

Combining Control by CTA and Dynamic Enroute Speed Adjustment to Improve Ground Delay Program Performance

James C. Jones, David J. Lovell

Department of Civil and Environmental Engineering and
Institute for Systems Research
University of Maryland
College Park, USA
jonesjc1@umd.edu, lovell@umd.edu

Michael O. Ball

Robert H. Smith School of Business and Institute for
Systems Research
University of Maryland
College Park, USA
mball@rsmith.umd.edu

Abstract—Over the past several years there have been proposals and discussions regarding a move from the use of controlled times of departure (CTDs) to controlled times of arrival (CTAs) for ground delay programs (GDPs) in the U.S. In this paper we show that, by combining control by CTA with the judicious use of en route speed control, significant improvements to GDP performance can be achieved. Our analysis of this problem includes both new GDP control procedures and also new flight operator GDP planning models. While the ability to achieve all the benefits we describe will require NextGen capabilities, substantial performance improvements could be obtained even with a near-term implementation.

Keywords- *Speed Control, Collaborative Decision Making, Traffic Management Initiatives, Integer Programming*

I. INTRODUCTION

Airports throughout the U.S. National Airspace System (NAS) regularly encounter elevated levels of congestion. When faced with inclement weather, the capacity of airports is often insufficient to accommodate demand and the potential effects on congestion are particularly burdensome. These capacity/demand imbalances will often lead the Federal Aviation Administration (FAA) air traffic managers to impose Ground Delay Programs (GDPs) at airports or Airspace Flow Programs (AFPs) in the airspace. Such initiatives reduce the rate of arriving flights into the airport or region of airspace to a level more compatible with the resource's reduced capacity. Specifically, flights receive later departure times from their origin airports commensurate with the rate of arrival desired by the affected airport or airspace region. These programs lower flight operator operational costs and reduce the workload imposed on air traffic controllers by mitigating the potential resulting airspace congestion.

Traffic Flow Management Initiatives (TMIs) in the U.S., such as GDPs and AFPs, operate under a philosophy known as Collaborative Decision Making (CDM). In this paradigm, decisions on flight release times are made by airlines and air navigation service providers (ANSPs) in a joint decision process. CDM processes involve both information sharing and specialized resource allocation mechanisms. [1,2].

GDP planning procedures universally start by considering constraints on airport arrival rates and set corresponding arrival flows, which generally are converted to specific flight arrival

times / CTAs. These in turn are converted to CTDs by subtracting an estimated flight time from the CTA. The reason this is done is that a CTD is much easier to monitor and implement than a CTA. Airport tower controllers can monitor flight departure times and prevent controlled flights from departing early. On the other hand, it can be difficult, costly and even dangerous to insure a set of airborne flights meet CTA constraints. Nonetheless, for many years CDM developers and researchers have expressed the goal of controlling based only on CTAs. This is attractive because it gives flight operators more flexibility and allows both for more dynamic flight planning and a greater ability to do system-wide tradeoffs. Even though control by CTA was an expressed goal early in CDM development, it has not yet been implemented. The reasons are perhaps a combination the implementation challenges and the lack of credible quantification of the benefits.

As both U.S. and worldwide air traffic flow management moves toward time-based metering and eventually trajectory based operations (TBO), the use of CTAs for a variety of goals should become routine. Thus, we argue that the implementation challenges just mentioned should decrease and eventually go away. The experimental results in this paper provide at least a first step in overcoming the second obstacle by demonstrating significant benefits.

CDM procedures were first described in the literature by Wambsganss and Chang et al. [1,3]. A body of research has developed around both GDP planning and CDM. One research stream has modeled GDP uncertainty and incorporated stochastic components into integer programming models [4,5,6] for GDP planning. GDP decision support tools now mitigate the impact of weather uncertainty by exempting flights whose origins are outside a computed radius [7]. The theory underlying this process and extensions are provided in [8]. Other work has formalized and then extended CDM resource allocation processes [9]; related work extended the compression algorithm to include a more comprehensive slot exchange process [10]. Vossen et al. [11] developed metrics of equity and then showed how CDM procedures could be improved to deliver a better level of equity. CDM was later extended for use in a multi-resource context by Fearing and Barnhart [12]. This philosophy has also been proposed in departure queuing by Briton et al. and Bhadra et al. [13,14].

Speed control has been widely studied for a variety of air traffic management applications. At the tactical level Neuman and Erzberger [15] described a variety of sequencing and spacing algorithms designed to reduce fuel consumption and en route/arrival delay. These algorithms laid the foundation for the Traffic Management Advisor (TMA) system currently used at many airports across the country to manage flights up to 200 nmi from the airport. An enhanced version of the system called The Terminal Area Precision Scheduling and Spacing System (TAPSS) was later developed [16]. The technology was also proposed for cooperative use in Traffic Flow Programs [17]. Carrier-centric approaches such as The Airline Based En Route Sequencing and Spacing tool have also been proposed. The tool sends speed advisories to the Airline Operations Centers (AOC)s to allow crews to more actively manage their speeds en route [18].

In recent years, the horizon for such air traffic management initiatives has also moved farther away from the airport. Airservices Australia developed the ATM Long Range Optimal Flow Tool (ALOFT) to allow pilots to control speeds up to 1000 nmi away from the airport. In so doing, they achieved an estimated fuel savings of nearly 1 million kg in 2008 [19]. Since then, they have also used additional metering fixes to better manage trajectory and arrival time uncertainty [20]. Delta Airlines achieved an estimated \$8 million in fuel savings over a 20-month period using a dispatch monitored speed control program known as Attila[21]. At Schiphol, a ground based planning system that interfaced with aircraft through datalink was used to remove vectoring in their nighttime operations [22]. Knorr et al. [23] identified substantial inefficiencies in the terminal phase of flights and characterized the benefit pool achieved by “transferring” terminal delays to the en route phase of flight. Jones et al. [24] developed a bi-criteria integer programming model to facilitate delay transfer away from terminal airspace and demonstrated that a substantial proportion of the potential delay transfer benefit could be realized through this approach.

Speed control measures have also been proposed for capacity allocation in GDPs. Delgado and Prats showed that it was possible to absorb some of the delay assigned to flights within the GDP en route and maintain the planned level of fuel consumption [25, 26, 27]. The authors also showed that by flying earlier and at a slower speed, a considerable portion of the imposed delay could be recovered in the event of an early GDP cancellation. Jones and Lovell showed that speed control could also be used to help curb the exemption bias in GDP slot assignments [28].

