

Analysis of Demand Uncertainty Effects in Ground Delay Programs

Michael Ball, Thomas Vossen
Robert H. Smith School of Business
and Institute for Systems Research
University of Maryland
College Park, MD 20742, USA

Robert Hoffman
Metron Aviation, Inc.
131 Elden St. Suite 200
Herndon, VA 20170, USA

October 15, 2001

Abstract

The Federal Aviation Administration and the aviation community within the U.S. have recently adopted new operational procedures and decision support tools for implementing and managing Ground Delay Programs (GDPs) based on the Collaborative Decision Making paradigm. In this paper, we investigate the impact these procedures have had on aircraft arrival time uncertainty during GDPs and, more generally, on the consequences of arrival demand uncertainty for GDP performance. Our analysis employs two models: a stochastic integer program and a simulation model; results are augmented with historical analysis. The integer program is interesting in its own right in because of an embedded binomial probability distribution. Our analysis produces three results. First, we compare the relative significance of some of the most prevalent sources of demand uncertainty. Second, we indicate how changes in the current practice for setting airport arrival rates can lead to significant benefits. Finally, we discuss how the combination of timely cancellation notices and the use of the Compression algorithm (an inter-airline slot exchange procedure that attempts to maximize utilization in the presence of delays and cancellations) may affect the uncertainty of flight arrival times during ground delay programs.

Introduction

The Federal Aviation Administration (FAA) and the aviation community within the U.S. have recently adopted Collaborative Decision Making (CDM) as a new approach toward air traffic flow management. CDM is based on the recognition that improved data exchange and communication between the FAA and the airlines will lead to better decision making. In particular, the CDM philosophy emphasizes that decisions with a potential economic impact on airlines should be decentralized and made in collaboration with the airlines whenever possible (see [5] and [8]).

While the CDM paradigm applies to a wide range of applications in Air Traffic Flow Management, its initial implementation has focused on the development of new operational procedures and decision support tools for implementing and managing Ground Delay Programs (GDPs). A GDP is a control action taken by the FAA to reduce arrival flow into an airport suffering from degraded arrival capacity or excess demand. Typically, capacity reductions are caused by bad weather. In a GDP, flights bound for congested airports are delayed on the ground so the arrival demand will match the arrival capacity. The numerous GDP enhancements implemented under CDM include improved data-exchange, better situational awareness tools, and increased flexibility for the airlines. The most significant enhancements, however, have been changes to resource

allocation procedures. Under CDM, arrival capacity is allocated to the airlines by a procedure called Ration-by-Schedule (RBS). RBS is based on the consensus recognition that airlines have claims on the arrival schedule based on the original flight schedules, and has removed disincentives airlines previously had for providing accurate information. In addition, CDM has introduced a new procedure for inter-airline slot exchange called Compression. This procedure aims to optimize arrival capacity usage in the presence of delays and cancellation, in a fair and equitable manner. The effects of these procedures has been significant: since initial implementation of GDP prototype enhancements in January of 1998, over six million minutes of assigned ground delay have been avoided (see [1], [2]).

In this paper, we take a closer look at the impact of these procedures, with a focus on their effects on demand uncertainty (e.g. the uncertainty of flight arrival times during GDPs) and, more generally, on the consequences of demand uncertainty for GDP performance. First, we present a simulation model to analyze how uncertainty reduction can lead to smaller airborne arrival queues needed to maintain high levels of airport utilization. We also use this model to highlight the relative significance of some of the most prevalent sources of demand uncertainty. Second, we present a stochastic integer programming model to determine optimal arrival rates in the presence of demand uncertainty. Computational results obtained with this model indicate that changes in the current practice for setting arrival rates might have significant benefits. Finally, we discuss how the combination of timely cancellation notices and the use of the compression algorithm has reduced the uncertainty of flight arrival times during ground delay programs, and we use our models to quantify the resulting benefits. To the best of our knowledge, our study is the first to analyze the effects of demand uncertainty on GDP performance: prior studies concentrated on stochastic capacity as the source of uncertainty (e.g., see [6], [4]).

