

Contingencies and Cancellations in Ground Delay Programs

Thomas R. Willemain, Ph.D.

Distinguished Visiting Professor, Federal Aviation Administration

and

Professor, Department of Decision Sciences and Engineering Systems

Rensselaer Polytechnic Institute

May 30, 2001

NEXTOR WP-01-3

Abstract

When weather or other circumstances reduce the landing capacity of an airport below the level of demand, arriving aircraft experience airborne holding. A ground delay program (GDP) controls the arrival flow by holding some aircraft at their departure airports for specified periods of time, converting airborne holding into less expensive and less risky ground holding. We developed a spreadsheet model of GDPs to provide an accessible and transparent platform for developing a better understanding of their properties. We used the model to address two questions. First, since GDPs are planned in anticipation of reductions in arrival capacity, when should uncertainty about those conditions affect the ground delays assigned to individual flights? Second, when is it advantageous to cancel flights rather than delay them, and which flights should be cancelled?

1. Introduction

Sometimes the capacity of an airport to accept arriving aircraft is reduced by weather or other circumstances below the level of demand. This reduced capacity causes airborne delays as aircraft are forced to queue for landing. A ground delay program (GDP) controls the arrival flow by holding some aircraft at their departure airports for specified periods of time, converting airborne delay into less expensive and less risky ground delay.

GDPs are implemented by the Air Traffic Control System Command Center (ATCSCC) after consultation with regional Federal Aviation Administration (FAA) centers and with airline operations centers. A GDP applies to a particular airport, has specified start and stop times, and sets an allowable arrival rate called a Program Airport Acceptance Rate (PAAR). Once a GDP is declared, airlines rearrange their schedules accordingly. The process by which GDPs are created is part of a program called Collaborative Decision Making (CDM) (Ball et al. 1999a). The reasoning behind CDM is attractive:

"...CDM creates an environment in which all the major stakeholders have the opportunity to view overall system performance. This information sharing greatly improves air traffic control system predictability allowing us to anticipate problems and mitigate them to improve service delivery to the customer." (http://ffp1.faa.gov/tools/tools_cdm.asp).

The implementation of GDPs is done under the Ground Delay Program Enhancement (GDP-E) component of CDM. GDP-E is organized around use of a software program called Flight Schedule Monitor (FSM). FSM is widely and routinely used to compare current and forecasted demand against airport capacity:

"Flight Schedule Monitor is used to monitor airport demand and capacity, detect any imbalance, enact traffic management programs, and analyze the effects of any traffic management action taken. FSM is currently in prototype operations at the FAA, NavCanada, and many airlines in support of the CDM Ground Delay Program Enhancements effort. As part of Ground Delay Program Enhancements effort, the FAA's Air Traffic Control System Command Center, which controls air traffic for the entire United States, uses Flight Schedule Monitor 24 hours a day, 7 days a week to manage aircraft arriving and departing U.S. airports." (<http://www.metsci.com/>)

An evaluation of CDM benefits has been completed by the FAA (Free Flight Phase 1 Metrics Team 2000) and by the National Center of Excellence for Aviation Operations Research

(Ball et al. 1998). The general impression in the aviation community is that CDM is a success and that GDPs are helpful interventions (Wambsganns 1997, Ball et al. 2000). Nevertheless, GDP-E is an evolving program, and almost every GDP encounters implementation problems. ATCSCC and contractor staffs perform retrospective analyses to document and remedy these problems.

The FSM software at the heart of GDP-E contains a mathematical programming system to compute nearly optimal assignment of ground delays to flights. There are a number of papers in the operations research literature developing theory and methods for this optimization problem (including Richetta and Odoni 1993, Vranas et al. 1994, Hoffman 1997, Rifkin 1998, Hoffman and Ball 1998, Ball et al. 1999b, Andreatta et al. 2000).

Most of these models are based on a type of integer program known as an assignment problem. The time horizon of concern is divided into uniform arrival slots in accordance with the specified airport acceptance rate. Flights are assigned to the time slots so as to minimize expected ground and airborne holding costs. Some of the models permit flight-specific ground and airborne delay costs; other models, more aggregate in nature, do not. Solutions are obtained by dynamic programming, converting integer programming formulations to network flow problems, or by using combinations of LP relaxations and heuristics. Since FSM assumes deterministic capacity, it is able to use of a simple greedy algorithm to solve the assignment problem.

The uncertainty of airport capacity in the ground holding problem has been addressed in several papers (Hoffman 1997, Ball et al. 1999b, Rifkin 1998, Inniss 2001). These models employ multiple scenarios of capacity profiles (sample paths) but are not oriented toward individual flights and precise estimation of delays. Rather, flights are regarded as interchangeable, and their airborne and ground delays are approximated using 15-minute time buckets. Using discrete time buckets and ignoring the identity of individual flights is a sensible adaptation to the implementation issues surrounding CDM, especially three: the need for fast computation using data from thousands of flights, uncertainties about actual takeoff times, and the prerogative of airlines to shift and substitute flights.

However, since it is not our goal to develop a system for routine production use, we need accept neither course granularity nor anonymity of flights. The novelty of our approach, outlined in the next section, is the use of traditional queueing theory and the retention of explicit information about individual flights. FSM has some simulation capabilities but is locked into a single capacity scenario and deterministic arrival times of flights. The purpose of our model is to develop insight into (a) the sensitivity of GDP design to uncertainty about future airport capacity and (b) the value of cancellation as an alternative to ground delay.

We implemented our model in an Excel spreadsheet, in part to make it easily accessible for use by FAA and airline personnel. The spreadsheet model is also much more transparent than a production model implemented in a specialized optimization system.

2 Methodology

Every model requires a choice of what to include and what to exclude. One hopes to achieve a reduction in complexity without losing contact with essential reality. One further hopes that the model can be made to quickly yield numerical solutions. Finally, one hopes that these solutions illuminate the original problem and develop insight.

2.1 Assumptions

We consider an airport with one runway devoted exclusively to arrivals and another to departures. We focus our attention on the arrival runway, so we model the airport as a single-server queue. We assume a maximum airport acceptance rate (AAR) of 20 aircraft per hour (3 minutes per aircraft) under ideal conditions and a reduced capacity at other times. We regard the size and duration of capacity reductions as random variables whose joint distribution may be adequately summarized by a small set of scenarios.

In the analysis in section 2.3, we consider operations conducted over a period of about one hour. During that time, we assume that 20 aircraft are expected to land. Thus, under ideal conditions, the arrival runway would not be stressed because the AAR equals demand. However, we assume that arrivals appear at somewhat irregular intervals due to deviations from planned

takeoff and air times, so there is arrival congestion even when the runway operates at full capacity, and severe congestion is possible when arrival capacity is reduced.

We regard an aircraft as appearing “at” the airport at an expected time plus any time imposed as ground delay. If the aircraft cannot land at that moment, we regard it as queueing in the vicinity of the airport.

This picture is oversimplified in three ways. First, it is not customary to have detailed knowledge of expected times at which flights approach the arrival airport, because actual takeoff times are uncertain. Second, airborne holding sometimes takes place some distance from the airport. Third, ground delays are not assigned in the form of real numbers matched to individual aircraft; rather, integer-valued blocks are assigned to airlines, which have flexibility to shift their flights among their blocks. Luckily, none of these complications appears to make any essential difference to our conclusions.

We calculate airborne holding by treating the arrival runway as a single server queue with first-come/first-served queue discipline, constant service time, and interarrival times based on the combined effects of the schedule plus any assigned ground delay. It would be relatively simple to add several complications to our spreadsheet (e.g., multiple arrival runways, random landing times), but it would not necessarily be helpful to do so.