In this paper, we consider replacing the use of a CTD with a CTA in GDP planning and control. The principal change is conceptually quite simple: flights and, by association, flight operators, are assigned CTAs rather than CTDs. When a GDP is revised, the assigned CTAs rather than the assigned CTDs are adjusted. Because of added flexibility provided by the use of CTAs, we also propose the elimination of GDP flight exemptions, instead allowing flight operators to effectively make exemption decisions regarding their own flights. To effect these changes we only need to make minor changes to the existing CDM/GDP allocation procedures. We propose a

new flight operator GDP planning model, specifically a scenario-based stochastic integer programming model that determines a cancellation and substitution plan for each carrier. The model matches the carrier’s flights to the assigned arrival capacity (CTAs). In doing this, it takes into account the ability to adjust flight speeds en route, e.g. the model might assign a flight an “early” departure time, consistent with a relatively slow speed but anticipate the ability for the flight to increase its speed should the weather clear at the destination and additional capacity be assigned to the carrier. The integer programming models builds on the prior literature on stochastic models for GDP planning and the use of speed control extends the work of Delgado and Prats. In Section II we provide a description of the CDM GDP process and our modifications to it. We also provide background on aircraft fuel burn characteristics and how they impact speed control processes. In section III we present our models along with our methodological assumptions. In Section IV we apply our models to a case study based on data obtained at Atlanta Hartsfield-Jackson Airport and demonstrate the ability of our models to improve flight operator performance metrics.

II. BACKGROUND

A. CDM Assignment Practices

The CDM resource allocation mechanism for GDP planning consists of three components: capacity allocation, schedule adjustments, and slot exchange. As discussed above, while control is executed based on a CTD, planning is done based on a CTA. Specifically, arrival capacity is allocated to carriers using a mechanism known as Ration-by-Schedule (RBS). In RBS, flights are assigned to arrival slots based on the order they appeared in the original schedule. This procedure provides an equitable initial allocation and removes incentives for carriers to report inaccurate information. Once capacity has been assigned, schedule adjustments are typically performed by allowing airlines to cancel and substitute flights based on their own priorities. To improve the overall footprint, an inter-airline substitution procedure known as compression is used facilitate trades. A notional diagram of the process is shown in Figure 1.




Figure 1: Flight assignment under a CDM framework

When a GDP is issued at an airport, air traffic managers at the Air Traffic Control System Command Center (ATCSCC) decide the planned capacity and duration of the GDP based on the predicted conditions over the course of the day. They also determine the radius of exemption for the GDP. This exemption radius defines the set of flights that will receive ground delays. Once this parameter has been determined there are three pools of flights that are affected. Flights inside the exemption radius receive ground delays based on their order in the schedule. Flights on the ground outside of the radius are exempted from the GDP and receive no delays. In addition, all

flights already in the air, regardless of their origin, are exempted from the GDP.

Figure 2a illustrates an example RBS allocation where the two exempt flights identified on the left are both airborne at the time of allocation. After the RBS allocation, carriers may freely substitute flights based on their own priorities. They may also choose to cancel flights and make substitutions using the vacated slots. A notional example of this process is shown in Figure 2b. Here, AA has chosen to cancel AA561 and move AA321 into its slot. AA alternatively could have chosen to swap the slots of the two flights. In either case, once the appropriate arrival changes were made, the arrival times (CTAs) would be converted to departure times (CTDs) and appropriate ground delays. Although DAL and UA both also have two slots, they are unable to make any changes since in each case one of their two slots is occupied by an airborne flight whose arrival time cannot be adjusted. If DAL and UA could reassign the slots of their airborne flights, then each airline could improve the number of flights arriving less than 15 minutes after their scheduled arrival time. An example of this exchange is shown in Figure 2c.




Figure 2a: An example of flight allocation in Distance Based Ration-by-Schedule. Exempt flights receive priority.

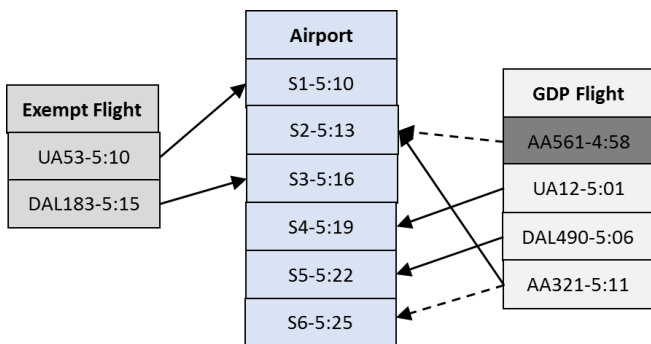


Figure 2b: Cancellation and Substitution process in current CDM framework.




Figure 2c: Cancellation and Substitution process without exemptions. Delta and United can substitute and improve their on-time performance.

The advantages illustrated by these examples underlie the one component of the benefits that can be achieved by combining control by CTA with dynamic speed adjustments. Similar improvements to the performance of the compression algorithm can be achieved by allowing adjustments to airborne flights. We note a second source of benefits have the same origin as those investigated by Delgado and Prats [25,27], namely the ability of airborne flights to more quickly react to increases in arrival capacity resulting from weather changes.

B. Fuel Burn Implications of Speed Changes

The relationship between fuel efficiency (specific range) and Mach number is illustrated in Figure 4. In this relationship, as the Mach number of the aircraft increases, its fuel efficiency will also increase to a point known as the maximum range, beyond which it begins to decline. In general, the shape of this curve in the vicinity of the optimum is relatively flat. The flatness of this curve implies that the speed of the aircraft can be adjusted (within a reasonable range) to absorb the intended delay with a minimal increase in fuel burn.

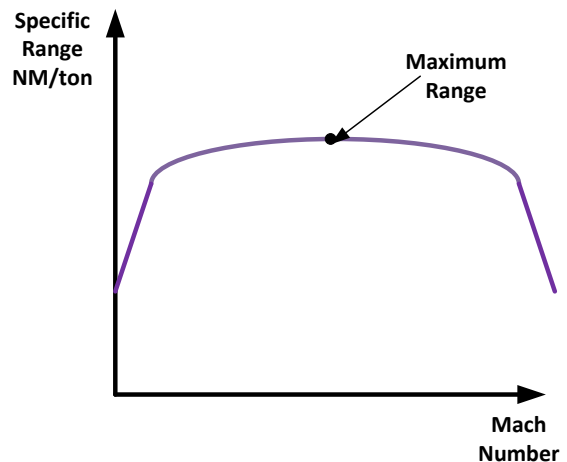


Figure 4: Variation in Aircraft fuel efficiency with speed.

