Tenth USA/Europe Air Traffic Management Research and Development Seminar (ATM2013)

En Route Speed Control Methods for Transferring Terminal Delay

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

Abstract— **In this paper, we present an approach for transferring delay away from the terminal to the en route phase of flight. We propose a multi-objective integer programming model designed to assign delays to flights well in advance of the terminal. The IP model weights an objective of fuel savings and throughput to assign controlled times of arrival to flights 500 nmi from the airport. A series of trade studies is performed to evaluate our concept. First, the model is tuned by developing a Pareto Frontier to identify weight factors on our objective function. We demonstrate that the model can effectively transfer delay en route. This transfer holds up even with relatively moderate carrier compliance. We go on to demonstrate that this delay transfer yields significant fuel savings benefits on a per flight basis.**

I. INTRODUCTION

In recent years heavy demand for air travel coupled with limited airport capacity has led to significant delays within the terminal airspace surrounding many airports across the country. Due in large part to the projected demand increases, this problem, if left unchecked, is likely to worsen in the coming years. Weather fluctuations introduce additional capacity constraints that exacerbate the issue. In the U.S. there is no operational coordination of the arrival times of flights, until the traffic management advisor (TMA) exercises control in the general vicinity of the airport (starting approximately 250 nmi out). As a result, individual flights may behave in a manner that in hindsight is quite inefficient, e.g. accelerate their routes to make their scheduled arrivals times only to be vectored off in terminal airspace to temporarily stem the flow of traffic into the destination airports.

The goal of our research is to develop mechanisms to adjust speed during the en route portion of a flight so as to reduce terminal area congestion and generally improve trajectory efficiency during the arrival phase of a flight. We do not seek to develop optimal arrival profiles but rather to provide coordination during the en route phase to allow the use of efficient arrival profiles. If one takes a long-term view, this challenge can be viewed as a component of time-based metering and more broadly trajectory-based-operations. We seek, however, a mechanism usable in the near term, e.g. say 5 years. Our focus is specifically toward the U.S. air traffic management (ATM) environment; however, many concepts should apply more broadly.

The use of speed adjustments to achieve fuel savings and throughput benefits has been studied for over two decades. Neuman and Erzberger [1] present a number of sequencing and spacing algorithms designed to reduce fuel consumption and en route/arrival delay. These algorithms are designed to work with the current TMA system by advancing the leading aircraft in a sequence while preserving the first-come-first-serve ordering over a sequence of aircraft. They also examine the effect of a slightly more aggressive constrained position shifting algorithm to achieve increases in runway capacity. Carr et al. [2] later studied the effect of a priority based scheduling algorithm in reducing the allocated deviations from the preferred airline arrival times. While these contributions demonstrated noticeable improvements in the capacity of the TMA system to improve fuel and throughput performance, they were operating under a distinct constraint that limited their effectiveness. Since the TMA system only operates out to a range of 250 nmi it has limited ability to impact the major portion of each trajectory and even less ability to coordinate large groups of flights. For example, controllers have limited ability to issue overtakes and must perform sequencing operations at lower altitudes where aircraft fuel burn rates are high.

To the extent that en route speed control measures can alter the flight arrival sequence, the widely studied aircraft sequencing problem is relevant to our work. The problem was first examined by Dear [3] who examined the effect of constraining the movements of aircraft through constrained position shifting (CPS). More recent work [4], [5] has resulted in efficient dynamic programming, integer programming, and heuristic approaches. Despite these advances, the focus of the aircraft sequencing problem has been toward eliminating delay. While such an objective is laudable due to the heavy degree of congestion in terminal airspace, it is not always possible to achieve an acceptable amount of delay and eliminate the need for vectoring by merely optimizing the fight arrival banks. In these cases it is often very beneficial to transfer delay to other phases of flight.

There have been a number of attempts to implement en route speed control programs within industry. United Airlines conducted operational trials into London Heathrow Airport demonstrating that en route speed control fuel efficiency gains saved an average of 45 kg of fuel per flight. This figure excludes any benefit achieved by reducing the overall distance traveled [6]. To deal with the morning rush at Sydney Airport, Airservices Australia developed the ATM Long Range Optimal Flow Tool (ALOFT) to allow pilots to control speed out to 1000 nmi away from the airport. In so doing, they achieved an estimated fuel savings of nearly 1 million kg in 2008 [7]. Delta achieved an estimated $8 million in fuel savings over a 20 month period using a dispatch monitored speed control program known as Attila[8].