2.2 Optimization Model

As in the literature cited above, we put the problem of designing a GDP into an optimization framework. The decision variables are the amounts of ground delay assigned to individual flights. The objective function is a weighted sum of ground delay and airborne holding and is to be minimized. The relative weight of airborne delay is sometimes taken to be twice that of ground delay (Inniss and Ball 2001); we explored the sensitivity of our results to this parameter. The objective function is nonlinear, since it requires a translation of aircraft arrival times, which depend on assigned ground delay times, into queueing delays at the arrival runway.

We accomplished this translation by embedding the queueing calculations directly into the spreadsheet.

The basic model can be expressed mathematically as follows. Let:

X_j = Ground delay time assigned to flight j , $j = 1..N$

S_j = Scheduled (expected) time at which flight j approaches the destination airport

A_j = Airborne holding of flight j

C = ratio of the cost of a minute of airborne holding to a minute of ground holding

T = time interval required to land, including final approach; assumed constant.

Then the basic problem of optimal ground holding is:

$$\text{Min } \sum (X_j + C \cdot A_j)$$

where the optimization is over $0 \leq X_j$ and the nonlinearity arises from the queueing model

$$A_j = A_j(X_1..X_N, S_1..S_N, T).$$

2.3 Optimal Ground Delays

We used the spreadsheet model to understand the basics of GDPs. Exhibit 1 shows an example of the problem of irregular arrivals creating airborne queues. The columns in the exhibit are as follows:

- Flight Number: We assumed 20 flights to keep the example simple.
- Expected Approach: The time (in minutes) at which each aircraft approaches the terminal airspace.
- Ground Delay: The amount of ground delay assigned to each aircraft before takeoff.
- Actual Approach: The sum of Expected Approach and Ground Delay, being the time at which the landing process would begin in the absence of airborne queueing.
- Airborne Delay: The amount of airborne holding caused by congestion at the arrival airport.
- Start to Land: The moment at which a given aircraft begins the landing process, i.e., begins its service time.

- Time to Land: The time interval from Start to Land to Time off Runway. A reduction in the AAR is reflected in an increase in the time to land from the normal value of 3 minutes.
- Time off Runway: The moment at which the aircraft exits the runway.
- Total Delay: The sum of Ground Delay and Airborne Delay.
- Contingency Adjustment: Used later to show changes from optimal single-scenario ground delays when there are multiple capacity reduction scenarios.

Examining Exhibit 1, we note that the expected approach times are irregular. This variation in interarrival times creates approach queues despite that fact that, in the aggregate, there is exactly enough landing capacity to handle 20 flights in one hour. Only the first flight encounters no airborne delay; the last flight encounters 14.47 minutes. Consider flight 2. It approaches the airport at time 2.42, shortly after flight 1. Since flight 1 arrives at time 2.32 and needs 3 minutes to land, flight 2 cannot use the runway until time 5.32 ($= 2.32 + 3$). Therefore, flight 2 could delay its arrival by 2.90 ($= 5.32 - 2.42$) minutes and convert its airborne holding into a ground hold.

Next we considered the effect of a reduction in AAR. For illustrative purposes, we supposed that the AAR, normally 20 aircraft per hour, is reduced to 10 aircraft per hour from time 30 minutes to time 60 minutes. Exhibit 2 shows the calculations. Note that the time to land increased from 3 minutes to 6 minutes for flights 11 to 15, which use the runway during the period of reduced capacity. As one would expect, the reduced AAR increased the airborne queueing, which reached its peak at 29.47 minutes for flight 20 and averaged 11.50 minutes per flight.

The justification for GDPs is that it is cheaper and safer to endure the queueing delays on the ground rather than in the air. Exhibit 3 shows the results of applying an optimal set of ground delays to each flight. As long as a minute of ground delay is considered less costly than a minute of air delay, the optimal solution to the optimization problem formulated in section 2.2 is simple. It is to arrange the ground delays to insure that each aircraft approaches the airport late enough so

that it does not have to wait for its predecessor to be processed. Exhibit 3 shows this result: the delays previously in the airborne column have been shifted to the ground delay column. For instance, by delaying the arrival of flight 2 for 2.90 minutes, it will arrive at the airport just as flight 1 clears the runway. (In practice, randomness in the times that flights actually approach the airport creates gaps in the arrival process that result in inefficient use of the runway. To counter this, actual GDPs often deliberately create a small airborne queue, called a managed arrival reservoir or MAR, to “keep pressure on the airport”.)

3 Planning for Contingencies

Any GDP is based on uncertain projections of the three GDP parameters: when a capacity reduction will begin, how deep it will be, and how long it will last. Since most capacity reductions are weather-related, it is obvious that there can be great uncertainty about these parameters. The uncertainty is amplified by the desire to plan GDPs as far ahead as possible, thereby stretching weather forecasts to their limit. Innis and Ball (2000) considered the problem of forecasting the duration of fog-related closures in San Francisco and documented the inherent uncertainties in that one piece of the problem. Ball et al. (1999b) presented a method to incorporate scenarios into the optimization framework used by FSM.

The comprehensive way to represent the uncertainty in the three parameters is to specify a 3-dimensional joint probability distribution. A less accurate but more accessible way is to develop a discrete set of scenarios embodying the most likely combinations. The scenario method is especially suited to spreadsheet modeling.

Any given assignment of ground delay yields different results in different scenarios. For each scenario, there is a weighted average of airborne and ground delays costs, computed using the cost ratio C . Combining these weighted averages using their respective scenario probabilities produced the overall performance measure to be minimized. This approach, evaluating each set of decision variables against a set of probabilistic conditions, is an example of stochastic programming (Birge and Louveaux 1997). Rather than making the best choices for a single

scenario, stochastic programming takes account of the uncertainty inherent in multiple scenarios, hedging against contingencies.

An important practical problem is to be able to identify situations requiring analysis beyond the most likely scenario. Birge and Louveaux (1997, p. 144) stated that there is not yet a good general answer to this question, although greater variability is often (but not always) indicative of a problem that would benefit from stochastic programming. Using the notion of stochastic programming in the spreadsheet model, we investigated the issue of how uncertainties should shape the design of a GDP. Our main result is that we discovered some simple mathematical conditions defining when a second scenario can or cannot be ignored when assigning ground delays to flights.

To keep the exposition simple, we primarily considered two scenarios. The first was the one analyzed above, in which the AAR dropped from 20 to 10 aircraft per hour for 30 minutes beginning 30 minutes into the arrival period. The second scenario was more severe, reducing the AAR to only 5 for 60 minutes. To begin with, we considered the first scenario three times more probable than the second.

Exhibit 4 shows a summary of the results and the details of each scenario. The top table in the exhibit summarizes the scenario definitions and results. The columns of the table are:

- Capacity Scenario: We used only two scenarios in our analysis. This is a smaller number than might be used in practice but sufficient to clarify the role of stochastic programming in the design of GDPs.
- Scenario Probability: We considered scenario 1 to be three times as likely as scenario 2. The extent to which the first, main, scenario dominates the contingencies turns out to be an important factor in deciding whether to pay attention to less likely scenarios.
- Start of Reduction: The time at which the AAR is first reduced. (All times are in minutes.)
- Duration of Reduction: The length of time that the AAR is reduced.
- Normal AAR: The AAR in the absence of capacity reductions.