In [25] it was suggested that GDP flights receive an early release time and fly at the minimum possible speed such that

their specific range could be maintained. In the event that the weather clears prior to the end of the GDP, the flight can get in the air earlier and can fly faster to recover some of the delay. In [27] the same authors examine the effect of the exemption radius on GDP performance when applying their original idea. In [28] speed control was applied to the exempt flights while flights inside the radius received only ground delays. Each of these studies examined trades between delaying flights in the air and delaying flights on the ground but in all cases the assignments were made by the ANSP. They serve as a significant first step, however, as these trade-offs may be more effective when they are performed by the carriers. If a flight on the ground receives a 15-minute reduction in delay at the expense of a flight in the air, the benefit to that flight is irrelevant if the carrier decides to cancel the flight on the ground. For reasons such as this we would like to propose a more carrier-centric alternative to these strategies. In the following section we present a model designed to allow carriers to control their own flights using a combination of ground delay, speed control and cancellations. We describe the framework for the models and context in which we envision the assignment process operating.

III. NEW MODELS AND ARCHITECTURE

In this section, we describe our CTA-based architecture (III-A), present our new airline optimization model to support the airline cancellation and substitution process (III-B) and describe the model used to represent compression and revisions in our experiments (III-C).

A. Architecture

The previous two sections have described in general terms the modifications that we envision to major components of the process. Here we more specifically define the architecture and explain some of the important changes. While the new process uses the RBS mechanism, the exemption radius is eliminated. Once capacity is allocated to carriers, each carrier can use both speed control and ground delays to manage their substitution and cancellation decisions. Since no exemption radius has been imposed, carriers must be more strategic about their substitution process because in the event of an early weather clearance they will want to take advantage of capacity increases, e.g. by speeding up airborne flights. Our scenario based stochastic model (Section III-B) is designed to facilitate that end. Each scenario accounts for a possibility of the weather clearing at different times and the associated increase in capacity. The goal is to position the flights in slots that allow them to make the best use of capacity under all scenarios. The fact that slot assignment and the effective use of speed control are key to evaluating the impact of this new approach implies that proper evaluation of its effectiveness requires experiments that involve GDP revisions. The manner in which we model revisions is discussed both later in this section and in Section III-C.

TABLE I gives the basic steps in our CTA-based architecture.

TABLE I: CTA-Based Flight Assignment Architecture

Step 1 [FAA].

1a: Assign a slot to each airborne flight based on the flight's expected time of arrival.

1b: Assign a slot to all flights on the ground using RBS.

1c: Create a list of slots (and CTAs) owned by each airline based on the allocation from both steps 1a and 1b.

Step 2 [Airlines]. Execute cancellation and substitution processes and adjust flight-to-CTA assignments. Assign a departure time to each flight

Step 3 [FAA]. Execute compression, adjusting assignments and filling any unusable slots.

This process looks almost identical to the existing process illustrated in Figure 1. However, there are some subtle differences. First, none of the flights on the ground at the start of the GDP are exempted. Second, when the airlines perform their cancellations and substitutions, and also when the FAA performs compression (steps 2 and 3), both airborne flights and flights on the ground should be considered. The consideration of flights in the air imposes a substantial new information requirement: the (possibly very tight) limits on the degree to which their arrival times can be adjusted. Third, today the assignment of a departure time (CTD) is performed by subtracting a nominal flight time from the CTA. Under this new approach the airlines have substantial flexibility in assigning the departure time, e.g. as in Delgado and Prats [25,27], assuming an initial "slow" speed while anticipating possible speed-ups if weather conditions change. This added airline flexibility implies that when the airlines perform their cancellation and substitution process, they have a rich set of alternatives to consider and the opportunity to improve performance. In the next section, we present an optimization model to address this new airline decision problem.

Another very important challenge associated with this new approach is the manner in which GDP controls are dynamically updated over time. Today a variety of possible GDP revisions might take place as weather conditions change at the destination airport. Perhaps the simplest is a cancellation of the GDP in the event of clearance of poor weather. If this occurs, all issued ground delays are immediately rescinded and the impacted flights can immediately take off. An equivalent action in a CTA-based architecture would be to allow flights on the ground to immediately depart and flights in the air increase their speed, to the extent feasible, in order to arrive at an earlier time if this is desired. It is difficult to assess a priori whether such a complete cancellation might ever be appropriate under a CTA-based system. However, it is clear that new GDP revision models and controls will be required. In particular, it is likely that "revisions" will be required not only based on major changes in conditions at the destination airport but also more minor disturbances that impact the flight times of en route flights. It is likely that such models could build on the recent experience with airborne speed control [17,18,19,20,21,22] and the growing body of research [23,24,25,26,27,28] on the topic. Of course, this also relates to current efforts on time-based metering and TBO.

In the current research we have not attempted to address all the nuances of GDP revisions under CTA controls. This would certainly represent another significant research contribution. Rather, to estimate the benefits of this new architecture, we evaluate a relatively simple scenario in which weather clears at a random time and use an optimization model [9] that represents the combined effect of RBS and compression in reassigning CTAs based on the newly available capacity. This model is described in Section III-C.

B. Model to Support Airline Substitution and Cancellation Process

Under the new architecture and considering both the possibility of en route speed adjustments and no flight exemptions, each airline has more control over the disposition of its own flights. Since GDPs are often cancelled prior to their planned end time, it behooves airlines to hedge between the prospect of early and on-time cancellation. Such hedging is effectively done today by the FAA through the exemption radius. The challenge for an airline lies in positioning flights in the appropriate slots to best deal with all possibilities. To do so we adapt stochastic models developed earlier [4,5,8] from an FAA/ANSP perspective to the perspective of a specific airline.

To understand this model, consider the deterministic case where the set of available slots, i.e. the CTAs assigned to that airline, is known with certainty, e.g. as described in [10]. This is a simple assignment problem where flights are assigned to slots allowing for the possibility that some flights may be canceled at a cost. Since this model is solved by a specific airline, we can assume the availability of a rich cost function that takes into account various factors regarding flight, crew and passenger status, passenger count, etc.