Other means of ATM management have also been proposed to alleviate the congestion imposed by terminal and en route traffic. The Terminal Area Precision Scheduling and Spacing System (TAPSS) system builds upon the FAA's Traffic Management Advisor (TMA) system [9]. This system enhances strategic and tactical planning through improved route prediction and constraint scheduling to allow air traffic controllers to optimize capacity and accommodate more fuel efficient maneuvers inside the terminal. The Airline Based En Route Sequencing and Spacing tool sends speed advisories to the Airline Operations Centers (AOC)s to allow crews to more actively manage their speeds in the en route phase of flight [10]. These systems could prove useful for trajectory and speed management; however, simpler alternatives may be possible through the assignment of Controlled Times of Arrival (CTAs).

Our work was specifically motivated by a desire to operationalize the benefits identified in Knorr et al. [11]. This research identified substantial inefficiencies in the arrival phase of flights, considering a set of major airports both in the U.S. and Europe. Further, the authors showed conceptually that much of these inefficiencies could be eliminated by "transferring" delays from the arrival phase of flights to the en route portion. While this research identified the opportunity, it did not explicitly describe any procedure for ATM-managed speed control.

The research described in this paper builds on the initial results described in Jones et al. [12]. This earlier work used a much simpler model of the airport and terminal area and based its experiments on a simpler and less realistic data set.

In this paper we propose mechanisms to issue speed control directives well in advance of reaching the terminal area, e.g. 500 nmi away from the airport. While our method (implicitly) sequences flights, unlike aircraft sequencing models, the primary purpose of our model is to transfer delay away from the terminal area. In section II we develop two operational approaches for transferring delay away from the terminal. Section III presents a formulation of a bi-criterion integer programming model designed to assign delay while achieving improved fuel savings and maintaining a high level of throughput. In section IV, we evaluate the operational impact of our model through a set of computational experiments. We tune this model by developing a Pareto Frontier and identifying effective objective function tradeoff coefficients. We then engage in a set of studies that yields several results. First, we identify the impact of this model by calculating the delay savings. Second, we construct an additional simulation to illustrate how the achieved delay savings affects aircraft fuel consumption. Third, we perform further trades to evaluate the impact of assigning CTAs at various distances from the destination airport. Finally, we investigate the effect of compliance on the performance of our solution.

II. METHOODLOGY

A. *Fuel Savings Assumptions*

The goal of our research is to devise a system that delivers flights to the terminal area in such a way that the descent phase of each flight is as efficient as possible. This can be viewed as the process of transferring delay from the terminal area to the en route portion of a flight. To understand the significance of this goal note that delays are typically taken in the terminal area by adding distance to a flight through a multitude of mechanisms including long "downwind" approach paths (also called "tromboning" – see Figure 1), vectoring and circular holding patterns. On the other hand, given sufficient advance notice, delays can be taken in the en route phase by simply reducing aircraft speed without increasing distance traveled.

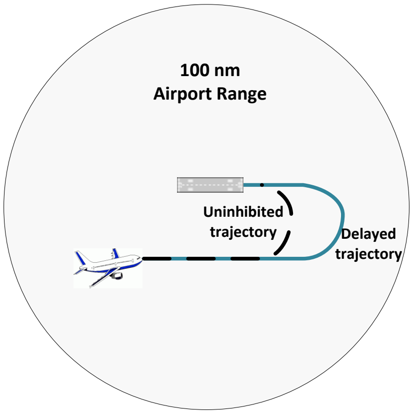

Figure 1: "Downwind" trajectory to absorb terminal area delay

Figure 2 illustrates the relationship between fuel efficiency (specific range) and Mach number. Note that as the Mach number of the aircraft increases, its fuel efficiency increases to a point known as the maximum range, beyond which it begins to decline. The shape of this curve, however, is relatively flat. The flatness of this curve implies that absorbing delay (within limits) during the en route portion of a flight is nearly costless from a fuel usage standpoint. Thus, transferring say 5 minutes of delay from the terminal area to the en route phase is approximately equivalent to reducing the length of the flight by a distance corresponding to 5 minutes in travel time.

Note also from Figure 2 that as altitude increases the specific range curves move decisively upward. Since the magnitude of the upward shift of the specific range is large relative to the increases along an individual curve at constant altitude, fuel efficiency at a high altitude is decisively greater regardless of whether the Mach number changes significantly. This implies that if, as is typical, excess distance in the terminal airspace is taken at lower altitudes, then the fuel burn rate is higher than would be the case for a similar distance at a higher altitude. Thus, there are two very strong effects at work that produce fuel cost savings when delay is transferred from the terminal area to the en route portion.