- Reduced AAR: The value to which the AAR is reduced by weather or other problems.
- End of Reduction: The time at which the AAR returns to its normal level, computed as the sum of the Start of Reduction and Duration of Reduction.
- Air:Gnd Cost Ratio: The ratio of the cost of a minute of airborne holding to a minute of ground delay. This ratio is assumed to be greater than unity. The ratio was another key factor in deciding whether to take contingencies into account when computing the ground delays to assign to each flight.
- Avg Gnd Delay: The average of the ground delays assigned to all flights.
- Avg Air Delay: The average amount of airborne holding arising after ground delays have been assigned.
- Wgt Avg Delay: The weighted average of airborne and ground holding, using the Air:Gnd Cost Ratio as the weight to assign to airborne delays.
- Runway Utilization: The percentage of time the runway was used, computed as the sum of all landing times divided by the time until the last flight has exited the runway. One can always reduce airborne delays by spacing arrivals, but too much spacing leads to inefficient use of the runway.

The results in Exhibit 4 apply to the ground delays shown in Exhibit 3, i.e., the delays computed with regard to scenario 1 only. The middle table of Exhibit 4 gives the results for scenario 1 previously shown in Exhibit 3. Thus, we see again the set of assigned ground delays averaging 11.5 minutes and yielding zero airborne holding. The bottom table in Exhibit 4 shows something new: how the ground delay assignments would work out if scenario 2 obtained. Here we note that, in scenario 2, flights 11 to 15 required 12 minutes instead of 3 minutes to use the runway. The ground delays assigned on the basis of scenario 1 sufficed to eliminate airborne holding for flights 1 to 12, but the severely reduced capacity starting at time 30 created airborne delays for all nine flights arriving after that point. Including the zero delays for the earlier flights, the average flight endured not only 11.5 minutes of ground delay but also 10.5 minutes of

airborne delay. This example makes clear that designing a GDP around one scenario can be inefficient if another scenario obtains.

It is natural at this point to ask whether a different assignment of ground delays could produce better overall performance. A little thought makes clear that the answer depends on two parameters: the relative probability of the most likely scenario and the relative cost of airborne holding. If the most likely scenario is highly probable, little will be lost by not considering remote contingencies. And since the cost of poorly chosen ground delays is airborne delays, if airborne delays are relatively inconsequential, one can ignore scenarios in which they arise.

We can see the bad effect of trying to improve the results in unfavorable conditions by comparing Exhibits 4 and 5. In Exhibit 4, flight 12 incurs 6 minutes of airborne holding if scenario 2 obtains, and succeeding flights endure even greater airborne delays. In Exhibit 5, we add 6 minutes of contingency ground delay to flight 12 to avoid the airborne delay it suffers in scenario 2. This contingency adjustment for flight 12 is a mistake in this case. Under scenario 2, the contingency adjustment does eliminate the 6 minutes of airborne delay for flight 12, reducing the weighted average delay under scenario 2 from 10.83 in Exhibit 4 to 10.73 in Exhibit 5. However, under scenario 1, the contingency adjustment creates additional airborne delay not only for flight 12 but also for all eight following flights. This raises the weighted average delay under scenario 1 from 3.83 in Exhibit 4 to 5.03 in Exhibit 5. The net effect, weighted by the probabilities of each scenario, is to increase the combined weighted average delay from 5.58 in Exhibit 4 to 6.46 in Exhibit 5. In this example, the contingency adjustment makes no sense, given the relatively high 75:25 ratio of scenario probabilities and relatively low 2:1 ratio of costs. Even worse, under the circumstances, would be to make contingency adjustments to the ground holds imposed on not only flight 12 but flights 13 to 20 as well. Doing so further increases the combined weighted average delay from 6.46 to 7.33.

These bad results from contingency adjustment are not universal. With different ratios of costs and probabilities, allowing for the possibility of scenario 2 can definitely improve the

performance of the GDP. Furthermore, if conditions are such that making a contingency adjustment to a single flight is helpful, then making contingency adjustments to all flights with airborne delay is optimal. For example, in the present case with a cost ratio of 2:1, adding contingency delays to flights 12 to 20 exactly matches or exceeds the delay performance of adding no contingency ground delays at all, provided the probability of scenario 1 is 0.5 or lower. Looked at another way, holding the assumed probability of scenario 1 at 0.75, a cost ratio of 4:1 leads to equal delay performance; see Exhibit 6. If scenario 1 were to have a lower probability than 75%, contingency adjustment would show a net benefit. However, as Exhibit 6 also shows, contingency adjustments that increase ground delays also lower the runway utilization, which argues further for applying contingency adjustments only when the main scenario is less dominant.

In all the numerical examples built around the arrival data used in Exhibits 1 to 6, there is a numerical rule that equates the delay performance of full contingency adjustment with that of no adjustment. Let P = probability of scenario 1 and C = ratio of cost of airborne to ground holding. Contingency adjustment provides benefits in terms of the combined weighted delay whenever $(1-P)C > 1$.

To see the origin of this rule, consider Exhibit 7. The top row of the exhibit shows the two types of delay when there is no contingency adjustment. At the left are the numerical values for our dataset. At the right is a generalization, using G = average ground delay and A = average airborne delay. The bottom row of the exhibit shows the same quantities when there is full contingency adjustment, which converts all the airborne delay under scenario 2 into ground delay. There is a benefit to contingency adjustment whenever:

$$P(G + C \times 0) + (1 - P)(G + C \times A) > P(G + A + C \times 0) + (1 - P)(G + A + C \times 0).$$

Algebraic manipulation reduces this relationship to:

$$(1 - P)C > 1.$$

For example, if $C = 2$, then adjustment makes sense if $P < 0.5$. Since we think of scenario 1 as the most likely scenario, this condition is impossible to meet. But if $C = 3$, then the condition becomes $P < 0.67$.

A similar relationship can be derived if the second scenario is more favorable rather than less favorable. See Exhibit 8. In this example, scenario 1 has a reduction in AAR from 20 to 10 lasting 60 minutes, whereas the reduction only lasts 15 minutes in scenario 2. Thus, if scenario 2 obtains, there is unnecessary ground holding, which can be adjusted by reducing the ground delay of selected flights. In this case, there is a benefit to contingency adjustment whenever:

$$P(G+A + C \times 0) + (1 - P)(G+A + C \times 0) > P(G+A + C \times A) + (1 - P)(G + C \times 0).$$

Algebraic manipulation reduces this relationship to:

$$P C < 1.$$

For example, if $C = 2$, then adjustment makes sense if $P < 0.5$, which is again problematic. If $C = 3$, the problem is worse, since it is required that $P < 0.33$. In fact, contingency adjustment makes sense when the contingency is favorable only if $1 < C < 2$. That is, when airborne holding is only slightly more costly than ground holding, we should be willing to accept some additional airborne holding under the main scenario in hopes that the contingency will be favorable.

It should be possible to develop similar analyses for situations with more than two scenarios (Hoffman 1997), although we have not done so in the interest of space and clarity.

4 Canceling Flights

Airborne delay arises when arrival demand approaches or exceeds landing capacity. Queues can develop even when the average arrival rate is less than the AAR, provided there is sufficient randomness in either the interarrival times or landing times or both. If capacity is regarded as fixed in the short run, then the only option for avoiding delay is, in essence, to change the schedule of flights. One can do this either by imposing ground delays to spread out the arrival process or by canceling flights to relieve the pressure on the airport. In section 3, we considered the problem of assigning ground delays. In this section, we consider cancellations of flights.

Canceling flights is not an inherently attractive option for either the airline or the passengers. However, it is one way to relieve pressure. Here, we investigate the question of optimal cancellations. Our investigation is limited relative to the full complexity of the problem. In practice, the decision to cancel a specific flight will depend on many factors not within the scope of our spreadsheet model. These include: the number of passengers on the cancelled flight, the availability of empty seats on later flights to the passengers' destinations, and the airline's operating philosophy (some are more loathe to cancel than others). What we do consider is the optimal sequence of flights to cancel.

