initiatives is defined. If only a single initiative is considered, then the value N_f^i is always empty, and none of this constraint set is present in the formulation.

$$x_{fs}^i - \sum_{\substack{t\in S^j:\\ \tau_t^j\in R_{fs}^{ij}}} x_{ft}^j \leq 0 \qquad \begin{aligned}&\forall f\in F, i\in V_f,\\ &j=N_f^i, s\in S^i:\\ &\tau_s^i\geq\alpha_f^i, N_f^i\neq\varnothing\end{aligned} \tag{9}$$

The maximum delay value should be small, as it defines the period over which the operator is indifferent to various slots. That is, the unit cost of these first few minutes of “airborne” delay is equal to the unit cost of ground delay. The maximum delay in this context is meant only to permit a bit of slack for the variations in spacing of slot times at each initiative.

The concept of constraint set (9) is illustrated in . Assume that there are two initiatives, and the nominal travel time between them is 60 minutes. A flight may travel somewhat slower than nominal between the two, up to ten extra minutes. Under this scenario, the feasible slots to use in the second initiative are shown for the 10:32 slot in the first.

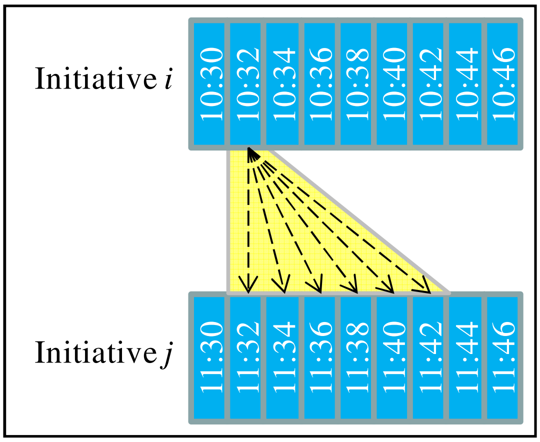

Figure 2 – Feasible range example

The range R_{fs}^{ij} used in (9) is defined in (10) as beginning at the sum of the time for slot *s* at initiative *i* and the inter-initiative travel time. The range ends after the maximum inter-initiative delay.

$$R_{fs}^{ij} = \left[\tau_s^i + \alpha_f^j - \alpha_f^i, \tau_s^i + \alpha_f^j - \alpha_f^i + m_f \right] \qquad (10)$$

The objective of this formulation, as with most of its ilk, is to minimize delays. It is very important, however, to consider precisely which delay is being minimized. Two alternate objectives are presented here, with results for each shown in the following section.

The difference between the two possible objectives is the scope of the delay summed. The first considers delays at each initiative independently while the second considers only those at a flight's final initiative. This first potential objective, considering the total amount of delay assigned, is shown in (11). Thus, some delays may be "double counted" according to (11) because they are counted twice but truly impact the flight only upon arrival to its destination.

$$\min z = \sum_{f \in F} \sum_{i \in I} \sum_{\substack{s \in S: \\ \tau_s^i \ge \alpha_f^i}} \left(\tau_s^i - \alpha_f^i \right)^{1+\varepsilon} x_{fs}^i \qquad (11)$$

The second objective, considering only the delay at the flight's final initiative, is shown in (12).

$$\min z = \sum_{f \in F} \sum_{i \in I} \sum_{\substack{s \in S: \\ \tau_s^i \ge \alpha_f^i \\ N_f^i = \varnothing}} \left(\tau_s^i - \alpha_f^i \right)^{1+\varepsilon} x_{fs}^i \qquad (12)$$

Because this model treats capacity deterministically and sufficient capacity at the origin airport is assumed, all delays are taken on the ground before departure. Thus, there is no need to consider a cost differential. A comprehensive model that schedules all points along a flights route, or one that considers capacity stochastically, however, would fail for this assumption.

B. *Formulation size*

Again, the size of the constraint matrix is considered. The theoretical worst case numbers of constraints and variables and several realistic numerical problems are shown in .