Capacity uncertainty is modeled using a set of scenarios: each scenario is characterized by the time at which that scenario becomes known, the revised set of slots, i.e. additional capacity represented by the augmentation of the existing slots with a set of additional slots, and a probability. An additional set of variables indicates how the initial assignment is adjusted when the new capacity becomes available. In defining the data underlying this model, the differences in constraints underlying airborne flights and flights on the ground must be taken into account. For example, if a flight was assigned a CTA of 4:00 and, at the time the new scenario was effective that flight was airborne, then the flight might be restricted to revised CTAs not earlier than 3:50 based on limitations on speedup options (no more than 10 minutes). On the other hand, if a flight on the ground was assigned a CTA of 4:00 and that flight still had one hour to serve on its ground delay, then that flight could be assigned any departure time within the next hour and in order to meet any new assigned CTA between 3:00 and 4:00. The air carriers should assign both a CTA and departure time to each flight. For the present experiments, we assume the departure time assigned is the earliest possible departure time that can meet the assigned CTA. This approach provides maximum flexibility where weather scenarios only allow for capacity increases. We recognize that ideally the optimization model should contain both departure time and arrival time variables – we leave such a model to future research.

This model certainly has some substantial data requirements, most notably the scenario information. There are two aspects to generating the slot lists for each scenario. The first is defining the set of *slots available to all carriers* and the second is *how those slots are assigned to each carrier*. There has been prior research on the first aspect, but this certainly would have to be adapted to this new context. For the purposes of this paper, we use representative/stylized information that captures the essential aspects of the problem setting. Regarding the second aspect, we use a basic RBS reallocation that (by necessity) cannot take into account the status (and slot assignment) of each carrier's flights. Thus, this reallocation must be viewed as an approximation; however, it only impacts cost assigned to the initial slot assignment and so it impacts only the quality of the solution and not its feasibility. We can judge the overall quality of our approach by the results of our simulation experiments. We also note that some air carriers might wish to use other processes; thus, this model could be viewed as a surrogate for any number of internal airline decision support processes.

The specific integer programming problem formulation is given below. Note that this model includes a subscript for airlines – in practice, each airline will solve its own model.

Parameters:

$F_a \equiv$ The set of all flights available to airline a

$A \equiv$ The set of all airlines

$S_a \equiv$ The set of all slots available to airline a

$E_{fa} \equiv$ The set of all slots available to flight f at stage 1 prior to first probable end of the GDP available to airline a

$P_{fa} \equiv$ The set of all slots available to flight f at stage 1 following the first probable end of the GDP available to airline a

$K_{fsq} \equiv$ The set of all slots available to flight f at stage 2 in scenario q from slot s available to airline a

$d_{jsa}^q \equiv$ Cost of delaying flight f to slot s owned by airlines a in scenario q

$c_{fa} \equiv$ The cost of cancelling flight f operated by airline a

$p_q \equiv$ The probability of scenario q occurring

Variables:

$$x_{fsa} = \begin{cases} 1 & \text{if flight } f \text{ of airline } a \text{ is assigned to slot } s \\ 0 & \text{otherwise} \end{cases}$$

$$y_{fa} = \begin{cases} 1 & \text{if flight } f \text{ of airline } a \text{ is cancelled} \\ 0 & \text{otherwise} \end{cases}$$

$$z_{jsa}^q = \begin{cases} 1 & \text{if flight } f \text{ of airline } a \text{ is assigned to slot } s \\ & \text{in scenario } q \\ 0 & \text{otherwise} \end{cases}$$

$$\min \sum_{q \in Q} P_q \sum_{f \in F_a, s \in S_a} d_{fsa}^q z_{fsa}^q + \sum_{f \in F_a} c_{fa} y_{fa} \quad (1)$$

$$\text{s.t.} \sum_{s \in S_a} x_{fsa} + y_{fa} = 1 \quad \forall f \in F_a \quad (2)$$

$$\sum_{f \in F_a} x_{fsa} \leq 1 \quad \forall s \in S_a \quad (3)$$

$$\sum_{f \in F_a} z_{fsa}^q \leq 1 \quad \forall s \in S_a, \forall s \in E_{fa}, \forall q \in Q \quad (4)$$

$$x_{fsa} \leq z_{fsa}^q \quad \forall f \in F_a, \forall s \in E_{fa}, \forall q \in Q \quad (5)$$

$$x_{fsa} \leq \sum_{k \in K_{fsqa}} z_{fka}^q \quad \forall f \in F_a, \forall s \in P_{fa}, \forall q \in Q \quad (6)$$

$$x_{fsa}, z_{fsa}^q, y_{fa} \in \{0,1\} \quad \forall f \in F_a, \forall a \in A, \forall s \in S_a \quad (7)$$

Constraint (2) ensures that for each airline, every flight is either assigned to a slot or cancelled. Constraint (3) ensures that no more than one flight is assigned to a slot in the first stage of the problem. Constraint (4) ensures that no more than one flight is assigned to a slot in the second stage of the problem for all scenarios. Constraint (5) ensures that if a flight is assigned to a slot in the first stage prior to the first feasible weather clearance time, it must be assigned to the same slot in the second stage for all scenarios. Constraint (6) ensures that if a flight is assigned to a slot in the first stage after the first feasible weather clearance time, it must be assigned to a slot that is reachable from that slot in the second stage. Note that this constraint, through the definition of K_{fsqa} , restricts the set of slots to which a flight can be reassigned based on flight status and the various timing restrictions. Constraint (7) reflects that our assignment variables are binary. Our objective is to minimize the expected cost of the flight delays over all scenarios plus the cost of flight cancellations.

C. Compression and GDP Revisions

To carry out our experiments, we must both execute compression as part of the initial allocation process (see Table I) and also perform a slot reallocation for the case of a GDP revision. Under current practices revisions are performed using a modified application of RBS that takes into account both flight status and the new set of available slots. Compression is also typically performed. Very often a combined RBS/compression process is executed called RBS⁺⁺. In [10], an optimization model is defined that provides both the functionality of compression and RBS⁺⁺. We use this model in our experiments for both the initial compression step and also the revision process. This model actually provides carriers with more flexibility in the application of compression. However, for our purposes here we only wish to mimic the basic processes. Specifically, the model employs a set of ‘‘goal slots,’’ with one such slot assigned to each flight to be assigned. To mimic compression the goal slot assigned to each flight is the RBS slot for that flight. Other assignments can be used by carriers to implement various flight prioritization schemes.