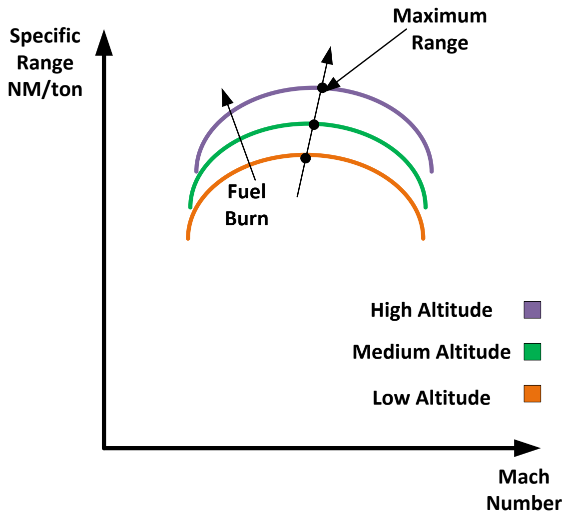

Figure 2: Notional variation in aircraft fuel efficiency with speed at various altitudes.

B. *Operational Concept*

In the proposed operational scheme, CTAs are assigned to flights once they reach an approximate distance of 500 nmi from the destination airport. Once a CTA has been assigned to a flight it notionally proceeds to the assigned metering fix 150 nmi from the airport while exercising the appropriate speed control guidelines. When the flight reaches the metering fix TMA would issue adjustments to controllers to effectively guide the flight on its assigned STAR trajectory. Under this concept the system does not require close coordinate with TMA.

It is important to recognize that this is not a static problem. The changing environmental conditions necessitate that any assignment algorithm make use of revised information as it is presented. As flights travel en route their estimated times of arrival (ETAs) are updated on an ongoing basis. As flights get closer to their destination these ETAs become increasingly reliable. The ETAs provide a forecast of the degree of congestion and the resultant excess flight time and maneuvering that will occur in the terminal area. The assignment of CTAs effectively adjusts the ETAs to provide a more orderly flow of traffic into the terminal area, thereby injecting an increased level of predictability into the flow of traffic. In the longer term, the 500 nmi horizon could be lengthened and also could vary by flight.

Under this approach the ANSP would update the list of flights that were available for scheduling every 15-30 minutes. At each period the ANSP would set the number of "slots" at the metering fix based on the capacity of the airport and each metering fix. When the number of slots has been determined an optimization model assigns a CTA to each flight once it reaches the 500 nmi boundary. These CTAs could be assigned using various communications tools discussed in the following section. When the pilot receives this CTA he/she would enter this time into the Flight Management System (FMS) on board the aircraft. The aircraft could then calculate the preferred route and speeds en route and proceed to the metering fix where it would then receive TMA-based controller instructions. It is important to note that the assignment process is iterative and dynamic. At each period a new set of flights between 1-30 minutes away from the 500 nmi boundary is evaluated by the assignment algorithm. Once the set of CTAs has been decided based on our model's logic the flights receive a CTA only once they approach the 500 nmi boundary. Note that there will generally be overlap between the set of flights considered from one iteration to the next as only the closest-in flights are given the computed CTAs. Thus, the CTAs computed for the further-out flights are temporary; these flights are included to provide an assignment procedure with a more global perspective of total flight demand.

C. *System Description*

For U.S. implementation, we anticipate that the Air Traffic Control Systems Command Center – ATCSCC – would have responsibility for determining the CTAs. It is also the case that the data required to support these decisions are already readily available to the Command Center. The traffic flow management system (TFMS) integrates real time flight information such as estimated arrival times, scheduled arrival times, landing times, flight, aircraft positions and flight cancelations. The Command Center also has rich weather feeds and through consultation with airport Air Traffic Control Towers (ACTs) and Terminal Radar Approach Control facilities (TRACONs), up-to-date information on airport and terminal area capacities.

In the longer term, CTAs would certainly be transferred to aircraft using datalink. However, this option will most likely not be possible in the shorter term. Thus, after examining the existing communications technology between pilots and command centers we see two options for assigning CTAs. In the first the Command Center would pass CTA assignments to the Air Route Traffic Control Center (ARTCC), who would inform the pilots of these assignment times via controllers / radio communication link. In the second approach the Command Center would communicate CTAs to appropriate airline operational control centers (AOCs). It is possible (at least in the longer term) that CTAs could be adjusted based on Command Center / AOC negotiation. Once a CTA was finalized the AOC would send it to the appropriate aircraft over the Aircraft Communications Addressing and Reporting System (ACARS). Notionally this approach has the advantage of very naturally supporting the inclusion of (future) Collaborative Decision Making (CDM) features.