TABLE 2 – LINKED INITIATIVES PROBLEM SIZE

	Constraints	Variables
Nominal	$\sum_{f \in F} \lvert v_f \rvert + \sum_{i \in I} \lvert S^i \rvert + \sum_{f \in F} \lvert v_f \rvert * \sum_{i \in I} \lvert S^i \rvert$	$\lvert F \rvert \sum_{i \in I} \lvert S^i \rvert$
Small $\lvert F \rvert = 10, \lvert I \rvert = 2,$ $\lvert S^i \rvert = 10$	440	200
Typical $\lvert F \rvert = 100, \lvert I \rvert = 3$ $\lvert S \rvert = 150$	135750	45000

5. COMPUTATIONAL RESULTS

To demonstrate the efficacy of this linked rationing initiative formulation, a realistic numerical case study has been developed. The schedule data are randomly generated, but represent a realistic situation such as is encountered during summer convective weather over the northeastern United States.

In this case study, there are two airports (B, C) for which capacity is being rationed, and one disruption (A) in the enroute airspace for which rationing must take place, as shown in .

To better visualize the problem setup, the aircraft are nominally grouped into flows, as labeled in this figure. Flow 1 comprises flights crossing the enroute disruption, but not travelling to either of the two disrupted airports. Flows 3 and 5 travel to airports B and C, respectively, but do not cross the enroute disruption. The most interesting

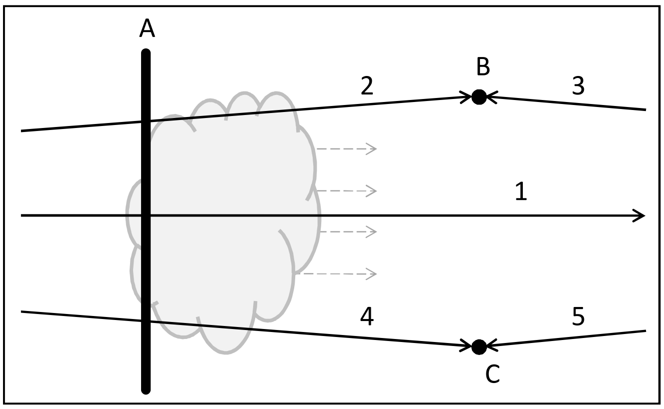

Figure 3 – Case study layout

flows, 2 and 4, cross the disrupted airspace before arriving at the disrupted airports (B and C, respectively). It is these two groups of flights that confound the traditional single resource ration-by-schedule methods employed. Of course in the model, the individual flights themselves are considered.

Nominally, each resource has a capacity of 60 flights/hour (or an interoperation time of 60 seconds), but under the reduced conditions in this case study will have half of that capacity available for a four hour period. The number of scheduled arrivals to each resource over the study period is shown in Figure 4. Because the travel time between the storm and each airport is one hour, the bars representing the schedule for flows 2 and 4 are simply shifted by one hour from their appearance in the resource A schedule to their appearances in the resource B and C schedules, respectively.

The schedule is assumed to terminate after the flights shown in Figure 4. While potentially unrealistic, this simplifies considerably the conditions surrounding the end of the program because the flights expected to arrive after the end of the program are not subject to rationing.

Because the capacity reduction is sufficiently extreme relative to the scheduled number of aircraft, the optimization model will assign flights to every slot. Thus, the time-varying profile of flights after the model has run will match precisely with the reduced capacity line until the entirety of the set of flights has been assigned.

This case study resulted in a constraint matrix with 24540 rows (constraints) and 189992 columns (variables). The model was run on a quad processor system with 16GB memory using the Xpress 2008A solver. Two objectives were considered. Using (11), the model solved to integer optimality with the linear programming (LP) relaxation in 38 seconds. With (12), the LP relaxation was not integer and so branch and bound was employed. After 60 minutes, the model had a 13.0% optimality gap, while after 240 minutes, this had narrowed to 2.5%. The obvious differences resulting from these two objectives warrant further investigation.