Parameters:

$F \equiv$ The set of all flights

$A \equiv$ The set of all airlines

$S \equiv$ The set of all slots

$T \equiv$ The set of all time periods

$F_a \equiv$ The set of all flights belonging to airline a

$I_a \equiv$ The set of all goal slots belonging to airline a

$R_f \equiv$ The set of slots within acceptable for flight f

$S_{ft} \equiv$ The set of all slots available to flight f in period t

$t_s \equiv$ The time corresponding to slot s

$\tau_f \equiv$ The time corresponding to the goal slot of flight f

Variables

$$x_{fsat} = \begin{cases} 1 & \text{if flight } f \text{ of airline } a \text{ is assigned to slot } s \\ & \text{in time period } t \\ 0 & \text{otherwise} \end{cases}$$

$$\min \sum_{\substack{f \in F_a, s \in S_a, \\ a \in A, t \in T}} (\tau_f - t_s)^{1+\epsilon} x_{fsat} \quad (8)$$

$$\text{s.t.} \sum_{s \in S_a, t \in T} x_{fsat} = 1 \quad \forall f \in F_a, \forall a \in A \quad (9)$$

$$\sum_{\substack{f \in F_a, a \in A, \\ t \in T}} x_{fsat} \leq 1 \quad \forall s \in S \quad (10)$$

$$x_{fsat} \in \{0,1\} \quad \forall f \in F_a, \forall a \in A, \forall s \in S, \forall t \in T \quad (11)$$

Equation (9) ensures that every flight is assigned to exactly one slot. Constraint (10) ensures that each slot is assigned to no more than one flight. Constraint (11) reflects that our assignment variables are binary. The objective of the model is to lexicographically minimize the distance of the flights from their goal slots. It accomplishes this minimization by using coefficients that grow super-linearly. This mimics the impact of compression, which seeks to find a slot as close as possible to the flight’s RBS slot in the case where that flight cannot be feasibly assigned to its RBS slot.

IV. RESULTS AND DISCUSSION

A. Experimental Description

To conduct our studies we selected data collected from Atlanta Hartsfield-Jackson Airport on May 1, 2011. The weather conditions were clear and sunny and all runways were active. The data were obtained from a file generated by TFMS in conjunction with an ASDX file, the combination of which listed flight numbers, carrier, collection time, ETA, scheduled time of arrival (STA), the origin airport, actual time of departure, aircraft position, aircraft type, arrival time.

The airport acceptance rates on an hour-by-hour basis varied from 56 to 101 flights per hour. Since this dataset was not taken on a day on which a GDP was issued, a hypothetical GDP was superimposed on the data. A 5 hour GDP was assigned to the airport over the hours of 16:00-21:00 GMT. Flights inside the exemption radius were assigned ground delays. Flights on the ground that originated from airports outside of the radius as well as flights in the cruise phase of flight at the start of the GDP were allocated slots over the range achievable by the aircraft. The model used flight trajectories observed in the data over the day of operations. Speed control directives were issued over the period of time that the aircraft reached an altitude of 35,000 ft. Based on these trajectories we calculated the distance traveled. As a baseline case flights were given a nominal cruise speed based

on the aircraft performance listed on the BADA database. This database was also used to derive a set of speeds at which each aircraft could fly. In general we used these speeds as guidelines; however, speeds on all aircraft were restricted to ± 0.02 of their performance maximum/minimum. Also, when aircraft were capable of flying above Mach 0.85 or below 0.72, aircraft speeds were restricted to a maximum of 0.85 or a minimum 0.72 respectively. CTAs for ground delayed flights could correspond to any time at or following the scheduled time of arrival of the flight.

A baseline run was used to evaluate the delay performance with no intervention. On these runs capacity was allocated to airlines using DB-RBS. A deterministic version of the substitution and cancellation model was used which did not account for the possibility of early clearance. A compression model was then adopted to improve throughput. To understand the full extent of the performance we tested the computation run time of each model using a dual core system with four Intel Xeon X5535 processors and 12 GB of memory in a 64 bit environment. The models were coded in Python 2.7 using a GUROBI solver.

B. The Cost of Delay and Cancellation

If this proposed scheme were implemented, each airline would compute the cost of delay based on their internal cost measures; however, to perform a computational experiment we needed to find a suitable proxy. In this paper we chose to start with the cost model presented in Vakili and Ball [29], which draws from ATA data and models from Metron Aviation. The model assumes that the direct operating cost per minute of block time is free during the first 15 minutes. After 15 minutes the cost jumps to \$64 in the air and \$32 on the ground. Since our airborne delay is essentially free from a fuel cost standpoint and fuel typically accounts for roughly half the delay costs we decided to use an equal cost for ground and air delay. Updating for yearly changes in delay costs we found the cost on both the ground and the air was \$40[30]. The Vakili and Ball approach also assumes that the cost per minute of passenger delay is \$34.88 per hour or $\$34.88/60 = \0.5813 per minute. Since the airlines do not suffer the same degree of impact as customers on a per minute basis the approach approximates the cost by multiplying passenger cost by 1/6 and uses a cost of \$0.1 per minute. Adopting the same process using 2013 passenger costs we find that the additional airline cost is \$0.125 per minute, per passenger. An expression for the cost function is show below:

$$C(x, P) = \begin{cases} 0 & x < 15 & (12) \\ (40 + 0.125P)(x - 15) & 15 \geq x \geq M_p & (13) \\ (40 + 0.125P)(M_p - 15) & x > M_p & (14) \end{cases}$$

where P is the number of passengers on the flight and M_p is the maximum amount of time before the delay cost levels off. When the cost levels off it does not matter whether the airline delays the flight an additional minute or a day. Thus we assume the cost to cancel a flight is the cost at level off. Aircraft specifications were used to determine the number of passengers on a given aircraft. Using the 2013 average reported in IATA our analysis assumed a load factor 0.8 on all flights [31]. We set M_p to a value of 90.

