A GENERAL APPROACH TO EQUITY IN TRAFFIC FLOW MANAGEMENT AND ITS APPLICATION TO MITIGATING EXEMPTION BIAS IN GROUND DELAY PROGRAMS

Thomas Vossen, University of Colorado, Boulder, CO

Michael Ball, University of Maryland, College Park, MD

Robert Hoffman, Michael Wambsganss, Metron Aviation, Herndon, VA

Abstract

A primary objective of the FAA's ATM functions is to provide fair and equitable access to the National Air Space. Traditionally, the FAA has interpreted fairness as prioritizing flights on a "first-come, first-served" basis. The allocation procedures introduced under Collaborative Decision Making (CDM), however, represent a departure from this paradigm: allocations are based on carriers' original flight schedules. Yet in spite of these changes, the concept of fairness under CDM is largely left implicit in the procedures. Different and even conflicting concepts are sometimes used to describe these procedures. Moreover, the achievement of equitable allocations is often complicated by practical considerations. This paper describes a general framework for equitable allocation procedures within the context of ATM, and illustrates its use in reducing certain systematic biases that exist under current procedures. We also discuss other applications of this approach, and summarize practical considerations and implementation issues.

1. Introduction

It has become increasingly apparent that demand for National Airspace System (NAS) resources has become very close to NAS capacity. Although measures are constantly being taken to increase capacity and, in spite of the recent demand decrease resulting from the events of September 11, it seems clear that demand and capacity will continue to remain very close into the foreseeable future.

The task of monitoring traffic demand-capacity imbalances in the U.S. falls on the FAA's air traffic flow management (TFM) specialists. Potential instances of excessive airborne delay are anticipated by TFM and met with various initiatives such as miles-in-trail or ground delay.

Each demand-capacity imbalance places TFM in a decision-making environment in which traffic flow specialists must allocate scarce resources. Which aircraft should receive priority? What is the basis for equity? How can the decision maker be certain that resources are being used wisely? TFM would like to make efficient use of resources while, at the same time, honoring its commitment to equitable treatment of NAS Users.

Currently, the most sophisticated resource allocation mechanisms are found in the context of ground delay programs (GDPs). A GDP is a traffic flow management initiative for addressing airport arrival capacity shortfalls: delays are applied to flights at their origin airports when they are bound for a common destination airport with reduced capacity or excessive demand. For allocation purposes, the time horizon of reduced capacity is divided into contiguous time intervals known as arrival slots.

The Collaborative Decision Making (CDM) program has established a highly successful paradigm for allocation of airport arrival slots. There are many subtleties associated with arrival slot allocation under CDM, but the essence of the allocation principle is "first-scheduled, first-served", meaning that the earlier arrival slots are generally awarded to the flights that are scheduled to arrive earlier. See [1], [2] for details on rationing.

The CDM experience has shown that not only is equitable treatment of carriers advisable, but it may be a necessary condition for efficient use of resources (see [1], [3]). Prior to CDM, effective GDP initiatives were based on dated flight data that did not reflect the airline's intentions upon the day of operation. This hindered accurate flow control and led to inefficiencies. It was found that the airlines felt that they were not being treated equitably and that the information they provided could be used to provide benefits to their competitors. As a result, they were

reluctant to provide up-to-date information. By instituting resource allocation methods that were based on an agreed upon standard (first-scheduled, first-served) and by allowing the airlines to derive benefits for the provision of cancellation information (the compression algorithm), the airlines were induced to provide up-to-date intent information through the CDM-Net. In addition, through regular meetings of the CDM working groups, a cooperative culture developed and joint problem solving was used to address open issues.

1.1 Exemptions within GDPs

Numerous program parameters must be chosen by TFM before a GDP can be implemented. Chief among these are the temporal scope of the program (start and end times), and the geographical scope. The temporal scope is usually determined by the start and end times of the degraded airport conditions.

Having set the temporal scope of the program, the amount of total delay that needs to be applied in order to smooth demand surges to capacity levels is easily computed. The crucial observation is that the total necessary delay that must be absorbed is virtually constant, and independent of which carriers receive the delay.

From an equity standpoint, it would be most desirable to spread this total necessary delay out over as many flights, or types of flights, as possible. But important, valid considerations motivate both limiting the scope of programs and also exempting certain flights within the scope from FAA-assigned delay.