We modified the model used in section 3 in several ways. First, we made the decision variables binary indicators of whether to cancel the corresponding flights. Second, we considered a cancellation to be equivalent to delaying a flight by 24 hours and assumed that there would be no delays when the flights resumed; thus, the only airborne holding would be that endured by the flights not cancelled on the day in question. Third, we modeled 50 rather than 20 flights to better simulate two arrival "rushes" or "banks", in which a large number of flights appear nearly simultaneously. Such rushes are a feature of the hub-and-spoke systems used by most major airlines. For simplicity, we kept the AAR constant at 20.

Exhibit 9 shows the cumulative approach of flights to the airport (top curve) and their subsequent times clearing the runway (bottom curve). The two arrival rushes begin at about 30 minutes and 75 minutes, respectively. The last flight approaches at 119 minutes, endures 38 minutes of airborne delay and 3 minutes of landing time, to clear the runway at 160 minutes. Flights 45 – 47 all endured more than 50 minutes of airborne delay, while the average flight had 22 minutes. Clearly, this scenario puts too much pressure on the runway.

Given this level of congestion, we determined the priority for canceling flights. The priority is highly dependent on the particulars of the times at which flights approach the airport. Exhibit 10 shows the average airborne delay over all noncancelled flights as a function of which flight might be cancelled first. Canceling flight 9 would reduce the average airborne delay the

most, down to below 20 minutes. Note that it is possible to cancel flights, such as flights 1 – 7, and actually increase the average delay. This happens when cancellation removes a flight with very little or no delay from the average.

Apparently, determining the optimal sequence of cancellations involves solving a huge combinatorial problem. With 50 flights, there are 2^{50} possible sets of cancellations. The Excel Solver failed to solve this problem using its branch-and-bound algorithm. Therefore, we proceeded heuristically, solving the problem sequentially. That is, we found the best flight to cancel first: flight 9. Then, conditional on having cancelled flight 9, we found the conditional best flight to cancel second: flight 10. This method reduced the problem to a manageable sequence of at most 50 sequential one-dimensional exhaustive enumerations, each with at most 50 flights to test. In fact, the process terminated after canceling 33 flights. Local two-dimensional search produced the same beginning sequence as local one-dimensional search, increasing our confidence in the heuristic search.

Exhibit 11 shows the flights, their times of scheduled approach to the airport, their priority for cancellation, and the average airborne delay conditional on canceling according to the priority list. For example, with no cancellation, the average delay was 22.0 minutes. Canceling flight 9 reduced the average delay to 19.9 minutes. Then also canceling flight 10 further reduced the average to 17.8 minutes. The horizontal lines in Exhibit 11 mark the times of the two arrival rushes. The optimal solution skipped the first flight in a rush but then cancelled a few flights early in the rush. The optimal sequence jumped between rushes and eventually outside the rushes, always whittling down the average airborne delay.

Exhibit 12 shows the sequential reduction in average delay. Clearly, the marginal value of successive cancellations decreases.

From the perspective of minimizing total delay, whether it is advantageous to cancel flights depends on the time cost associated with cancellation. In the left half of Exhibit 13, we see that canceling flight 9 cuts the average delay for the noncancelled flights from 22.0 to 19.9

minutes. In this case, the cancellation would result in an overall decrease in delay only if cancellation were the equivalent of no more than 125 minutes of airborne delay, since $(49/50) \times 19.9 + (1/50) \times 125 = 22.0$. Exhibit 13 shows the calculations of this breakeven level of delay for canceling various numbers of flights. If it were possible to guarantee that no cancellation would cause a delay of more than 2 hours, then the cancellation of flights 9, 10 and 11 would appear to be justified. Of course, in practice, the time savings from cancellation would have to be greater than the breakeven level to offset the loss of good will associated with cancellations.

Exhibit 13 also shows the results of a second trial with a different sample of approach times. This replication gives some feeling for the variability in our results, which depend on the particular sequence of approach times. In the second trial, equating a cancellation with 2 hours of delay would justify canceling only flight 10. The optimal sequence in the second trial has the same general character as its counterpart in the first trial: the cancelled flight tends to be from early in a rush but not the very beginning, and the selection jumps back and forth from one rush to the other.

While the best flight to cancel depends on the specifics of the approach times, the choice has some robustness relative to the AAR. We re-examined the second trial in Exhibit 13, varying the AAR to see how this changed the choice of the best flight to cancel. Exhibit 14 shows the results. Usually, the best flight to cancel was the same over a range of AAR values. However, the general trend was to cancel earlier flights as the AAR decreased.

5 Summary and Conclusions

Ground delay programs (GDPs) are a response to conditions that lower the capacity of the destination airport. GDPs convert airborne holding into safer and less costly ground holding. The planning and execution of GDPs is evolving, as airlines and the FAA learn to collaborate and as more sophisticated software tools become available.

Current planning of GDPs focuses on a primary scenario, taking account of other possible scenarios in an informal way. We developed a spreadsheet model to explicitly treat a secondary scenario. We developed simple rules to determine when it would be advantageous to adjust ground delays in a way that hedges against the possibility that the secondary scenario would be the one to occur. These rules depend on the relative probability of the two scenarios and the relative cost of airborne and ground delays. The practical estimation of these parameters is an open research topic. However, the Quality Assurance department of the ATCSCC, which performs next-day analysis of GDPs for quality control, could likely use such methods to develop a policies or guidelines for GDP specialists confronted with uncertain airport capacities.

Another way to respond to excess demand is to cancel flights. We used the spreadsheet model to determine which flights to cancel in which order, conditional on the sequence of times at which the flights approach the destination airport. In our examples involving two arrival rushes, the cancellation sequence began early in the rushes and jumped back and forth between rushes. In general, the lower the airport acceptance rate (AAR), the earlier the cancellation. The choice of which flight to cancel was often stable over a range of AARs. We also used the spreadsheet model to calculate whether cancellation would be worthwhile in the first place. Generally, the first flight cancelled offers the greatest benefits and might be justified if passengers can be accommodated with less than, say, two hours of cancellation-related delay.

In an effort to understand GDP dynamics, we have relaxed the operational constraint that only the carriers can make cancellation decisions. One possible use of our method is to implement in FSM a cancellation indicator, which would point out to the carriers which flights could offer the most global delay reduction benefits. The carrier could then decide if this is a worthwhile tradeoff. Perhaps some form of credit could be issued to the carrier as an added incentive.

Acknowledgements

The author gratefully acknowledges the assistance of Mike Ball, Bob Hoffman, and Tasha Inniss.

References

Andreatta, G., Brunetta, L., and Guastalla, G. 2000. From ground holding to free flight: An exact approach. *Transportation Science* 34(4), 394-401.

Ball, M., Hoffman, R., Hall, W. and Muharremoglu, A. 1998 Collaborative decision making in air traffic management: a preliminary assessment, NEXTOR Research Report #RR-99-3

Ball, M., Chien, C-Y., Hoffman, R., and Vossen, T. 1999a. Collaborative decision making in air traffic management: Current and future research directions. *Proceedings of ATM '99, Workshop on Advanced Technologies and Their Impact on Air Traffic Management in the 21st Century*.

Ball, M., Hoffman, R., Odoni, A., and Rifkin, R. 1999b. The static stochastic ground holding problem with aggregate demands. NEXTOR Report # RR-99-1. University of Maryland. (To appear in *Operations Research*).

Ball, M., Hoffman, R., Knorr, D., Wetherly, J., and Wambsganns, M. 2000. Assessing the benefits of collaborative decision making in air traffic management. Third USA/Europe Air Traffic Management R&D Seminar. Naples, Italy. June 13-16.