C. Effect on Airlines Metrics with no GDP cancellation

To better study the distribution of delay we reduced the number of flights to consider just the 5 largest carriers. A baseline run was performed using a conventional GDP procedure. Capacity was allocated with DB-RBS and cancellations and substitutions were made using a deterministic model. The resulting performance for a GDP with a Planned Airport Arrival Rate (PAAR) of 40 is shown in TABLE II. The percentage of cancellations remained relatively consistent across carriers ranging between 25% and 33.33%. Delta and AirTran, however, both exhibit stronger delay performance in traditional GDPs. This is understandable as Delta and AirTran both control a larger pool of exempt flights than regional carriers and those with a smaller presence at the airport.

To evaluate CTA-based architecture and new planning modes, we used 5 capacity profiles. The set consisted of a complete GDP and weather clearances of 15, 30, 45 and 60 minutes early. Each scenario was assumed to be equally probable. The results of our test are shown in TABLE III. The tests yielded noticeably different results relative to the baseline. All carriers reduced their number of cancellations except for American Airlines, which only controlled 3 flights.

TABLE II: AIRLINE PERFORMANCE WITH A CONVENTIONAL GDP MODEL

Airline	Percentage of Flights Cancelled	Passenger Delay	Number of Flights
Delta (DAL)	27.78	11.14	108
AirTran (TRS)	32.14	12.37	28
American Southwest Airlines (ASQ)	27.59	18.29	58
American (AAL)	33.33	46.50	3
Pinnacle (FLG)	25.00	29.25	4

TABLE III: AIRLINE PERFORMANCE WITH CTA-BASED ARCHITECTURE

Airline	Percentage of Flights Cancelled	Passenger Delay	Number of Flights
Delta (DAL)	24.07	25.40	108
AirTran (TRS)	28.57	18.02	28
American Southwest Airlines (ASQ)	17.24	19.24	58
American (AAL)	33.33	25.00	3
Pinnacle (FLG)	0	37.75	4

The performance data suggests that airlines will approach the two GDP procedures in remarkably different fashion. In the current framework carriers are more likely to cancel flights to create additional capacity and flexibility as well as reduce delay. In our modification carriers have more opportunity for intra-airline substitution both through speed control and the lack of an exemption radius and are also accounting for the benefits achieved in the event of an early cancellation. This is not something that is assumed in the deterministic planning case. Thus carriers will choose to keep a greater portion of their slots. Since there are far fewer cancellations the carriers are less affected by actions of other carriers during compression. This allows carriers to have more direct control over their performance.

While the example above reveals some information regarding the relative effect of our CDM modification, it does not provide us with a sense of how strong the possibility of early clearance needs to be to affect the decision. We ran the model with another set of scenarios in which the early clearance intervals were only 7.5 minutes apiece. The resulting performance of both models is shown in Figures 7a-b. In nearly all cases the prospect of an early clearance reduced the number of cancellations while increasing the passenger delay carriers with more long haul flights. The magnitude of the reduction is not quite as prominent, however, as that of the 15 min scenarios.

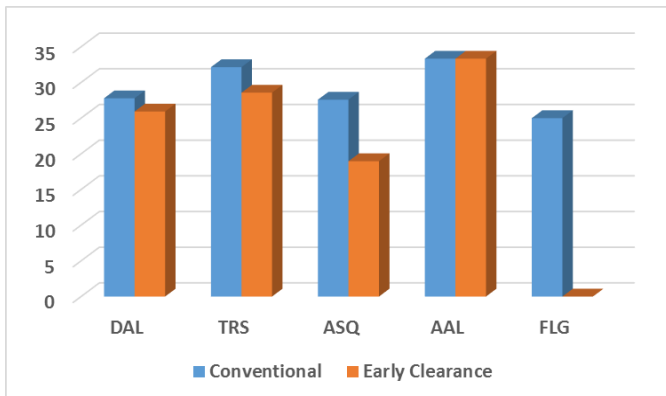


Figure 7a: Percentage Flight cancellation level of Airlines with Conventional and Early Clearance.




Figure 7b: Passenger Delay of Airlines for Conventional and Early Clearance Models in minutes.

While the previous graphs demonstrate significantly different behavior on the part of airlines, it is unclear what portion of the change is attributable to the possibility of early cancellation vs. the lack of an exemption radius. To isolate the effect we tested our models both with and without a radius. In the former case RBS was used to generate capacity while the later used the DB-RBS algorithm. The performance is shown in Figure 8. The results suggest that when a radius is present large carriers such as Delta will reduce the number of cancellations they impose on their flights; this is also the case with Air Tran. This likely attributable to the significantly larger number of exempt eligible flights they have relative to other carriers. Regional carriers such as American Southeast Airlines are negatively affected by the presence of the radius and are forced to cancel more flights to create substitution opportunities.

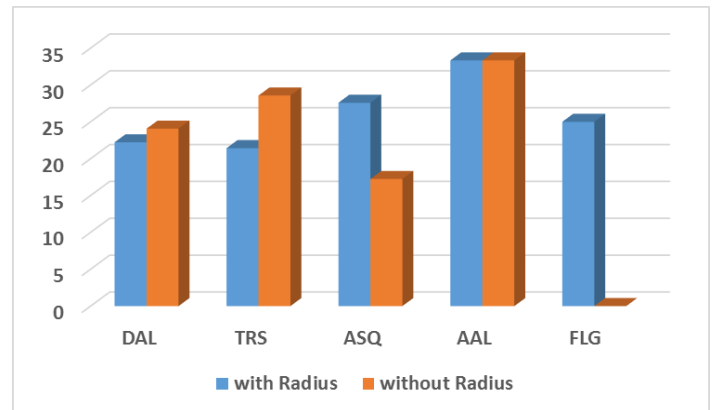


Figure 8a: Effect of the Exemption Radius on Percentage of Flight Cancellations.