There are two major categories of exemptions. First, flights that are already in the air when the program is filed cannot be assigned ground delay and, by necessity, are exempt. A second type of exemption is much more at the discretion of the TFM specialist designing and implementing the GDP: geographical exemption (see [4]). Clearly, FAA-assigned ground delay at the origin airport must be served before a flight departs, which in the case of coast-to-coast flights, can be 5 or 6 hours in advance of arrival. TFM is not always confident enough in a weather forecast to assign irrevocable ground delay that many hours in advance. For if the weather clears earlier than expected or, worse yet, does not materialize at all, then many of the flights will have endured (in hindsight) unnecessary ground delay.

One way that TFM mitigates the effects of capacity uncertainty is to limit the geographical scope of the program. Generally, flights originating at distant airports are exempted from FAA-assigned ground delay in a GDP. The price paid for exempting large numbers of flights is that the total necessary delay in a GDP is distributed over a smaller collection of flights, thus driving up the maximum delay and average delay per (non-exempted) flight. Thus, efficiency and equity are at odds with each other.

We take as a given in this paper the need to restrict the geographical scope of a GDP and that TFM exercises this option at their discretion on a program-by-program basis (see [5], [6], [7] for a treatment of stochastic planning issues in GDPs). Instead, our interest lies in the ramifications of the exemptions. If programs tend to be implemented with tight geographical scope, then carriers with predominantly shorter stage lengths may be receiving an undue amount of delay. Section 3 addresses possible systematic biases against such carriers. This assumes, of course, that an ideal allocation can be agreed upon as the standard of fairness, which is addressed in Section 2. Section 2 also proposes a general approach to minimizing the deviation between an actual allocation and an ideal standard. In Section 4, we apply the method given in Section 2, to offset the inequities imposed by exemptions. Lastly, in Sections 5 and 6, we discuss implementation issues and possible extensions of the approach to more general settings.

2. A General Approach to Equitable Resource Allocation

In approaching the problem of how to equitably allocate a resource, two basic problems usually arise. The first involves determining a standard or policy that constitutes an ideal fair allocation. This can be challenging, as there are usually multiple possible alternatives and the ultimate criterion for a good standard is that it be accepted by the parties involved. Typically, equity is not the sole criterion by which an allocation is judged. The second problem is to construct the resource allocation in such a way that takes into account the ideal allocation but also considers other objectives and constraints. We address each of these problems in the following subsections.

2.1 Defining an Ideal Allocation

The general acceptance of the CDM GDP procedures leads one to believe that the CDM resource allocation algorithms must be based on sound principles. To understand the fundamental

nature of CDM resource allocation we consider the "unconstrained" version of the CDM ration-by-schedule (RBS) algorithm, which is based on the first-scheduled, first served principle. Here by unconstrained, we mean that no flights are exempt from FAA-assigned ground delay. Recent research (see [2],[8]) has shown unconstrained RBS closely related to well-established equitable allocation concepts. Specifically, it is a strict priority method, where a flight's priority is based on its OAG time, i.e. a flight with an earlier OAG time has priority over a flight with a later OAG time. Furthermore, it has other important properties. The result of RBS is to assign to each flight, f, a controlled time of arrival, $CTA(f)$. This translates into assigning a delay, $d(f)$, to flight, f, which is given by $d(f) = CTA(f) - OAG(f)$, where $OAG(f)$ is the scheduled arrival time for f. All time values are rounded to the nearest minute under RBS, hence, each delay value $d(f)$ is integer. If we let D equal the maximum delay assigned to any flight and let $a_i = |\{f : d(f) = i\}|$ for $i = 0, 1, 2 \ldots, D$, then the important RBS properties can be defined by,

Property 1: RBS minimizes total delay $= \Sigma_f\, d(f)$.

Property 2: RBS lexicographically minimizes $(a_D, \ldots, a_1, a_0)$, i.e. a_D is minimized; subject to a_D being fixed at its min value, a_{D-1} is minimized; subject to (a_D, a_{D-1}) being fixed at its lexicographic min value, a_{D-2} is minimized, etc.

Property 3: For any flight f, the only way to decrease a delay value, $d(f)$, set by RBS is to increase the delay value of another flight g to a value greater than $d(f)$.

Property 3, which follows directly from Property 2, expresses a very fundamental notion of equity that has been applied in a number of contexts (see [9]). In some sense, it is remarkable that the procedures developed on very practical war-gaming and consensus building exercises have very well-studied properties. On the other hand, this result may not be surprising in that these properties represent the basis for consensus being reached.

We should note however, that the equity notions described in Properties 2 and 3 effectively treat flights as independent entities. However, the majority of flights are associated with particular airlines. The implicit definition of an airline priority is that the airline's priority is the "sum" of the priorities of all of its flights. One might ask if alternate solutions should be considered that are more reflective of an "airline-centric" point of view.