The first approach offers a significant advantage from a compliance standpoint. Since assignments are issued directly by air traffic controllers they will likely be taken quite seriously. Further this approach would by necessity offer a degree of coordination between the CTA directives and other controller directives, e.g. those emanating from TMA. This approach could, however, impose an additional workload burden on some air traffic controllers and increase training needs at certain control centers. The second approach minimizes the burden on ATC staff by limiting their direct

---

This work was supported by NEXTOR-II, under a Contract with the Federal Aviation Administration; sponsor agreement number: `DTFAWA11D00017`. Any opinions expressed herein do not necessarily reflect those of the FAA or the U.S. Department of Transportation.

involvement in the assignment process. Although the ATC staff may issue resource capacity guidelines and updates to command centers to inform them in their decision making, the assignments would be made jointly by carriers and the Command Center. This process allows the carrier to actively voice their priorities during the assignment process and potentially adjust their assignments through CDM mechanisms. The price of such accommodation, however, may be borne at the expense of operational effectiveness. If compliance is sufficiently low it will likely prove quite challenging to realize a substantial portion of the potential benefit pool. Thus it is critical that carriers actively enforce CTAs on their flights. Figures 3a and 3b illustrate the flow of data between systems and stakeholders.

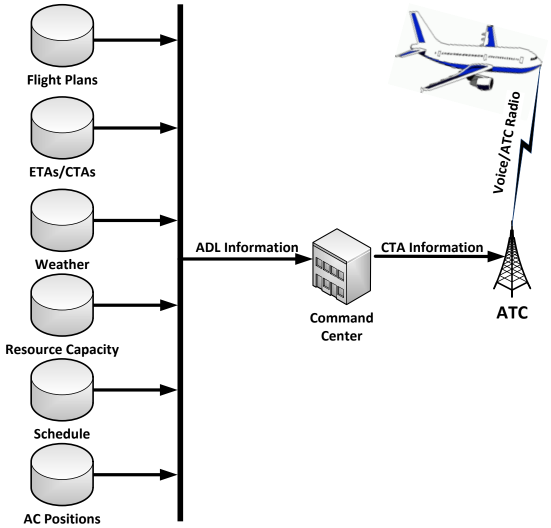

Figure 3a: Information flow between databases aircrafts and command centers under a centralized approach.

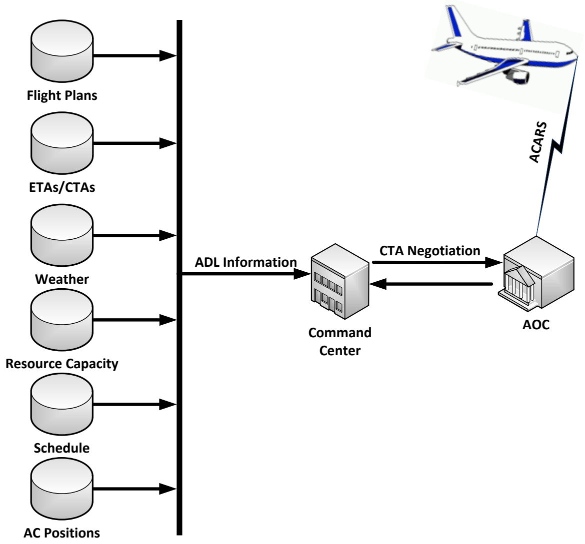

Figure 3b: Information flow between databases aircrafts and command centers with collaboration from carriers.

III. MODEL DESCRIPTION

In this section we develop the model structure used to coordinate and assign delays to flight arrival banks. The model assumes a multi-resource framework and dynamically resolves the problem to accommodate the changing conditions within the airspace. It aims to transfer delay away from the terminal and reduce fuel consumption while exploiting opportunities to improve throughput within flight streams.

There are a number of significant scheduling limitations with implementing the reordering of airborne flights. While flights can be delayed on the ground for long stretches of time, we cannot impose the same delay lengths in the air due to fuel limitations. Moreover, issuing large airborne delays through vectoring imposes a considerable burden on air traffic controllers. In addition it is often advantageous from a system perspective to speed up flights when other flights are immediately behind them as it reduces the size of the arrival queue. These speed-ups do not always yield universal benefits as they can add to the fuel costs of the flights being sped up. We have tried to consider all of these factors in our model development.