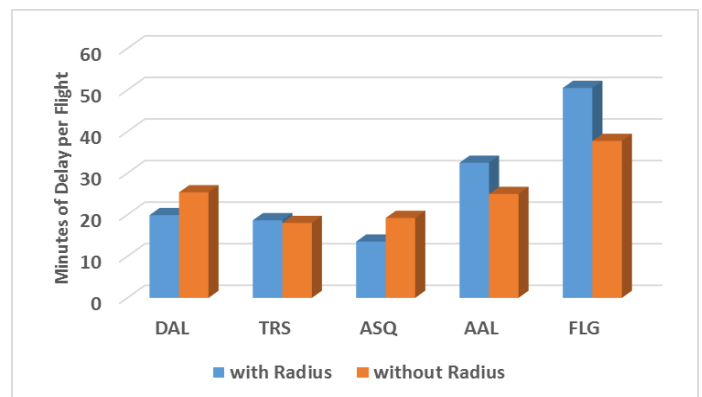


Figure 8b: Effect of the Exemption Radius on Passenger Delay

D. Delay Recovery with GDP cancellation

In addition to the effect our GDP modifications have on the propensity of airlines to cancel flights, we also wanted to study the potential benefit we could achieve in delay recovery in the event of an early GDP cancellation. To test our model we used 5 scenarios in which we assumed early clearance times of 0, 15, 30, 45 minutes and 1 hour. The performance in each case is shown in Figure 9 below. Delta and AirTran both

experience noticeable reduction in the overall delay as the extent of the early cancellation reaches one hour. This is not entirely surprising in the case of Delta because they have a greater number of cross-continental and international flights than regional carriers such as American Southeast Airlines and operate more flights than American and are in a better position to recover the delay in the event of cancellation.

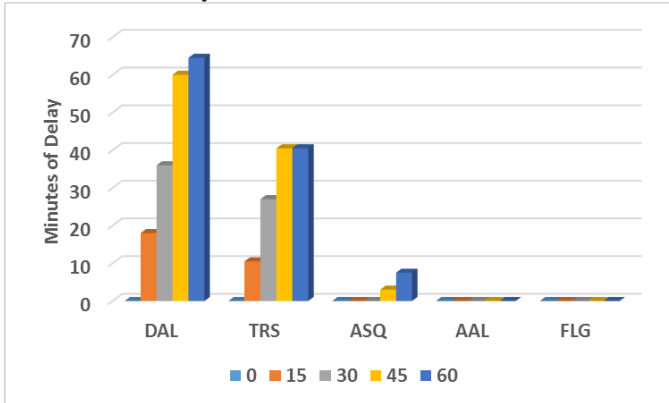


Figure 9: Minutes of Passenger Delay Recovered in each Scenario

V. CONCLUSIONS AND PERSPECTIVES

In this paper we proposed a new strategy for managing ground delay programs. The strategy incorporated both Controlled Departure and Arrival Times as well as en route speed control. It also eliminated the use of an exemption radius which provides incentives for carriers to create their own hedging strategies. To model performance under our new framework we adapted a stochastic model to account for airline hedging. Our analysis suggests that under our new set of GDP controls airlines are significantly less likely to cancel flights because they hope to recover delay in the event of early cancellation. Below we discuss implementation and also suggest implications on NextGen.

A. Near-Term Implementation

The two biggest challenges to near-term implementation are i) insuring CTA compliance (as was mentioned at the outset) and ii) modifying the various GDP procedures to support the proposed architecture. Two types of enforcement can be envisioned. First, violations could be monitored and flight operators with poor records penalized in various ways. Second, as time-based metering methods are implemented CTA information could be communicated to these systems so that they could be “CTA-aware” and aid in insuring compliance. Regarding ii), the research in this paper as well as the work on various speed control measures could be adapted to provide revision and dynamic CTA adjustment methods applicable to this context. It is probably safe to say there are no major road blocks, just the requirement for further development and experimentation with the existing concepts. The research in this paper also should provide a starting point for airline decision support models. A variety of approaches (some simpler, some more complex) are possible. There will be new information exchange requirements including the need for information on the limits to which CTAs can be changed

for airborne flights. Of course integration with time-base metering tools would also induce new information requirements.

It should be admitted that in the near term the full benefits envisioned could not be achieved as they require complete flexibility on the part of each flight to independently adjust its speed. This limitation suggests certain NextGen goals as discussed below.

B. Far-Term Implementation and Implications for NextGen

NextGen and Sesar both express a TBO vision in which flight timing will be closely monitored and controlled. Implicit in this vision is the ability to insure some degree of CTA compliance. In fact, one can view the architecture we have described as a (partial) vision of how GDPs would be migrated to a TBO-based NAS. NextGen and Sesar technologies also should provide the ability for flights to more independently adjust their speeds. This in turn should allow for the benefits described in this paper to be more completely realized. It is perhaps instructive to consider the underlying operational concept of our architecture. Note that, while there is a high degree of control over en route flight timing, there is also an assumption of a high degree of flexibility. This is not compatible with a TBO vision in which a 4D trajectory is set at the time of flight departure and then rigorously adhered to for the remainder of the flight. Our vision calls for a high degree of control and system-wide coordination among 4D trajectories coupled with the ability to dynamically adjust those trajectories to achieve flight operator and ANSP objectives. We feel it is important to incorporate this vision into future TBO architectures.

VI. ACKNOWLEDGMENTS

This work was supported in part by the FAA through the NEXTOR-II Consortium.

REFERENCES

- [1] M. Wambsgans, “Collaborative decision making through dynamic information transfer,” *Air Traffic Contr Quarterly* 4 (1996), pp. 107–123.
- [2] Ball, M. O., Hoffman, R. L., Knorr, D., Wetherly, J., & Wambsgans, M. (2001). Assessing the benefits of collaborative decision making in air traffic management. *Progress In Astronautics and Aeronautics*, 193, pp. 239-252.
- [3] Chang, K., Howard, K., Oiesen, R., Shisler, L., Tanino, M., & Wambsgans, M. C. (2001). Enhancements to the FAA ground-delay program under collaborative decision making. *Interfaces*, 31(1), 57-76.
- [4] Richetta, O.; Odoni, A. Solving optimally the static ground-holding policy problem in air traffic control. *Transportation Science* 1993, 27, 228-238.
- [5] Ball, M.O., Hoffman, R., Odoni, A. and Rifkin, R., 2003, “A Stochastic Integer Program with Dual Network Structure and its Application to the Ground-Holding Problem”, *Operations Research*. 51, 167 – 171
- [6] Mukherjee, A. and M. Hansen, 2007, "A dynamic stochastic model for the single airport ground holding problem," *Transportation Science*, vol. 41, 444-456.
- [7] Ball, M.O. and G. Lulli, 2004, “Ground Delay Programs: Optimizing over the Included Flight Set Based on Distance”, *Air Traffic Control Quarterly*, 12, 1 – 25.