One limitation of the RBS-based approach to equity is *there is an implicit assumption that all flights are activated*. In fact, for many GDPs where airport capacity is degraded for several hours, airlines typically cancel large numbers of flights. When delays for a flight become very large, e.g. greater than two hours, it may be impractical or uneconomical to operate the flight. An airline with slots toward the end of a long, severe GDP would receive very few "usable" slots.

A possible alternative to RBS is proportional allocation. In general (see [8], [9]), a group of claimants each have a claim of a certain size. The resource in question is divided up in proportion to the size of each claim. Implicit in this approach is that the size of the resource to be allocated is smaller than the sum of the sizes of the claims (otherwise it is possible to give all claimants an allocation exactly equal to their respective claims). The most obvious application in the GDP context would be one where the claimants are airlines and the size of each claim is the number of flights scheduled during the GDP. There are some problems associated with the direct application of this idea. In particular, some slots might not be usable by an airline even though it has an implicit claim to these slots. For instance, an airline could not use a slot in the first hour of a GDP if its earliest flight cannot arrive until the second hour of the GDP. We have developed procedures that address the problem just described, and can be viewed as a compromise between RBS and a "pure" proportional approach (see [8]). These procedures will be discussed and analyzed in future papers.

In the primary analyses described in the rest of this paper we will use "unconstrained" RBS as an ideal allocation. We will also report on some experiments related to the use of a proportional allocation.

2.2 Minimizing the Deviation between the Actual and Ideal Allocations

As mentioned above, it is rare in a practical setting that the ideal allocation is feasible and/or even desirable. Taking the GDP example, the actual arrival slot allocation can be far different from the "ideal" RBS allocation because of the sometimes-extensive set of exempt flights. There are other incidents that lead to deviations from the ideal. We note two significant cases below.

A small number of flights typically must be inserted into the arrival stream after the GDP has been planned. These so-called *popup flights,* by occupying space in the arrival sequence and possibly causing unanticipated airborne delays to other flights,

are implicitly assigned arrival slots by a process totally independent from the original RBS allocation.

Occasionally an airline will be assigned a slot early in the GDP that it cannot use because its flights have earliest arrival times later than the slot. These incidents typically occur due to flight cancellations or when flights are delayed for reasons not having to do with the GDP allocation process. Such slots are currently reassigned by the *compression algorithm.*

The primary purpose of this paper is to propose a method for mitigating inequities that arise due to exemptions. We also discuss, both here and in later sections, how the approach represents a general approach to addressing problems of equitable allocation not only for GDPs but also for other air traffic flow management problems.

The roots of the general approach we employ lie in problems arising in balanced just-in-time (JIT) scheduling (see for example [10], [11]). In this application a factory must produce multiple products. Associated with each product is an ideal rate of production. Of course, there are typically many constraints regarding the manner in which production can occur within a factory. The models developed for this class of problems output production schedules that minimize the "deviation" between the actual production rates for each product and the ideal rates. A key issue in any such model is how deviation is defined. Absolute value and squared deviation measures are often used. For one general class of models, deviation is measured as a function of the time between when products are produced and the ideal production time. Another class measures deviation as a function of the number of products produced by a certain time and the ideal such number. In the remainder of this paper, we use the fist class of measures.

For the GDP problem, we define:

n = the number of slots to be allocated;

m = the number of airlines;

b_a = the number of flights associated with airline a.

Each slot j has a slot time s_j with $s_1 < s_2 < ... < s_n$. Each airline must be allocated b_a slots, denoted by $j'(a,1), j'(a,2), ...,j'(a,b_a)$. This ordering is consistent with the slot time ordering so that $s_{j'(a,1)} < s_{j'(a,2)} < ... < s_{j'(a,ba)}$. We call $j'(a,k)$ the k^{th} slot allocated to airline a. When developing an analogy with the balanced JIT production scheduling problem, airlines correspond to products. Rather than defining an ideal "rate" of slots each airline should receive, we define an ideal allocation over time based on RBS. That is,

an "unconstrained" version of RBS is executed, which produces an allocation of slots to flights. None of the issues given earlier in the section (i.e. popups, cancellations, and exemptions) are taken into account when this version of RBS is executed. The union of the slots allocated to the flights of each airline under RBS is the ideal allocation for that airline. The result of this step can be characterized by a 0/1 vector, p, defined by:

$p_{jak} = 1$ if j is the ideal value for $j'(a,k)$; 0 otherwise.

We wish to define an optimization model that produces an allocation with minimum deviation from this ideal allocation. Depending on the precise circumstances a variety of constraints could be imposed on that allocation. Specifically, we define:

$e(a,k)$ = the earliest slot that can be designated as the k^{th} slot for airline a.