Model Description

The implementation of GDPs consists of the assignment of ground delays to individual flights in accordance with a temporarily reduced airport arrival capacity (airport acceptance rate – *AAR*). The process can be viewed as one of assigning




Figure 1: Queueing model representation of flight arrivals

landing time slots to flights. Once a flight receives its arrival time slot, a corresponding ground delay can be assigned. Each GDP produces a *planned* arrival sequence, which states for each flight the time it is supposed to arrive at the airport. The *actual* arrival sequence, however, may differ substantially: flights can experience delays or be cancelled, while other flights can arrive unexpectedly. Such perturbations of the arrival sequence introduce an element of unpredictability for the FAA’s traffic flow specialists, which may result in airborne holding delays or airport under-utilization. This is perhaps clearer if we view the arrival of flights during a GDP as a single-server queueing system, as shown in Figure 1. Note that there are three key measures of performance: ground delay, airborne delay and airport throughput. The traffic flow specialists have direct control over the ground delay assigned. However, the resulting airborne delay and throughput depend on the assigned ground delay together with a variety of stochastic elements including those that perturb demand and those that perturb capacity. The effective planning and control of a GDP involves delicately balancing these three performance measures. An aggressive policy would assign relatively small amounts of ground delay, which would insure close to maximum throughput at the expense of large airborne delay. A conservative policy would assign a relatively large amount of ground delay, which would result in little or no airborne delay and possible underutilization of airport capacity.

In the models in this paper, we evaluate the performance of a GDP according to the three measures: ground delay, airborne delay, and uti-

lization. Here utilization equals (actual arrival rate)/(arrival capacity). We note that in practice, utilization is not precisely defined; airport capacity depends on a variety of subtle factors, such as aircraft arrival mix and the current rate of departures.

Model Assumptions

To model a GDP as a single-server queueing system, we assume that both the number of time periods T and the start time S of the GDP are known, and that the airport acceptance rate (AAR) is deterministic and constant in all periods. While these assumptions are a simplification we feel our model accurately represents the essential GDP tradeoffs. Furthermore, we assume there are n flights, f_0, \dots, f_n controlled by the GDP, and that each flight f_i has an en-route time ete_i , and an airline-scheduled arrival time oag_i , usually taken from the Official Airline Guide (OAG). For purposes of our analysis, we consider the number of flights vying for resources to be sufficiently large to warrant a GDP.

Under this representation, a GDP is completely specified by the following three parameters:

1. the service time $s = \frac{1}{AAR}$;
2. the inter-arrival time a (the time between consecutive flight arrivals);
3. the flight sequence f_1, \dots, f_n .

Once the air traffic specialist has determined appropriate parameters, each flight is assigned an arrival slot (i.e. flight f_i is assigned slot i). Flight f_i is scheduled to arrive at time $t_i = S + ia$. This fixes flight f_i 's total planned delay, and therefore its ground delay, at $S + ia - oag_i$ (we note that sometimes the ground delay assigned is adjusted depending on factors such as changes to flight plan, conditions at the origin airport, etc.). The actual arrival time, of course, is in general different from the scheduled arrival time. It can be represented as $t'_i = S + ia + w_i$, where w_i is a random variable that can be positive or negative.

As the service time is determined by the actual airport capacity AAR , the implementation of a GDP requires the following two decisions: the appropriate inter-arrival time a , and the flight sequence f_0, \dots, f_n . Flight sequence decisions affect only the *distribution* of delay among flights, and are based on equity principles embedded in the Ration-By-Schedule algorithm (see [7] for an overview). The distribution of delays is beyond

the scope of this paper, and we shall assume a predetermined flight sequence. Instead, we concentrate on the effect of the inter-arrival times on *overall* delay changes in the presence of demand uncertainty. Inter-arrival times are determined by the planned airport acceptance rate $PAAR$, that is, $a = \frac{1}{PAAR}$.