Birge, J., and Louveaux, F. 1997. *Introduction to Stochastic Programming*. Springer: New York, NY.

Free Flight Phase 1 Metrics Team. 2000. An Operational Assessment of Collaborative Decision Making in Air Traffic Management: Measuring User Impacts through Performance Metrics. Final Report. February 2000. Document Control Number: R90145-05.

Hoffman, R. 1997. Integer programming models for ground-holding in air traffic flow management. PhD thesis, University of Maryland.

Hoffman, R. and Ball, M. 1998. A comparison of formulations for the single-airport ground holding problem with banking constraints. University of Maryland Institute for Systems Research Technical Report TR 98-44.

Inniss, T. and Ball, M. 2001. Estimating one-parameter airport arrival capacity distributions for air traffic flow management. Submitted for publication.

Richetta, O. and Odoni, A. 1993. Solving optimally the static ground-holding policy problem in air traffic control. *Transportation Science* 27(3), 228-238.

Rifkin, R. 1998. The static stochastic ground-holding problem. SM thesis, Massachusetts Institute of Technology.

Vranas, P., Bertsimas, D. and Odoni, A. 1994. The multi-airport ground-holding problem in air traffic control. *Operations Research* 42(2), 249-261.

Wambsganss, M. 1997. Collaborative decision making through dynamic information transfer. *Air Traffic Control Quarterly* 4(2), 107-123.

Exhibit 1: Airborne holding in the absence of a reduction in arrival capacity

Flight Number	Expected Approach	Ground Delay	Actual Approach	Airborne Delay	Start To Land	Time To Land	Time Off Runway	Total Delay	Contingency Adjustment
1	2.32	0.00	2.32	0.00	2.32	3.00	5.32	0.00	0
2	2.42	0.00	2.42	2.90	5.32	3.00	8.32	2.90	0
3	3.46	0.00	3.46	4.86	8.32	3.00	11.32	4.86	0
4	5.95	0.00	5.95	5.37	11.32	3.00	14.32	5.37	0
5	7.47	0.00	7.47	6.85	14.32	3.00	17.32	6.85	0
6	10.00	0.00	10.00	7.32	17.32	3.00	20.32	7.32	0
7	15.42	0.00	15.42	4.90	20.32	3.00	23.32	4.90	0
8	16.53	0.00	16.53	6.79	23.32	3.00	26.32	6.79	0
9	17.36	0.00	17.36	8.96	26.32	3.00	29.32	8.96	0
10	22.63	0.00	22.63	6.69	29.32	3.00	32.32	6.69	0
11	27.76	0.00	27.76	4.56	32.32	3.00	35.32	4.56	0
12	32.50	0.00	32.50	2.82	35.32	3.00	38.32	2.82	0
13	33.38	0.00	33.38	4.93	38.32	3.00	41.32	4.93	0
14	37.93	0.00	37.93	3.39	41.32	3.00	44.32	3.39	0
15	39.41	0.00	39.41	4.91	44.32	3.00	47.32	4.91	0
16	40.46	0.00	40.46	6.86	47.32	3.00	50.32	6.86	0
17	43.34	0.00	43.34	6.98	50.32	3.00	53.32	6.98	0
18	43.59	0.00	43.59	9.73	53.32	3.00	56.32	9.73	0
19	44.56	0.00	44.56	11.75	56.32	3.00	59.32	11.75	0
20	44.85	0.00	44.85	14.47	59.32	3.00	62.32	14.47	0

Exhibit 2: Airborne holding in the presence of a reduction in AAR

Flight Number	Expected Approach	Ground Delay	Actual Approach	Airborne Delay	Start To Land	Time To Land	Time Off Runway	Total Delay	Contingency Adjustment
1	2.32	0.00	2.32	0.00	2.32	3.00	5.32	0.00	0
2	2.42	0.00	2.42	2.90	5.32	3.00	8.32	2.90	0
3	3.46	0.00	3.46	4.86	8.32	3.00	11.32	4.86	0
4	5.95	0.00	5.95	5.37	11.32	3.00	14.32	5.37	0
5	7.47	0.00	7.47	6.85	14.32	3.00	17.32	6.85	0
6	10.00	0.00	10.00	7.32	17.32	3.00	20.32	7.32	0
7	15.42	0.00	15.42	4.90	20.32	3.00	23.32	4.90	0
8	16.53	0.00	16.53	6.79	23.32	3.00	26.32	6.79	0
9	17.36	0.00	17.36	8.96	26.32	3.00	29.32	8.96	0
10	22.63	0.00	22.63	6.69	29.32	3.00	32.32	6.69	0
11	27.76	0.00	27.76	4.56	32.32	6.00	38.32	4.56	0
12	32.50	0.00	32.50	5.82	38.32	6.00	44.32	5.82	0
13	33.38	0.00	33.38	10.93	44.32	6.00	50.32	10.93	0
14	37.93	0.00	37.93	12.39	50.32	6.00	56.32	12.39	0
15	39.41	0.00	39.41	16.91	56.32	6.00	62.32	16.91	0
16	40.46	0.00	40.46	21.86	62.32	3.00	65.32	21.86	0
17	43.34	0.00	43.34	21.98	65.32	3.00	68.32	21.98	0
18	43.59	0.00	43.59	24.73	68.32	3.00	71.32	24.73	0
19	44.56	0.00	44.56	26.75	71.32	3.00	74.32	26.75	0
20	44.85	0.00	44.85	29.47	74.32	3.00	77.32	29.47	0

Exhibit 3: Optimal ground delays when the AAR is reduced

Flight Number	Expected Approach	Ground Delay	Actual Approach	Airborne Delay	Start To Land	Time To Land	Time Off Runway	Total Delay	Contingency Adjustment
1	2.32	0.00	2.32	0.00	2.32	3.00	5.32	0.00	0
2	2.42	2.90	5.32	0.00	5.32	3.00	8.32	2.90	0
3	3.46	4.86	8.32	0.00	8.32	3.00	11.32	4.86	0
4	5.95	5.37	11.32	0.00	11.32	3.00	14.32	5.37	0
5	7.47	6.85	14.32	0.00	14.32	3.00	17.32	6.85	0
6	10.00	7.32	17.32	0.00	17.32	3.00	20.32	7.32	0
7	15.42	4.90	20.32	0.00	20.32	3.00	23.32	4.90	0
8	16.53	6.79	23.32	0.00	23.32	3.00	26.32	6.79	0
9	17.36	8.96	26.32	0.00	26.32	3.00	29.32	8.96	0
10	22.63	6.69	29.32	0.00	29.32	3.00	32.32	6.69	0
11	27.76	4.56	32.32	0.00	32.32	6.00	38.32	4.56	0
12	32.50	5.82	38.32	0.00	38.32	6.00	44.32	5.82	0
13	33.36	10.93	44.32	0.00	44.32	6.00	50.32	10.93	0
14	37.93	12.39	50.32	0.00	50.32	6.00	56.32	12.39	0
15	39.41	16.91	56.32	0.00	56.32	6.00	62.32	16.91	0
16	40.46	21.86	62.32	0.00	62.32	3.00	65.32	21.86	0
17	43.34	21.98	65.32	0.00	65.32	3.00	68.32	21.98	0
18	43.59	24.73	68.32	0.00	68.32	3.00	71.32	24.73	0
19	44.56	26.75	71.32	0.00	71.32	3.00	74.32	26.75	0
20	44.85	29.47	74.32	0.00	74.32	3.00	77.32	29.47	0