$l(a,k)$ = the latest slot that can be designated as the k^{th} slot for airline a.

$FA(a)$ = the set of slots whose allocation is fixed a-priori to airline a.

The model variables are:

$x_{jak} = 1$ if j is the k^{th} slot allocated to airline a; 0 otherwise.

We can now formulate an assignment model that assigns slots to airlines. On the lefthand/supply side of the model are nodes that correspond to the available slots. On the righthand/demand side of the model are nodes representing (k,a) pairs (k^{th} slot for airline a). The model minimizes the squared deviation between the slots allocated and their ideal locations subject to constraints on fixed slots and earliest and latest slot time.

Min Σ_{jak} $(j-p_{jak})^2$ x_{jak}

s.t.

Σ_{ak} x_{jak} $\leq$ 1   for all slots j where

        $j \notin FA(a)$ for some a,   (1)

Σ_k x_{jak} $=$ 1 for all a and $j \in FA(a)$,   (2)

$\Sigma_{j:\ e(a,k)\ ≤j\ ≤l(a,j)}$ x_{jak} $=$ 1 for all a and k,   (3)

$x_{jak} \in \{0,1\}$ for all a, k and j.

Constraints (1) and (2) insure that each slot is assigned to at most one (a,k) pair and also that each fixed slot is allocated to the appropriate airline. Constraint (3) insures that each (a,k) pair is assigned a slot between its earliest and latest available times.
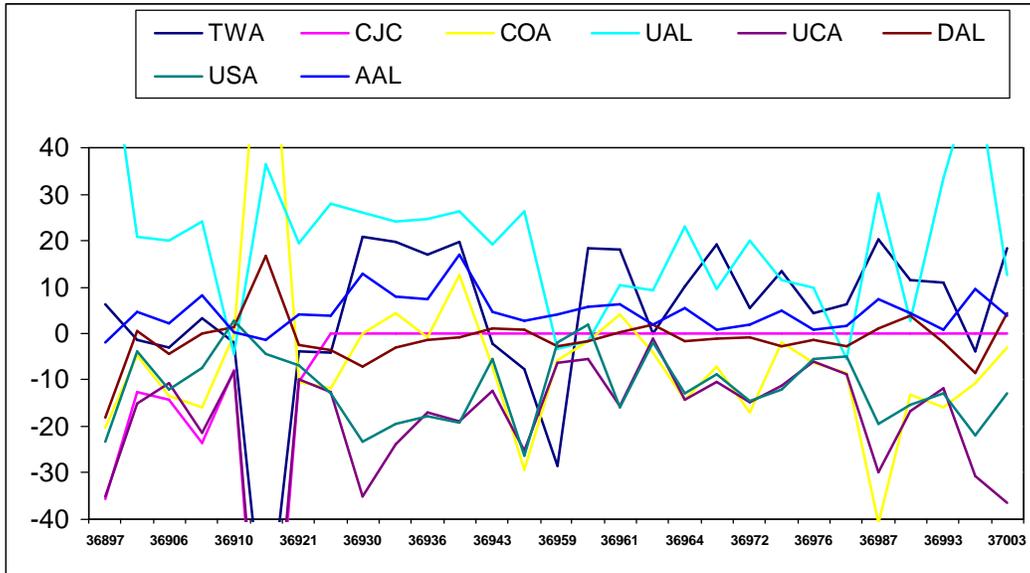
4

Figure 1: Historical Allocation, Logan Airport Boston

We have constructed alternative models and algorithms as well [8]. In particular, a more complex flow model can be used in certain cases to minimize alternative deviation measures. Also, in some cases it is possible to solve the problem directly using a simple "greedy" algorithm. These will be described in another paper.

This is a relatively simple model but it has very general applicability. In particular, a variety of applications are possible depending on how $e(a,k)$, $l(a,k)$ and $FA(a)$ are defined.

3. GDP Inequities Arising from Exemptions

Currently, exempted flights are assigned slots first, followed by the allocation of the remaining slots to the non-exempted flights. The manner in which flight exemptions are incorporated, however, can have a significant impact on the distribution of delays among airlines. To illustrate this, we analyzed the impact of flight exemptions using historical data.

For eight GDPs at Boston's Logan Airport during the first 4 months of 2001, we determined the delays for each airline with and without the exemptions that occurred during that day. The results are shown in Figure 1, which depicts for a selected number of airlines, the difference between an airline's average delay under the unconstrained RBS allocation and under RBS with exemptions included (a negative number on the vertical axis means the

airline would have received less delay if exemptions were not made). Each demarcation point on the horizontal axis corresponds to a date on which a GDP was executed at Logan.