We emphasize the distinction between airport acceptance rates (AARs) and planned acceptance rates (PAARs). An AAR reflects arrival capacity; it is the number of arrivals per unit time that can be accommodated by the airport, given physical constraints such as weather, runway configuration, and the number of departures. A PAAR reflects a control decision on the part of the FAA; it is the number of arrival slots per unit time allocated by the traffic flow specialist. The PAAR may be set higher than the AAR to allow for future cancellations, or set lower than the AAR to accommodate demand not known at the time of GDP planning. Both AARs and PAARs may fluctuate over the life of a GDP.

Sources of Demand Uncertainty

Deviations from planned arrival times may create demand uncertainty in the form of inter-arrival time variation. This may occur when the deviations are unknown at the time the GDP is planned or, more generally, whenever a flight's updated information cannot be used anymore for a possible mitigating response [Tom: I don't follow this last comment]. Historic analyses of GDPs indicate that the primary sources of demand uncertainty are flight cancellations, unexpected arrivals ("pop-ups"), and aircraft arrival time drift (see [3]).

Cancellations Cancellations cause unexpected gaps in the arrival sequence when they are not anticipated at the time the GDP was implemented. Typically, this occurs when an air carrier fails to notify the FAA of a flight cancellation, but other reasons (e.g. mechanical and upstream delays, flight diversions) are also possible.

In our models, each flight will be canceled (not show up) with a probability p_{cnx} . If cancellation were the sole source of uncertainty, the inter-arrival times would follow a geometric distribution with mean $2/(1 - p_{cnx})$ and standard deviation $2p_{cnx}/(1 - p_{cnx})^2$.

Pop-ups Pop-ups are flights that arrive at the airport but were not expected at the time the

GDP was implemented. Pop-ups are usually general aviation aircraft, military aircraft, or last-minute flights created by scheduled carriers.

In our analysis, we assume that the stream of pop-ups follows a Poisson process, that is, their inter-arrival times are exponentially distributed with mean λ .

Drift Aircraft drift represents the situation in which a flight deviates from its assigned arrival time. The primary causes for this are en-route congestion and late departure from their origin airports. Although flights can arrive early, they tend to arrive later than their assigned arrival time. Significant drift is commonly caused by en-route congestion or departure delays (e.g. taxi-out and gate delays).

In our analysis, we represent drift by a uniformly distributed displacement. More specifically, flight f_i will arrive at the airport at time $t_i = S + ia + u_i$, with $u_i \approx U(-5, 15)$. The parameters -5 and +15 have been set to correspond to the FAA’s current policy of controlled departure time adherence: no earlier than five minutes prior to the controlled departure time and no later than fifteen minutes after the controlled departure time.

While we believe these distributions provide a reasonable approximation of the uncertainty during GDPs, further empirical studies are necessary (and currently underway) to determine more precise models. We remark that the parameters used in the empirical studies that follow are based on initial analyses of the frequency of cancellations and pop-ups, and of the ranges of drift (e.g., see [3]).

Effects of Demand Uncertainty

To analyze the effects of uncertainty in the arrival stream, we performed a simulation study of the single-server queueing model and the influence of the inter-arrival time decisions by ATCSCC. In all cases, we used an arrival rate $AAR = 30$ and $T = 4$ periods. All simulation averages were based on 1000 replications.