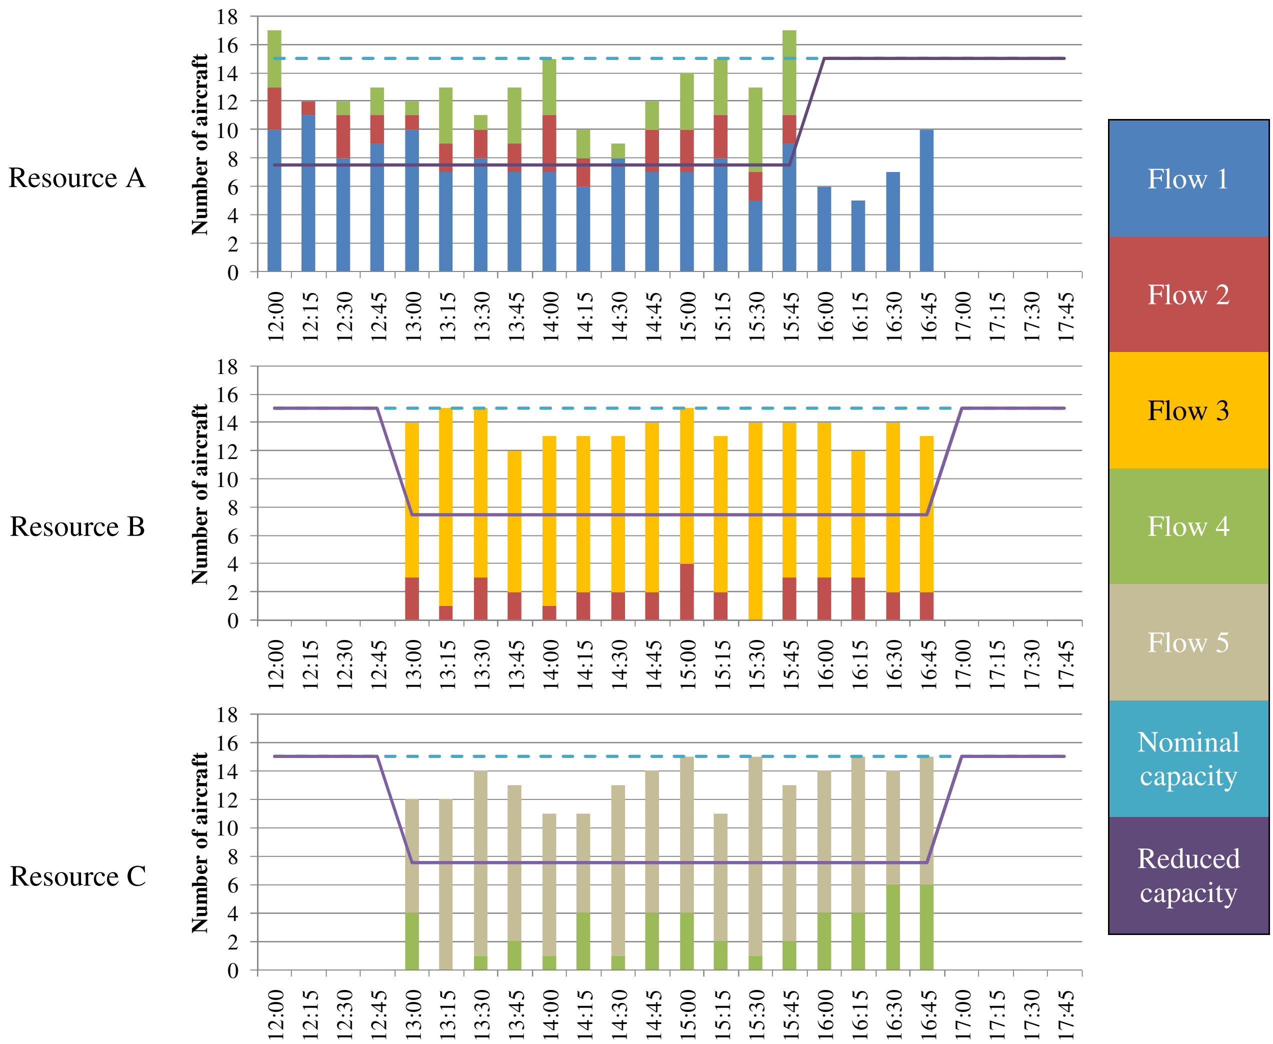

Figure 4 – Nominal resource schedules

The results for the two objectives and one heuristic approach are shown in Table 3. These results will be discussed in response to three issues regarding delay distribution:

1. Within the total delay optimal results, but between the five flows: Does the model favor one flow or type of flow?
2. Between total delay and final initiative optimal results: How do these objectives affect each flow?
3. Between the total delay optimal results and a priority scheme mimicking current practice: How is the optimal solution different from the heuristic one?