Clearly, flight exemptions have a significant impact on the distribution of delays. Moreover, we found that exemptions may introduce a systematic bias in favor of or against certain airlines. For example, for one major carrier (USA) the use of flight exemptions increased its delay by an average of 11.7 minutes per flight (over its delay under the unconstrained RBS allocation), while for a smaller carrier (UCA) the increase in delay was 18.2 minutes per flight.

4. Using the Model to Mitigate Exemption Bias

From an airline-centric point of view, the current procedures to deal with flight exemptions can introduce significant inequities. Here, we analyze the extent to which the optimization model defined in Section 2 can mitigate these biases. In addition, we consider how the optimization model impacts the distribution of delays within an airline, that is, the individual flight delays for each airline.

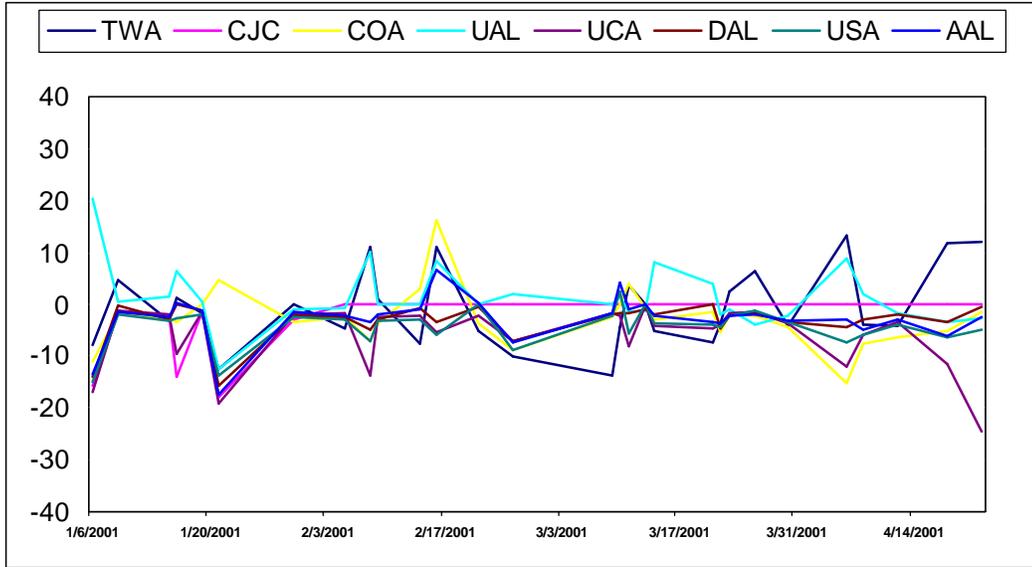We first must define how to apply the optimization model to this setting.

Figure 2: Optimization Model Impact, Logan Airport Boston

In addition to basic information on slots and flights, three key data sets must be defined: $e(a,k)$, $l(a,k)$ and $FA(a)$. The earliest arrival time for the k^{th} flight of airline a, $e(a,k)$, is set equal to the k^{th} earliest OAG arrival time amongst all of airline a's flights. The latest arrival time for the k^{th} flight of airline a, $l(a,k)$, is not defined, i.e. it is assumed to be infinite. Finally, the key dataset for this application is the set of fixed slots for airline a, $FA(a)$. This is set equal to the set of slots occupied by airline a's exempt flights. This insures that all the slots occupied by exempt flights will be assigned to the corresponding airline. However, by including these slots within the model, theses slots "count against" the allocation received by the respective airline. ***By doing this, if a large number of an airline's flights are exempted, then that airline will receive a correspondingly higher allocation of delay to its flights that are not exempt.***

We applied the model to all of the GDPs from Figure 1. To evaluate the effectiveness of the model, we compared the delay obtained under the unconstrained RBS allocation and the delay that would have been obtained using the optimization model. The results for Boston's Logan airport are shown in Figure 2. Figure 2 depicts, for a selected number of airlines, the difference between an airline's average delay under RBS without exemptions and under the optimization model (a negative number again means the airline would have been allocated less delay if exemptions were not taken into account).

It is instructive to compare Figure 2 with Figure 1, which shows the difference in delay between the unconstrained RBS allocation and RBS with exemptions (the current procedure). Clearly, the optimization model has a significant impact and is able to reduce the biases substantially. For example, for the major carrier mentioned earlier (USA), the use of the optimization model reduced the average (per flight) difference with the unconstrained RBS allocation from –11.7 minutes to –3.4 minutes, while for the smaller carrier (UCA) the exemption bias was reduced from -18.2 to -5.4 minutes per flight. These reductions are very significant and would have a major impact on the performance of the airlines involved.