As a first step, we studied the effect of increasing PAARs on GDP performance (airborne holding and ground holding) by running the simulation model with a cancellation probability $p_{cnx} = 0.15$ as the single source of uncertainty. The results are




Figure 2: Impact of increasing PAAR.

shown in Figure 2. In Figure 2(a), the solid lines represent the average ground delay, the dashed lines represent the average airborne delay, and the dotted lines represent average overall delay. Figure 2(b) shows the utilization. As expected, in the presence of unexpected cancellations, it is possible to reduce overall delay by raising the PAAR (sending more aircraft). The price paid, however, is an increase in airborne holding. Note that this is true even if we raise the PAAR above the expected number of arrivals in each period. This effect tapers off once the airport is close to fully utilized; further PAAR increases simply convert ground holding into airborne holding. Next, we studied the “marginal” effects of the various sources of uncertainty, that is, the effect each source has on GDP performance in the presence of the other uncertainty sources. The graphs in Figure 3 illustrate the variation in airborne delay as a function of the levels of different sources of demand uncer-

tainty. In each graph, the vertical axis measures airborne delay in average minutes per flight, while the horizontal axis measures cancellation probability, pop-up interarrival rate, and drift, respectively. Each curve on a graph represents a different level of airport utilization.




Figure 3: Marginal effects of uncertainty sources.

The results show that of the three uncertainty sources, drift has the greatest impact on airborne queue size. This highlights the importance of recent efforts by the FAA to reduce drift by removing en-route restrictions imposed on aircraft bound for a GDP airport and by tightening the window in which controlled flights can depart.

Optimal Planned Acceptance Rates

The planned airport acceptance rate is a nominal figure used by the FAA in resource allocation to realize an actual acceptance rate they have set as a goal. Higher planned rates lower the total amount of FAA-assigned ground delay, but risk incurring larger amounts of airborne holding (if more flights approach the airport than can be accommodated). Lower planned rates increase total ground delay at the risk of airport under-utilization. The simulation study takes these decisions as an input, and furthermore assumed that the planned rates were constant for the duration of the program. We developed a stochastic integer program to determine an “optimal” planned rate for each period of the GDP in the presence of demand uncertainty. By optimal, we mean rates that minimize the overall amount of airborne holding, given a constraint on the capacity utilization. Essentially, the GDP is modelled as a Markov-decision process. Under this approach, it is difficult to take into account the effects of drift. We have not yet found an efficient way to incorporate this source of uncertainty into the model.

Integer Programming Model

The formulation of an integer programming model for determining optimal planned acceptance rates proceeds in two steps. First, we develop a general formulation with non-linear constraints, which may be difficult to solve. We then derive a second model in which the original constraints are linearized.

In the first formulation, we have the following decision variables.

- $x_{a,t} \in \{0, 1\}$; $x_{a,t} = 1$ if the acceptance rate during period t equals a , and 0 otherwise.
- $q_{j,t} \geq 0$; $q_{j,t}$ represents the probability that there are j flights in the queue at the start of period t .

The model constraints are as follows.

- In each period, exactly one planned acceptance rate is selected, that is,

$$\sum_a x_{a,t} = 1, \quad \text{for all } t \in 1, \dots, T. \quad (1)$$

- The queue probabilities during each period sum to one, e.g.,

$$\sum_j q_{j,t} = 1, \quad \text{for all } t \in 1, \dots, T+1. \quad (2)$$

In addition, we have $q_{0,1} = 1$, that is, the airborne queue is empty initially.

- Transition probabilities are defined as follows.

$$q_{j,t+1} = \sum_{a,i} p_{i,j|a} q_{i,t} x_{a,t}, \quad \text{for } j, t \in 1, \dots, T. \quad (3)$$

$p_{i,j|a}$ represents the probability that there will be j flights in the queue at the end of a period, given that there are i flights in the queue at the start of the period and the acceptance rate equals a . Given the cancellation and/or pop-up probabilities, these coefficients are easily calculated.

- The utilization required may be expressed as

$$\sum_{a,i} e_{i,a} q_{i,t} x_{a,t} \leq \epsilon \quad \text{for } t \in 1, \dots, T. \quad (4)$$

$e_{i,a}$ represents the expected number of unutilized slots during period t , given that there are i flights in the queue at the start of the period and the acceptance rate equals a . Again, these coefficients are readily determined. ϵ represents the maximum number of unutilized slots allowed in each period.