Exhibit 4: An analysis of delays under two scenarios

Scenario Definitions								Scenario Results				
Capacity Scenario	Scenario Prob.	Start of Reduction	Duration of Reduction	Normal AAR	Reduced AAR	End of Reduction	Air/Gnd Cost Ratio	Avg Gnd Delay	Avg Air Delay	Wgt Avg Delay	Runway Utilization	
1	75%	30	30	20	10	60	2	11.50	0.00	3.83	87%	
2	25%	30	60	20	5	90	2	11.50	10.50	10.83	98%	
Combined	100%	n/a	n/a	n/a	n/a	n/a	2	11.50	2.63	5.58	97%	

Details of Scenario 1									
Flight Number	Expected Approach	Ground Delay	Actual Approach	Airborne Delay	Start To Land	Time To Land	Time Off Runway	Total Delay	Contingency Adjustment
1	2.32	0.00	2.32	0.00	2.32	3.00	5.32	0.00	0
2	2.42	2.90	5.32	0.00	5.32	3.00	8.32	2.90	0
3	3.46	4.86	8.32	0.00	8.32	3.00	11.32	4.86	0
4	5.95	5.37	11.32	0.00	11.32	3.00	14.32	5.37	0
5	7.47	6.85	14.32	0.00	14.32	3.00	17.32	6.85	0
6	10.00	7.32	17.32	0.00	17.32	3.00	20.32	7.32	0
7	15.42	4.90	20.32	0.00	20.32	3.00	23.32	4.90	0
8	16.53	6.79	23.32	0.00	23.32	3.00	26.32	6.79	0
9	17.36	8.96	26.32	0.00	26.32	3.00	29.32	8.96	0
10	22.63	6.69	29.32	0.00	29.32	3.00	32.32	6.69	0
11	27.76	4.56	32.32	0.00	32.32	6.00	38.32	4.56	0
12	32.50	5.82	38.32	0.00	38.32	6.00	44.32	5.82	0
13	33.38	10.93	44.32	0.00	44.32	6.00	50.32	10.93	0
14	37.93	12.39	50.32	0.00	50.32	6.00	56.32	12.39	0
15	39.41	16.91	56.32	0.00	56.32	6.00	62.32	16.91	0
16	40.46	21.86	62.32	0.00	62.32	3.00	65.32	21.86	0
17	43.34	21.98	65.32	0.00	65.32	3.00	68.32	21.98	0
18	43.59	24.73	68.32	0.00	68.32	3.00	71.32	24.73	0
19	44.56	26.75	71.32	0.00	71.32	3.00	74.32	26.75	0
20	44.85	29.47	74.32	0.00	74.32	3.00	77.32	29.47	0

Details of Scenario 2									
Flight Number	Expected Approach	Ground Delay	Actual Approach	Airborne Delay	Start To Land	Time To Land	Time Off Runway	Total Delay	Contingency Adjustment
1	2.32	0.00	2.32	0.00	2.32	3.00	5.32	0.00	0
2	2.42	2.90	5.32	0.00	5.32	3.00	8.32	2.90	0
3	3.46	4.86	8.32	0.00	8.32	3.00	11.32	4.86	0
4	5.95	5.37	11.32	0.00	11.32	3.00	14.32	5.37	0
5	7.47	6.85	14.32	0.00	14.32	3.00	17.32	6.85	0
6	10.00	7.32	17.32	0.00	17.32	3.00	20.32	7.32	0
7	15.42	4.90	20.32	0.00	20.32	3.00	23.32	4.90	0
8	16.53	6.79	23.32	0.00	23.32	3.00	26.32	6.79	0
9	17.36	8.96	26.32	0.00	26.32	3.00	29.32	8.96	0
10	22.63	6.69	29.32	0.00	29.32	3.00	32.32	6.69	0
11	27.76	4.56	32.32	0.00	32.32	12.00	44.32	4.56	0
12	32.50	5.82	38.32	6.00	44.32	12.00	56.32	11.82	0
13	33.38	10.93	44.32	12.00	56.32	12.00	68.32	22.93	0
14	37.93	12.39	50.32	18.00	68.32	12.00	80.32	30.39	0
15	39.41	16.91	56.32	24.00	80.32	12.00	92.32	40.91	0
16	40.46	21.86	62.32	30.00	92.32	3.00	95.32	51.86	0
17	43.34	21.98	65.32	30.00	95.32	3.00	98.32	51.98	0
18	43.59	24.73	68.32	30.00	98.32	3.00	101.32	54.73	0
19	44.56	26.75	71.32	30.00	101.32	3.00	104.32	56.75	0
20	44.85	29.47	74.32	30.00	104.32	3.00	107.32	59.47	0

Exhibit 5: Effect of adding extra ground delay to flight 12 when conditions are not favorable

Scenario Definitions							Scenario Results				
Capacity Scenario	Scenario Prob.	Start of Reduction	Duration of Reduction	Normal AAR	Reduced AAR	End of Reduction	Air/Gnd Cost Ratio	Avg Gnd Delay	Avg Air Delay	Wgt Avg Delay	Runway Utilization
1	75%	30	30	20	10	60	2	11.80	1.65	5.03	90%
2	25%	30	60	20	5	90	2	11.80	10.20	10.73	98%
Combined	100%	n/a	n/a	n/a	n/a	n/a	2	11.80	3.79	6.46	92%

Details of Scenario 1									
Flight Number	Expected Approach	Ground Delay	Actual Approach	Airborne Delay	Start To Land	Time To Land	Time Off Runway	Total Delay	Contingency Adjustment
1	2.32	0.00	2.32	0.00	2.32	3.00	5.32	0.00	0
2	2.42	2.90	5.32	0.00	5.32	3.00	8.32	2.90	0
3	3.46	4.86	8.32	0.00	8.32	3.00	11.32	4.86	0
4	5.95	5.37	11.32	0.00	11.32	3.00	14.32	5.37	0
5	7.47	6.85	14.32	0.00	14.32	3.00	17.32	6.85	0
6	10.00	7.32	17.32	0.00	17.32	3.00	20.32	7.32	0
7	15.42	4.90	20.32	0.00	20.32	3.00	23.32	4.90	0
8	16.53	6.79	23.32	0.00	23.32	3.00	26.32	6.79	0
9	17.36	8.96	26.32	0.00	26.32	3.00	29.32	8.96	0
10	22.63	6.69	29.32	0.00	29.32	3.00	32.32	6.69	0
11	27.76	4.56	32.32	0.00	32.32	6.00	38.32	4.56	0
12	32.50	11.82	44.32	0.00	44.32	6.00	50.32	11.82	6
13	33.38	10.93	44.32	6.00	50.32	6.00	56.32	16.93	0
14	37.93	12.39	50.32	6.00	56.32	6.00	62.32	18.39	0
15	39.41	16.91	66.32	6.00	62.32	3.00	65.32	22.91	0
16	40.46	21.86	62.32	3.00	65.32	3.00	68.32	24.86	0
17	43.34	21.98	65.32	3.00	68.32	3.00	71.32	24.98	0
18	43.59	24.73	68.32	3.00	71.32	3.00	74.32	27.73	0
19	44.56	26.75	71.32	3.00	74.32	3.00	77.32	29.75	0
20	44.85	29.47	74.32	3.00	77.32	3.00	80.32	32.47	0