It is also important to consider potential "side effects" of the model, including the impact on the distribution of delays among flights. The optimization model's impact on the distribution of delays is shown in Figure 3. Figure 3 depicts, for a selected number of airlines (which constitute over 95% of all flights during the GDPs), both the distribution of delays that was obtained under current RBS procedures and the distribution of delays that would have been obtained by the optimization-based approach. In these graphs the horizontal axis gives delay values in minutes and the vertical axis the number of flights having the corresponding value.
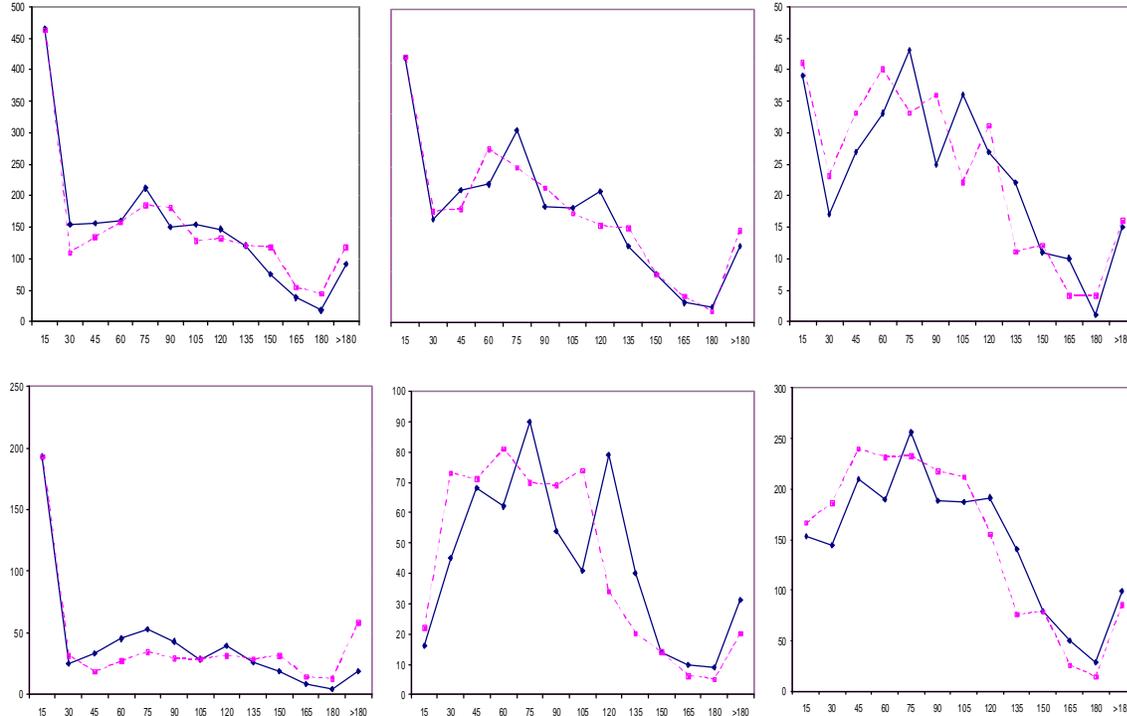
6

Figure 3: Delay Distribution Impact, Logan Airport Boston

In each of the graphs, the solid lines represent the distribution of delays under current RBS procedures, whereas the dashed lines represent the distribution of delays that would have been obtained using the optimization-based approach. To determine these distributions, we used the same GDPs at Boston's Logan airport as in the previous experiments. The results in Figure 3 indicate that, on the aggregate, the use of the optimization model appears to have a relatively minor influence on the distribution of delays. The impact is most severe for one medium-sized carrier, which would see a sizeable increase in large delays ($\ge$ 180 minutes). Intuitively, the reason for this is that this carrier has a high percentage of flight exemptions (approximately 60% of its flights are exempt), while it has a relatively small number of flights in a GDP (approximately 18 flights per GDP). As such, there will be little opportunity to shift the delay "benefits" absorbed by the exempt flights to its non-exempt flights. We note that if it was desirable to put a limit on maximum delay allocated then this could easily be incorporated into the model.

5. Other Applications

We have focused on a particular type of equity compensation mechanism, which might be used to offset the deviation of actual allocation from a baseline ideal. As mentioned earlier, there are other forms of equity disturbances in a GDP.