Finally, the objective function is defined as

$$\min \sum_{j,t} j q_{j,t}, \quad (5)$$

that is, we wish to minimize the expected amount of airborne delay.

Constraints (3) and (4) contain non-linear terms. To linearize these terms, we introduce a new set of variables:

- $\hat{q}_{j,t,a} \geq 0$; $\hat{q}_{j,t,a}$ represents the probability that there are j flights in the queue at the start of period t if the planned acceptance rate during that period is a , and equals 0 otherwise.

Using these variables, constraint set (2) can be expressed as

$$\sum_j \hat{q}_{j,t,a} = x_{a,t}, \quad \text{for all } a, t \in 1, \dots, T+1.$$

Constraint (3) can be replaced by the following pair of constraints

$$\sum_{a'} \hat{q}_{j,t+1,a'} - \sum_i p_{i,j|a} \hat{q}_{i,t,a} \leq 1 - x_{a,t},$$

$$\sum_{a'} \hat{q}_{j,t+1,a'} - \sum_i p_{i,j|a} \hat{q}_{i,t,a} \geq 1 - x_{a,t},$$

for all $a, j, t \in 1, \dots, T$.

Similarly, the utilization requirement (4) will be replaced by

$$\sum_{a,i} e_{i,a} \hat{q}_{i,t,a} \leq \epsilon \quad \text{for } t \in 1, \dots, T,$$

and the final objective function is expressed as

$$\min \sum_{j,t,a} j q_{j,t,a}.$$

Observe that the resulting formulation now represents a linear integer program.

Results

While the resulting IP formulation may take considerable time to generate an optimal solution and does not allow the incorporation aircraft drift as a source of uncertainty, empirical results obtained with the model exhibited insightful patterns, which could have a significant impact on policies for setting planned rates during GDPs. To illustrate this, we first note that traffic flow specialists use hourly acceptance rates that are more or less constant during the course of a GDP (see Figure 4 for typical patterns). This is driven largely by the anticipation of a fixed runway configuration, which carries with it a well-established arrival capacity. Traffic flow specialists recognize the danger of under-utilization due to gaps in the arrival stream. Sometimes, their mechanism for coping with this stochastic element is to plan a higher rate in the first one or two periods of the GDP, this creating an airborne queue, known as a managed arrival reservoir (MAR).

In contrast to the deterministic policy of uniform arrival rate, the optimal solution to our integer programming model consistently showed a staircase pattern of rising and falling arrival rates. Sample results are shown in Figure 5. These results not only validate the wisdom behind a MAR, they take it one step further. Rather than starting the GDP with one big MAR, it may be best to generate a series of smaller MARs. Intuitively, this says that if airport utilization is of paramount







Figure 4: Typical PAARs used by ATCSCC.

Figure 5: Sample Paars from IP model

interest, then the best way to guard against demand uncertainty is to maintain a steady arrival queue by sending periodic bursts of aircraft to the GDP airport.

Discussion: Effect of CDM Procedures

The results discussed in the previous sections show that uncertainty in the arrival stream may have a significant impact on GDP performance. In particular, our results highlight the influence of arrival time drift, which reinforces current efforts by the FAA to impose stricter controls on controlled departure times. Moreover, the results we obtained using the integer programming model indicate that airborne holding may be reduced by more creative policies for setting planned acceptance rates.