- [8] Ball, M.O., R. Hoffman and A. Mukherjee, 2010, "Ground Delay Program Planning under Uncertainty Based on the Ration-by-Distance Principle", *Transportation Science*, **44**, 1-14.
- [9] T. Vossen, and M. Ball, "Optimization and Mediated Bartering Models for Ground Delay Programs," *Naval Research Logistics*, vol 53, 2006, pp. 75 – 90.
- [10] T. W. Vossen, and M. O. Ball (2006). Slot trading opportunities in collaborative ground delay programs. *Transportation Science*, *40*(1), 29-43.
- [11] Vossen, T., Ball, M.O., Hoffman, R. and Wambsganss, M., 2003, "A General Approach to Equity in Traffic Flow Management and its Application to Mitigating Exemption Bias in Ground Delay Programs", *Air Traffic Control Quarterly*, **11**, 277-292. (a preprint was published in *Proceedings of 5th USA/Europe Air Traffic Management R&D Seminar*, 2003).
- [12] Fearing, D., & Barnhart, C. (2011). Evaluating air traffic flow management in a collaborative decision-making environment. *Transportation Research Record: Journal of the Transportation Research Board*, *2206*(1), 10-18.
- [13] Brinton, C., Provan, C., Lent, S., Prevost, T., & Passmore, S. (2011, June). Collaborative departure queue management. In *Ninth USA/Europe Air Traffic Management Research and Development Seminar*. Berlin, Germany.
- [14] Bhadra, D., Knorr, D. A., & Levy, B. (2012). Benefits of virtual queuing at congested airports using ASDE-X: A case study of JFK Airport. *Air Traffic Control Quarterly*, *20*(3), 225.
- [15] F. Neuman and H. Erzberger, "Analysis of delay reducing and fuel saving sequencing and spacing algorithms for arrival traffic," NASA Technical Memorandum, 103880, 1991.
- [16] H. N. Swenson, J. Thipphavong, A. Sadovsky, L. Chen, C. Sullivan, and L. Martin, "Design and evaluation of the terminal area precision scheduling and spacing system," 9th USA/Europe Air Traffic Management R&D Seminar, Berlin, Germany, June 2011.
- [17] S. Grabbe, B. Sridhar, A. Mukherjee, and A. Morando, "Traffic Flow Management Impact on Delay and Fuel Consumption: an Atlanta Case Study," *Air Traffic Control Quarterly*, Vol, *20*(3) pp. 203-224, 2012.
- [18] P. M. Moertl, "Airline based en route sequencing and spacing field test results: Observations and lessons learned for extended metering," 9th USA/Europe Air Traffic Management R&D Seminar, Berlin, Germany, June 2011.
- [19] Airservices Australia, "Annual Report: 2007-2008," Airservices Australia, Tech. Rep., 2008.
- [20] G. McDonald, and J. Bronsvort. "Concept of Operations for Air Traffic Management by Managing Uncertainty through Multiple Metering Points." *Air Transport and Operations: Proceedings of the Third International Air Transport and Operations Symposium 2012*. IOS Press, 2012.
- [21] J. Leib, "Flights flow gets innovative fix", Denver Post, Denver, CO, May 26, 2008.
- [22] D. Nieuwenhuisen, and N. de Gelder. "Optimizing Nightly Schiphol Traffic through Time Based Operations." *Air Transport and Operations* (2012): 205.
- [23] D. Knorr, X. Chen, M. Rose, J. Gulding, P. Enaud, and H. Hegendoerfer, "Estimating ATM Efficiency Pools in the Descent Phase of Flight," in 9th USA/Europe Air Traffic Management Research and Development Seminar (ATMf2011), Berlin, Jun. 2011.
- [24] J. C. Jones, D. J. Lovell, and M. O. Ball, "En Route Speed Control Methods for Transferring Terminal Delay," in 10th USA/Europe Air Traffic Management Research and Development Seminar (ATMf2013), Chicago, Jun. 2013.
- [25] L. Delgado and X. Prats, "En-Route Speed Reduction Concept for Absorbing Air Traffic Flow Management Delays," *Journal of Aircraft*, vol. 49, no. 1, pp. 214–224, Jan. 2012
- [26] X. Prats and M. Hansen, "Green delay programmes, absorbing ATFMdelay by flying at minimum fuel speed." in 9th USA/Europe Air Traffic Management Research and Development Seminar (ATMf2011), Berlin, Jun. 2011.
- [27] L. Delgado and X. Prats, "Effect of Radii of Exemption on Ground Delay Programs with Operating Cost Based Cruise Speed Reduction Case study: Chicago O'Hare International Airport," in 10th USA/Europe Air Traffic Management Research and Development Seminar (ATMf2013), Chicago, Jun. 2013.
- [28] J. C. Jones, and D. J. Lovell. "Methods for Curbing the Exemption Bias in Ground Delay Programs Through Speed Control." in *Transportation Research Record*, presented at 93rd Annual Meeting of Transportation Research Board, Washington DC, USA, January, 2014.
- [29] Pourtaklo, Nasim Vakili, and Michael Ball. "Equitable allocation of enroute airspace resources." *Eighth USA/Europe Air Traffic Management Research and Development Seminar, Napa, CA*. 2009.
- [30] <http://airlines.org/data/per-minute-cost-of-delays-to-u-s-airlines/>.
- [31] <http://www.iata.org/pressroom/pr/Pages/2014-02-06-01.aspx>.

AUTHOR BIOGRAPHIES

James C. Jones is a PhD student in the civil engineering department at the University of Maryland. He holds a B.S. in engineering from Swarthmore College, an M.S. in electrical engineering from the University of Maryland and an M.S. in systems engineering from The George Washington University. Prior to pursuing his PhD, James held positions at Northrop Grumman and the RAND Corporation. His research focuses on applying operations research systems engineering methods to solve air traffic management related problems.

David J. Lovell is an Associate Professor in the Department of Civil and Environmental Engineering at the University of Maryland. He holds a joint appointment with the Institute for Systems Research and an affiliate appointment with the Applied Mathematics, Statistics, and Scientific Computing Program. Dr. Lovell received his B.A. degree in Mathematics from Portland State University, and his M.S. and Ph.D. degrees in Civil Engineering from the University of California, Berkeley

Michael Ball holds the Dean's Chair in Management Science in the Robert H. Smith School of Business at the University of Maryland. He also has a joint appointment within the Institute for Systems Research in the Clark School of Engineering is co-Director of NEXTOR-II, an FAA consortium in aviation operations research.. Dr. Ball received his Ph.D. in Operations Research in 1977 from Cornell University.