The first disturbance, which is currently addressed very well through the compression algorithm, is the occurrence of flight cancellations. In keeping with the equity principles espoused earlier, we would say that when an airline cancels flight generating a slot it cannot use, it should be compensated. Of course, this is exactly what the compression algorithm, which was an original fundamental component of CDM, does. In fact, the model given in Section 2 can be adapted to produce results essentially equivalent to compression (see [2], [8]). While using the model in this way does not produce any new functionality, it is important and noteworthy that the algorithm of Section 2 provides a unified approach to both RBS and compression.

A second disturbance is the popup flight described earlier. Current procedures allocate a delay to a popup flight that is approximately equal to the delay received by comparable flights, meaning flights that were scheduled to arrive at about the same time. No explicit carrier-level controls or checks are considered. Thus, certain carriers could systematically "generate" many more popups than other carriers and, as a consequence, gain an

7

advantage. In fact, it is easy to see that popups are essentially equivalent to exemptions. That is, they are slots arbitrarily allocated to specific flights. Thus, the model we have just described could easily be applied to popups, thereby mitigating any potential biases associated with the degree to which an airline generated popups.

Finally, we mention a third disturbance: the recently implemented arrival slot change procedure, which states that when a flight cannot meet its controlled time of departure (hence will miss its designated arrival slot), it must dynamically obtain a new arrival slot from TFM. In this case, dynamic reassignment of an arrival slot might be done according to slot availability and deviation from the ideal slot allocation.

It seems clear that the approach we have described provides a general setting to mitigate biases from a wide range of disturbances. The fact that it is so general is extremely appealing. This implies that by including it in the CDM decision support tools a range of situations can be addressed. This generality also would seem to hold promise for addressing other contexts where equitable allocation is of interest, for example in the en route airspace. Potentially, this could be done by extending the analogy with balanced JIT problems to more complex versions of the problem (see [12] for a description of variations of balanced JIT problems). Moreover, we note that it is relatively straightforward to apply the model with alternative fairness standards (e.g. the proportional allocation discussed in Section 2; see also [8] for computational results).

6. Practical Considerations and Implementation Issues

Before the concepts presented in this paper can be implemented, several implementation issues must be resolved, which we list here.

(1) Part of the appeal of an equity compensation mechanism is that it would (partially) alleviate TFM from making equity-based decisions, thus allowing them to focus more on their forte, demand-capacity balancing. But extreme behavior, such as choosing an unusually small GDP geographical scope, may lead to inequities so large that they cannot be effectively restored. The limits of an equity restoration mechanism would have to be well understood. On a more general level, the choice of geographical scope poses an interesting yet complex trade-off: a large scope (i.e. fewer exemptions) may lead to an inequitable distribution of (in hindsight) unnecessary delay, if the predicted reduction in capacity does not materialize, while a small scope may lead to an inequitable distribution of actual delays if the capacity reductions do materialize. As such, integrating the approach described here with GDP models that incorporate uncertainty ([4], [6]) presents a topic for further research.

(2) We have discussed equitable allocation largely in the context of a static planning scenario. The reality of GDP planning, however, is that programs are frequently revised and/or extended due to weather uncertainty and its impact on arrival capacities. As time progresses, flights are being launched toward the GDP airport and the ability to control them (via ground delay) is lost. These dynamics naturally dictate that the ability of an allocation scheme to restore equity is degraded over time as well. It should be noted, however, that the current approaches do not take into account equity restoration during GDP revisions. Both the TFM specialists as well as the decision support tools must be able to handle the dynamics of GDP planning.

(3) Under current GDP allocation procedures, each carrier reserves the right to reallocate among its flights the slots that have been allocated to it. This is done in a cancellation and substitution process. In particular, they reserve the right to trade arrival slots (times) between long-haul and short-haul flights. So, one must fully think through the interactions between the air carrier cancellation-substitution process and any equity scheme that is based in any way on flight stage length. Since air carrier substitution practices are highly dynamic and carrier-specific, this might require human-in-the-loop exercises with the air carriers.

(4) Allocation procedures are constantly evolving, as the aviation community addresses new equity concerns. How robust are the algorithms and procedures proposed here with respect to these types of changes? (We note that an appeal of the approach presented is that it appears to be much more general and robust than current approaches).

(5) For one airport, our analyses have shown that current procedures have a systematic bias against ce certain carriers. But due to operational practices, a carrier may receive an advantage at one airport and a disadvantage at another. Do these factors average out over time, when GDPs at all airports are taken into account? Community acceptance of an equity compensation mechanism might hinge on the results of a more extensive analysis.