Details of Scenario 2									
Flight Number	Expected Approach	Ground Delay	Actual Approach	Airborne Delay	Start To Land	Time To Land	Time Off Runway	Total Delay	Contingency Adjustment
1	2.32	0.00	2.32	0.00	2.32	3.00	5.32	0.00	0
2	2.42	2.90	5.32	0.00	5.32	3.00	8.32	2.90	0
3	3.46	4.86	8.32	0.00	8.32	3.00	11.32	4.86	0
4	5.95	5.37	11.32	0.00	11.32	3.00	14.32	5.37	0
5	7.47	6.85	14.32	0.00	14.32	3.00	17.32	6.85	0
6	10.00	7.32	17.32	0.00	17.32	3.00	20.32	7.32	0
7	15.42	4.90	20.32	0.00	20.32	3.00	23.32	4.90	0
8	16.53	6.79	23.32	0.00	23.32	3.00	26.32	6.79	0
9	17.36	8.96	26.32	0.00	26.32	3.00	29.32	8.96	0
10	22.63	6.69	29.32	0.00	29.32	3.00	32.32	6.69	0
11	27.76	4.56	32.32	0.00	32.32	12.00	44.32	4.56	0
12	32.50	11.82	44.32	0.00	44.32	12.00	56.32	11.82	6
13	33.38	10.93	44.32	12.00	56.32	12.00	68.32	22.93	0
14	37.93	12.39	50.32	18.00	68.32	12.00	80.32	30.39	0
15	39.41	16.91	56.32	24.00	80.32	12.00	92.32	40.91	0
16	40.46	21.86	62.32	30.00	92.32	3.00	95.32	51.86	0
17	43.34	21.98	65.32	30.00	95.32	3.00	98.32	51.98	0
18	43.59	24.73	68.32	30.00	98.32	3.00	101.32	54.73	0
19	44.56	26.75	71.32	30.00	101.32	3.00	104.32	56.75	0
20	44.85	29.47	74.32	30.00	104.32	3.00	107.32	59.47	0

Exhibit 6: Equivalent performance from full contingency adjustment

No contingency adjustment for flights 12-20:

Scenario Definitions							Scenario Results				
Capacity Scenario	Scenario Prob.	Start of Reduction	Duration of Reduction	Normal AAR	Reduced AAR	End of Reduction	Air:Gnd Cost Ratio	Avg Gnd Delay	Avg Air Delay	Wgt Avg Delay	Runway Utilization
1	75%	30	30	20	10	60	4	11.50	0.00	2.30	97%
2	25%	30	60	20	5	90	4	11.50	10.50	10.70	98%
Combined	100%	n/a	n/a	n/a	n/a	n/a	4	11.50	2.63	4.40	97%

Full contingency adjustment for flights 12-20:

Scenario Definitions							Scenario Results				
Capacity Scenario	Scenario Prob.	Start of Reduction	Duration of Reduction	Normal AAR	Reduced AAR	End of Reduction	Air:Gnd Cost Ratio	Avg Gnd Delay	Avg Air Delay	Wgt Avg Delay	Runway Utilization
1	75%	30	30	20	10	60	4	22.00	0.00	4.40	64%
2	25%	30	60	20	5	90	4	22.00	0.00	4.40	98%
Combined	100%	n/a	n/a	n/a	n/a	n/a	4	22.00	0.00	4.40	73%

Exhibit 7: Analysis of conditions favoring contingency adjustment to ground delay times when second scenario is worse than first

Delays with no contingency adjustment:

	Avg Gnd Delay	Avg Air Delay
Scenario 1	11.50	0.00
Scenario 2	11.50	10.50

	Avg Gnd Delay	Avg Air Delay
Scenario 1	G	0
Scenario 2	G	A

Delays with full contingency adjustment:

	Avg Gnd Delay	Avg Air Delay
Scenario 1	22.00	0.00
Scenario 2	22.00	0.00

	Avg Gnd Delay	Avg Air Delay
Scenario 1	G + A	0
Scenario 2	G + A	0

Exhibit 8: Analysis of conditions favoring contingency adjustment to ground delay times when second scenario is better than first

Delays with no contingency adjustment:

	Avg Gnd Delay	Avg Air Delay
Scenario 1	13.00	0.00
Scenario 2	13.00	0.00

	Avg Gnd Delay	Avg Air Delay
Scenario 1	G + A	0
Scenario 2	G + A	0

Delays with full contingency adjustment:

	Avg Gnd Delay	Avg Air Delay
Scenario 1	12.10	0.90
Scenario 2	12.10	0.00

	Avg Gnd Delay	Avg Air Delay
Scenario 1	G	A
Scenario 2	G	0

Exhibit 9: Cumulative arrivals and departures for 50 flights with two arrival rushes

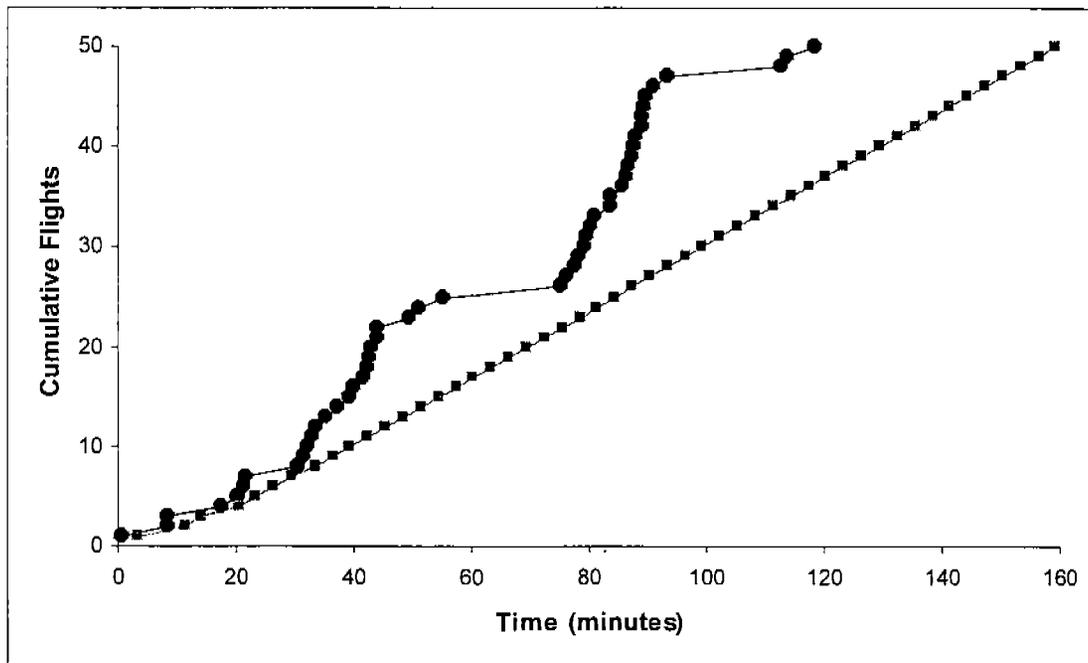


Exhibit 10: Average airborne delay as a function of which flight is cancelled first