A. Issue 1: Fairness between flows

Using the total delay objective, the first issue examined is whether the model favors one flow or type of flow over another. The concern in this case is that flights using multiple initiatives may be unduly prioritized or penalized. A *t*-test at 5% significance suggests that the average delay assigned to each flow is not significantly different (at a 5% level) than the overall average delay for flows 1-4. The null hypothesis of average delay equal to overall average delay cannot be rejected for flow 5 however. There is no obvious structural explanation for this, so further investigation with other case studies is warranted.

Table 3 – Summary of average delays (in minutes) assigned to each flow (standard deviations shown in parentheses)

Flow	**Flight count**	**Total delay optimal**	**Final init. optimal (2.5% gap)**	**Heuristic**
1	155	66.9 (26.9)	55.7 (26.9)	64.7 (25.5)
2	70	74.3 (31.3)	194.1 (62.4)	78.3 (33.5)
3	183	76.5 (34.1)	43.0 (22.5)	76.5 (33.9)
4	92	74.9 (26.8)	161.1 (78.7)	74.6 (28.2)
5	166	65.3 (31.9)	13.4 (7.3)	65.9 (31.8)

B. Issue 2: Differences between objectives

To examine the differences between the two objective functions considered, the average delays assigned to each flow are compared pair-wise. Clearly these results are different, particularly for the delays assigned to the flights in flows 2 and 4 – those that use two resources. A two sample *t*-test at 5% confirms that each of the average delays assigned to these pairs of flows are significantly different. Although it is clear that the solutions are different, the reasoning behind these differences is not immediately clear and further investigation is warranted on this issue, particularly in light of the differences in solution times and the potential utility of the second objective function.

C. Issue 3: Comparison to heuristic

To examine the utility of the optimization approach, a heuristic solution is also considered. This heuristic mimics current practice in that the schedule is used as the basis for rationing and conflicts between initiatives are resolved by prioritizing the airport initiative. Thus, flights using more than one resource are granted an exemption through their non-airport resource.

Because, for this case study, the total amount of delay assigned by the optimization model is the same as that assigned by the heuristic, comparisons in this case are made between each pair of flows. Comparing the heuristic and the total delay objective with a two sample *t*-test at 5% significance, the null hypothesis of equal means cannot be rejected for any matched pair of flows. This suggests that the distribution of delays across flows assigned by the optimization model is not statistically different from that assigned by the heuristic.

While this result is troubling for the application of the optimization model, it is actually quite useful. Because the optimization models may take considerable time to solve depending on the objective and other considerations used, it is useful to have heuristics that can ably mimic their results. Thus, in a laboratory environment, the optimization models can be used as the standard by which heuristics may be judged. The simple heuristic described here seems to perform well according to these criteria, but further investigation on this, and other related heuristics, is required to make a decisive determination on these issues.

6. CONCLUSIONS

In this paper, a model that coordinates multiple independent resource rationing initiatives employed in air transportation was shown. The problem that this model solves lies between those currently considered in the literature, both with respect to complexity and scope. An optimization-based model was defined, its computational properties explored, and a realistic case study outlining its utility shown.

The results of the case study suggest that further exploration is needed to understand the different possible objectives employed in optimizing the distribution of delay. Clearly the solution resulting from the second objective would be unacceptable in practice because of the tremendous inequity placed upon a subset of flights. The initial results of the case study, however, suggest that it is possible to develop a heuristic approach that mimics the solutions derived from the optimization model.