(6) A basis for ideal allocation has been established, which relies heavily on associating

each flight with an air carrier. But general aviation flights are usually independently owned and operated. How will arrival slots be allocated to general aviation flights? Is it reasonable to treat the set of general aviation flights as a single entity.

Acknowledgements

8. References

[1] Wambsganss, M., 1996, "Collaborative Decision Making Through Dynamic Information Transfer", *Air Traffic Control Quarterly*, **4**, 107-123.

[2] Vossen, T. and M.O. Ball, 2001, "Optimization and Mediated Bartering Models for Ground Delay Programs", submitted for publication, available at: http://bmgt2-notes.umd.edu:8080/Faculty/KM/papers.nsf

[3] Ball, M.O., R. Hoffman, D. Knorr, J. Wetherly and M. Wambsganss, 2000, "Assessing the Benefits Of Collaborative Decision Making", in Air Traffic Management", in *Proceedings of 3rd USA/Europe Air Traffic Management R&D Seminar.*

[4] Ball, M.O. and G. Lulli, 2002, "Ground Delay Programs: Optimizing over the Included Set Based on Distance", submitted for publication, available at: http://bmgt2-notes.umd.edu:8080/Faculty/KM/papers.nsf

[5] Richetta, O. and A. Odoni, 1993, "Solving Optimally the Static Ground Holding Problem", *Transportation Science*, **27**, 228-238

[6] Ball, M.O., R. Hoffman, A. Odoni and R. Rifkin, 2003, "A Stochastic Integer Program with Dual Network Structure and its Application to the Ground Holding Problem", to be published in *Operations Research*, available at: http://bmgt2-notes.umd.edu:8080/Faculty/KM/papers.nsf

[7] Innis, T. and Ball, M.O., 2002, "Estimating One-Parameter Airport Arrival Capacity Distributions for Air Traffic Flow Management", submitted for publication, available at: http://bmgt2-notes.umd.edu:8080/Faculty/KM/papers.nsf

[8] Vossen, T., 2002, *Fair Allocation Methods in Air Traffic Management*, PhD Dissertation, U. of Maryland.

[9] Young, H.P., 1994, *Equity in Theory and Practice*, Princeton University Press, Princeton, NJ.

[10] Miltenberg, J., 1989, "Level Schedules for Mixed Model Assembly Lines in Just-in-Time Production Systems", *Management Science*, **35**, 192-207.

[11] Balinski, M.L. and N. Shihidi, 1998, "A Simple Approach to the Product Rate Variation Problem via Axiomatics", *Operations Research Letters*, **22**, 129-135.

[12] Kubiak, W, 1993, "Minimizing Variation of Production Rates in Just-In-Time Systems: A Survey", *European Journal of Operations Research*, **66**, 259-271.

Biographies:

Michael O. Ball: Michael Ball is the Orkand Corporation Professor of Management Science in the Robert H. Smith School of Business at the University of Maryland. He also holds a joint appointment within the Institute for Systems Research (ISR) in the Clark School of Engineering. Dr. Ball received his Ph.D. in Operations Research in 1977 from Cornell University. His research interests are in network optimization and integer programming particularly as applied to problems transportation systems. He is co-Director of NEXTOR, the National Center of Excellence for Aviation Operations Research, and he leads the NEXTOR Collaborative Decision Making project.

Robert Hoffman: Robert Hoffman is a senior analyst with Metron Aviation, Inc.. He received his Ph.D. in Mathematics from the University of Maryland in 1997, with a concentration in Operations Research. Dr. Hoffman holds a visiting faculty research position with the Institute for Systems Research at the University of Maryland, and is affiliated with NEXTOR, acting as an interface between academic and private sector research efforts. His current work includes statistical analyses of air traffic flow management initiatives, modeling traffic into major airports and through sectors, metric development for aviation performance, and resource allocation models for collaborative routing efforts.

Thomas Vossen: Thomas Vossen is an assistant professor of operations management and information systems in the the Leeds School of Business at the University of Colorado at Boulder. He received his Ph.D. in Management Science in 2002 from the

University of Maryland. His research interests focus on real-time transportation and logistics operations, in particular in air traffic flow management.

Mike Wambsganss: Mike Wambsganss is President and CEO of Metron Aviation, Inc., a lead developer in CDM decision support systems, including Flight Schedule Monitor (FSM), which is currently being used for CDM Prototype Operations. He began his career as an officer in the U.S. Air Force and has worked as an analyst at TWA, at the Mitre Corporation, and at Quantech Statistical Consultants, before joining Metron. Mr. Wambsganss is actively involved in the FAA's Collaborative Decision Making program, serves as co-chair of the CDM analysis working group, and manages MAI's transportation research and development programs.