In light of these results, a logical question -

which we are currently investigating - is how the procedures instituted under CDM have affected demand uncertainty, and consequently overall GDP performance. So far, the benefits of CDM have been measured primarily in terms of the reduction in assigned ground delay (e.g., by compressing flights into vacated arrival slots); however, analysis of historical data shows that CDM appears to have had an equally important effect on the predictability of the arrival sequence. Figures 6 and 7 show two such effects. Figure 6 depicts the distribution of cancellation notice times before and after the introduction of CDM (relative to the original estimated time of departure of the flights). After the introduction of CDM, flight cancellations were on average known an hour and a half earlier than before CDM was introduced. Figure 7 shows the average drift during the course of a single CDM-based GDP. The graph depicted in Figure 7 represents the average absolute difference between a flight's current estimated arrival




Figure 6: Cancellation notice times before and after CDM




Figure 7: Drift evolution during GDPs

time of the flights and the time associated with the slot it has been assigned to (in minutes), at the given point in time. The averages are taken over the flights that were not cancelled and had not yet taken off, and the horizontal axis represents the time from the start of the GDP. The vertical lines represent times at which the compression algorithm was executed. Observe that each time compression was executed, drift was reduced significantly. As such, the improved information exchange and ability to dynamically readjust the arrival sequence of CDM appears to have had a significant effect on demand uncertainty.

One question is how these changes have impacted overall GDP performance. To properly assess the benefits, one must acknowledge that the introduction of CDM may have altered the FAA's decision process (e.g. the planned arrival rates that were used before and after the start of CDM). At one extreme, planned arrival rates may not have changed, and compression of arrival slots may have increased airport utilization. If so, then the primary benefit of CDM has been

a reduction in FAA-assigned ground delay. At another extreme, airport utilization may have remained constant, while planned arrival rates have been reduced. In this case, the primary benefit of CDM has been a conversion of airborne delay into ground delay. Our current research efforts focus on determining which of these cases might have occurred, by using historical data to determine the changes in planned arrival rates and utilization since the start of CDM.

Acknowledgments

This work was supported by NEXTOR, the National Center of Excellence for Aviation Operations Research, under Federal Aviation Administration Research Grant Number 96-C-001 and contract number DFTA03-97-D00004. Any opinions expressed herein do not necessarily reflect those of the FAA or the U.S. Department of Transportation.

References

- [1] Ball, M.O., R. Hoffman, W. Hall, and A. Muharremoglu, 1998, "Collaborative Decision Making in Air Traffic Management: A Preliminary Assessment", NEXTOR technical report RR-99-3.
- [2] 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.
- [3] Gallus, D., C. Ermatinger and M. Venanzi, 2000, "Airport Under-delivery Project", presented at the September 2000 CDM meeting.
- [4] Ball, M.O., Hoffman, R., Odoni, A. and Rifkin, R., 2000, "Efficient Solution of a Stochastic Ground Holding Problem", *Working Paper*.
- [5] Hoffman, R., W. Hall, M. Ball, A. Odoni, and M. Wambsganss, 2000, "Collaborative Decision Making in Air Traffic Flow Management," *Working Paper*.
- [6] 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.
- [7] Vossen, T. and M.O. Ball, 2001, "Optimization and Mediated Bartering Models for Ground Delay Programs", *Working Paper*.

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

Biographies

Michael Ball is the Orkand Professor 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 PhD in Operations Research in 1977 from Cornell University. He is currently Director of Research for the Smith School and Associate Director of NEXTOR, the National Center of Excellence for Aviation Operations Research. Dr. Ball's research interests are in network optimization and integer programming particularly as applied to problems in transportation systems. Dr. Ball is, or has been, area editor for Optimization for Operations Research and associate editor for Networks, IEEE Transactions on Reliability, Operations Research and Operations Research Letters

Thomas Vossen is a doctoral student in Management Science at the Robert H. Smith School of Business. His dissertation considers the development of resource allocation procedures in a collaborative air traffic management environment.

Robert Hoffman is a senior analyst with Metron Aviation, Inc., in Herndon VA. 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 the National Center of Excellence for Aviation Operations Research (NEXTOR), acting as an interface between the 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, the development of metrics for aviation performance, and resource allocation models for collaborative routing efforts between air carriers and the FAA.