Several enhancements are possible to extend both the realism and utility of this formulation. The first such extension is the inclusion of additional input data and slightly modified constraints to allow for the consideration of flight exemptions, both for flights already en route, or for those exempted for other reasons. In addition, more sophisticated methods for examining problem feasibility may be included with the addition of a single slot at a time much later than the last scheduled flight to capture flights that cannot otherwise find a feasible assignment.

In addition, the nature of the two different objective functions must be explored, aiming toward explaining what structural difference resulted in such a wide variance in solution times and delay distribution. Further, it may be interesting to explore other objective functions that make the trade between equity and efficiency differently from that shown here.

Finally, considerable attention will be paid to developing a formulation of this model that incorporates stochastic recourse. Of course this is a considerably more complex problem that departs from practice, but there is previous research on including stochastic information in similar models has shown that great potential for insight exists.

ACKNOWLEDGMENT

The authors gratefully acknowledge the support of the National Aeronautics and Space Administration Airspace Systems Program under ARMD NRA: NNH06ZNH001 and Metron Aviation for assistance with data.

REFERENCES

[1] Odoni, A.R., "The flow management problem in air traffic control," in *Flow Control of Congested Networks*, Amedeo R Odoni and G Szego, Eds. Berlin: Springer-Verlag, 1987.

[2] Terrab, M. and A.R. Odoni, "Strategic flow management for air traffic control," *Operations Research*, vol. 41, no. 1, pp. 138-152, 1993.

[3] Ball, M.O., R.L. Hoffman, A.R. Odoni, and R. Rifkin, "A stochastic integer program with dual network structure and its application to the ground-holding problem," *Operations Research*, vol. 51, no. 1, pp. 167-171, 2003.

[4] Mukherjee, A. and M. Hansen, "A dynamic stochastic model for the single airport ground holding problem," *Transportation Science*, vol. 41, no. 4, pp. 444-456, 2007.

[5] Vranas, P.B., D.J. Bertsimas, and A.R. Odoni, "The multi-airport ground-holding problem in air traffic control," *Operations Research*, vol. 42, pp. 249-261, 1994.

[6] Vranas, P.B., D.J. Bertsimas, and A.R. Odoni, "Dynamic ground-holding policies for a network of airports," *Transportation Science*, vol. 28, no. 4, pp. 275-291, 1994.

[7] Brennan, M., "Airspace flow programs - a fast path to deployment," *Journal of Air Traffic Control*, vol. 49, no. 1, pp. 51-55, 2007.

[8] Krozel, J., R. Jakobovits, and S. Penny, "An algorithmic approach for airspace flow programs," *Air Traffic Control Quarterly*, vol. 14, no. 3, pp. 203-230, 2006.

[9] Wambsganss, M., "Collaborative decision making through dynamic information transfer," *Air Traffic Control Quarterly*, vol. 4, no. 2, pp. 109-125, 1997.

[10] Ball, M.O., R.L. Hoffman, and A. Mukherjee, "Ground delay program planning under uncertainty based on the ration-by-distance principle," *Transportation Science*, 2009.

[11] Bertsimas, D.J. and S. Stock Patterson, "The air traffic flow management problem with enroute capacities," *Operations Research*, vol. 46, no. 3, pp. 406-422, 1998.

[12] Bertsimas, D.J. and S. Stock Patterson, "The traffic flow management rerouting problem in air traffic control: a dynamic network flow approach," *Transportation Science*, vol. 34, no. 3, pp. 239-255, 2000.

[13] Lulli, G. and A.R. Odoni, "The european air traffic flow management problem," *Transportation Science*, vol. 41, no. 4, pp. 431-443, 2007.

[14] Bertsimas, D.J., G. Lulli, and A.R. Odoni, "The air traffic flow management problem: an integer optimization approach," in *13th International Conference, IPCO 2008*, Bertinoro, Italy, 2008, pp. 34-46.

[15] Churchill, A.M., D.J. Lovell, and M.O. Ball, "Evaluating a new formulation for large-scale traffic flow management," in *Proceedings of the 8th USA/Europe Air Traffic Management R&D Seminar*, Napa, California, 2009.