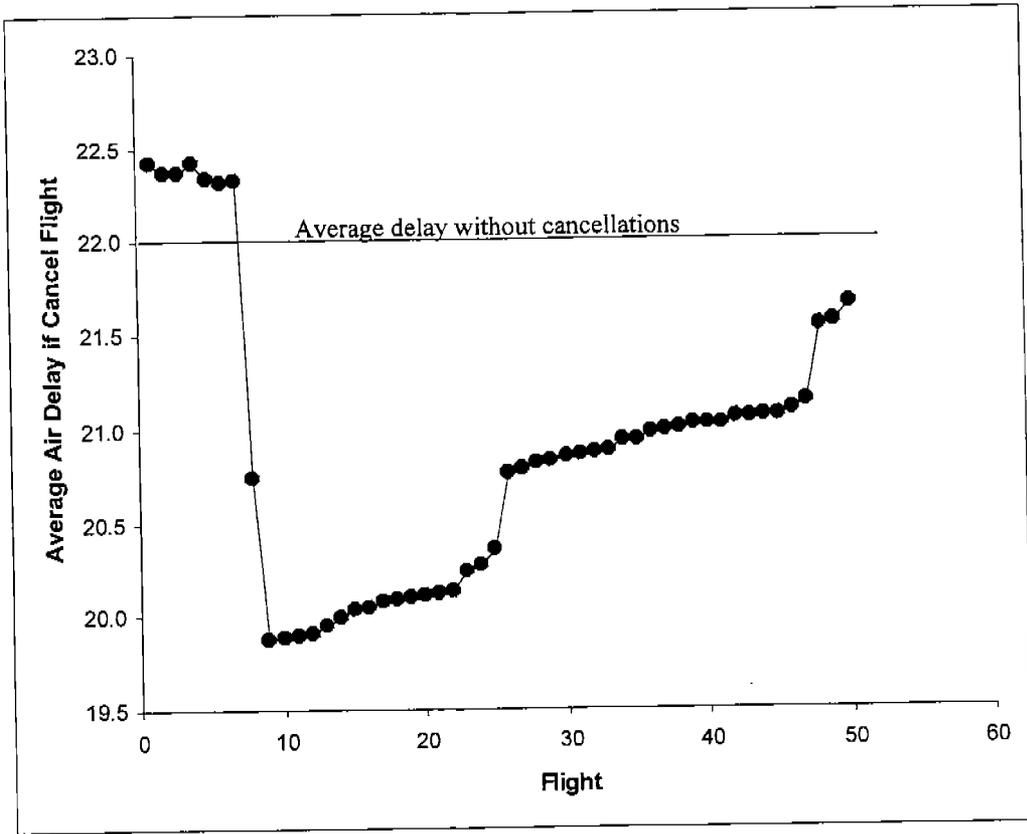


Exhibit 11: Priority for canceling flights

Flight	Scheduled Approach	Cancel Priority	Resulting Air Dly
		none	22.0
1	0.5		
2	8.4	26	0.4
3	8.6		
4	17.6		
5	20.5	28	0.3
6	21.5	24	0.7
7	21.7		
8	30.5		
9	31.5	1	19.9
10	32.2	2	17.8
11	32.9	3	15.6
12	33.4	29	0.2
13	35.3		
14	37.3	12	6.6
15	39.2		
16	40.1	14	5.2
17	41.8	16	3.9
18	42.3	32	0.0
19	42.6	17	3.4
20	43.0	19	2.4
21	43.9	22	1.2
22	44.1		
23	49.4	31	0.1
24	51.0		
25	55.1		
26	75.2		
27	76.1	4	14.4
28	77.6	5	13.3
29	78.4		
30	79.5	6	12.2
31	79.7	7	11.1
32	80.3	8	10.0
33	81.1	27	0.4
34	83.6	9	9.1
35	83.7		
36	85.6	10	8.2
37	86.3	11	7.4
38	86.9	33	0.0
39	87.5	13	5.8
40	87.7	15	4.5
41	88.0	18	2.8
42	89.1	20	1.9
43	89.3	21	1.5
44	89.6	23	0.9
45	89.8		
46	91.3	25	0.5
47	93.5		
48	112.7	30	0.1
49	113.9		
50	118.7		

Exhibit 12: Sequential reduction in average airborne delay

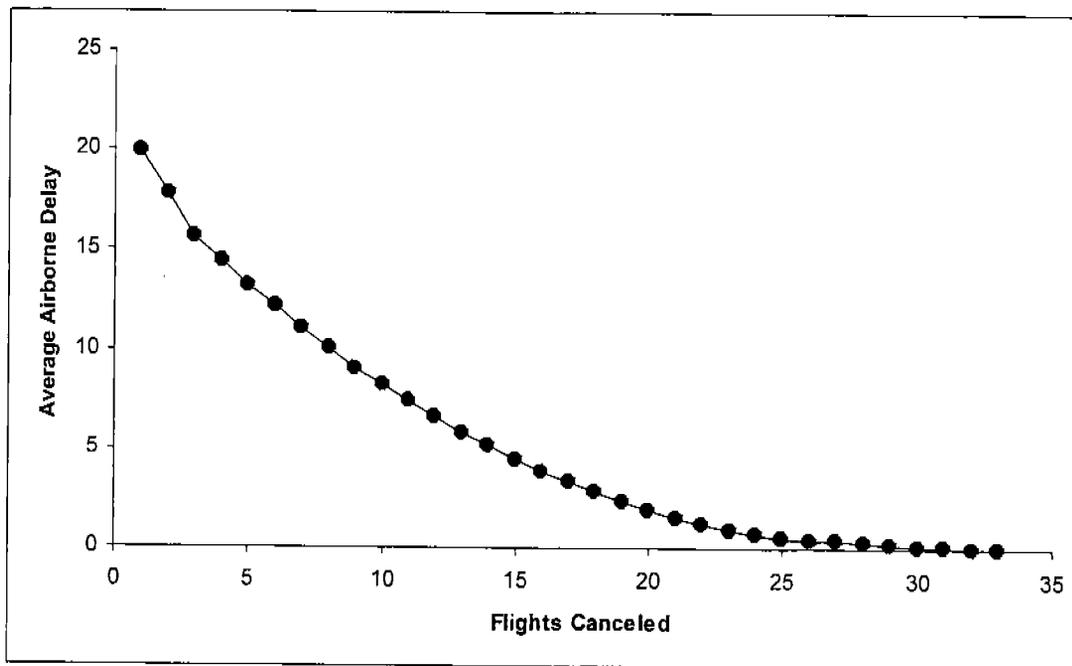


Exhibit 13: Cost of cancellation required to justify a given number of cancellations (2 trials)

Priority to Cancel	Flight to Cancel	Trial 1		Flight to Cancel	Trial 2	
		Resulting Air Delay	Breakeven Time		Resulting Air Delay	Breakeven Time
n/a	none	22.0	n/a	none	21.1	n/a
1	9	19.9	125.0	10	19.0	122.4
2	10	17.8	123.1	12	16.9	119.6
3	11	15.6	121.3	26	15.8	103.9
4	27	14.4	108.8	28	14.6	95.2
5	28	13.3	100.4	29	13.5	89.3
6	30	12.2	94.0	30	12.4	84.8
7	31	11.1	89.0	31	11.3	81.1
8	32	10.0	84.8	32	10.2	77.9
9	34	9.1	80.8	33	9.2	75.0
10	36	8.2	77.1	35	8.3	72.3
11	37	7.4	73.7	36	7.3	69.8
12	14	6.6	70.7	37	6.4	67.4
13	39	5.8	68.0	14	5.6	65.0
14	16	5.2	65.2	38	4.7	63.0
15	40	4.5	62.8	40	4.1	60.5
16	17	3.9	60.4	15	3.5	58.4
17	19	3.4	58.1	41	3.0	56.1
18	41	2.8	56.0	16	2.5	54.0
19	20	2.4	54.0	42	2.1	51.9
20	42	1.9	52.1	18	1.7	50.1
21	43	1.5	50.2	19	1.4	48.2
22	21	1.2	48.4	43	1.1	46.5
23	44	0.9	46.8	20	0.8	44.8
24	6	0.7	45.0	5	0.6	43.2
25	46	0.5	43.4	21	0.5	41.6
26	2	0.4	41.9	9	0.4	40.1
27	33	0.4	40.4	1	0.3	38.8
28	5	0.3	39.0	23	0.2	37.5
29	12	0.2	37.7	44	0.0	36.3
30	48	0.1	36.5	27	0.0	35.1
31	23	0.1	35.4	48	0.0	34.0
32	18	0.0	34.3	17	0.0	32.9
33	38	0.0	33.3			

Exhibit 14: Flight to cancel as a function of AAR