A. Model Objectives

There are three key performance criteria we consider in modeling this problem;

1) *Fuel usage and delay transfer: the principle motivation for the overall procedure is to allow efficient and unimpeded trajectories in the terminal area. To accomplish this, we seek to transfer delay from the terminal area to the en route portion of the flight, where it can be more efficiently absorbed.*

2) *Arrival throughput: it is important to maintain a high throughput into the airport, while effectively transferring delay. Ideally any model formulation will increase arrival throughput.*

3) *Equity: in assigning ETAs, it is inevitable that the natural order of the arriving flights will be perturbed. It is important that any flight prioritization be carried out in an equitable manner. Employing mechanisms designed to promote a more equitable outcome will also incentivize compliance among carriers.*

While each objective represents a relatively distinct measure of our system's operational effectiveness, the mathematical expressions associated with each metric can take a variety of forms. To provide some additional clarity to our metrics we briefly discuss our approach for incorporating each term. Fuel usage can be represented on an aircraft-by-aircraft basis using a piecewise linear function to represent aircraft fuel costs. These costs will vary by aircraft type and speed. The appropriate costs can be selected using integer variables by assigning each flight a specific arrival time.

While there are several proxies commonly employed to represent throughput, including makespan and arrival rate, we adopt the metric of system delay. Specifically we seek an objective that acts to minimize the total system delay. An expression of this objective is shown in equation (1). To formulate the expression we define the following parameters:

F – The set of all flights

A – The set of all Airlines

t_{ia} – The arrival time assigned to flight i of airline a

e_{ia} – The earliest possible arrival time that can be assigned to flight i of airline a

$$D = \sum_{i \in F, a \in A} \left(t_{ia} - e_{ia} \right)^{+} \qquad (1)$$

System equity can serve two functions within an algorithmic model: (1) It serves to ensure that some measure of fairness towards carriers is present within the model and (2) to promote compliance among affected flights. Widely used mechanisms for equity often include first-scheduled-first-served algorithms used in Ground Delay Programs such as Ration-by-Schedule (RBS) and the first-come-first-served algorithm currently used in runway sequencing. Constrained Position Shifting (CPS) has often been used to limit the amount of reordering in airborne sequencing problems. While we eventually envision a scheme that weights the objectives through proportional allocation for each airline we do not explicitly incorporate any equity mechanisms into the objective function of our model. The constraints on slot eligibility do, however, impose some limit on the amount of inequity that can be introduced by bounding the range of possible reassignment times. While we seek to reorder the flights to facilitate our other objectives, due to the physical dynamics of the environment we can only assign flights up to 15 minutes of delay per flight. While we also permit overtakes, speed-ups can only occur over a range of 5 minutes ahead of the estimated time of arrival.

B. *Model Formulation*

The changing environmental conditions necessitate that we consider an approach designed to adaptively reflect the most current information. A dynamic framework was adopted to incorporate periodic information updates into our model. We envision a framework in which TFMS data is continually provided to our model to inform us of the most up-to-date ETA and aircraft position related information. The model would then take the information from the two most recent periods to assign CTAs for flights within 30 minutes of the 500 nmi boundary. These CTAs would be issued and constrain further solutions and would be used along with new ETAs to issue CTA assignments in the next period. This process is illustrated in Figure 4.

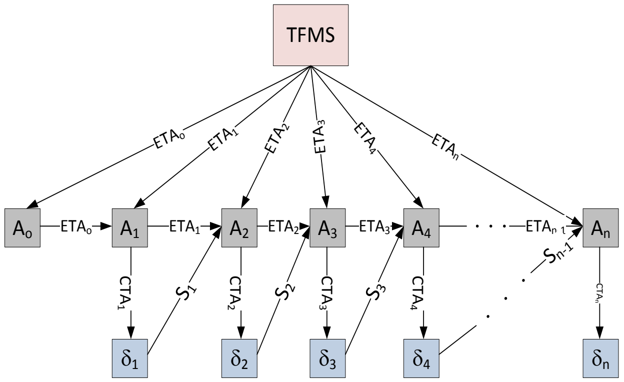

Figure 4: An illustration of the model's dynamical time slot allocation over a two period rolling horizon. TFMS provides ETA and position information to our model on an iterative basis. The algorithm uses this ETA information to assign CTAs. These CTA then form the basis for future assignments.

A multi-objective integer programming model was developed to incorporate our two primary objectives: fuel savings and throughput. This model considers flight assignments over a rolling two-period horizon by discounting the second period to account for a lower degree of confidence in more distant events. In order to limit the number of constraints we imposed certain restrictions on some of our sets. Since it is unreasonable for flights to periodically change their approaching corner post we restricted each flight to its planned fix at 500 nmi from the airport. We also restricted the range of slots over which each flight could be assigned to correspond to +5 to -15 minutes. We define our variables and parameters as follows:

F –set of all flights

S_{if} – set of all slots for flight f at fix i

S_{rf} – set of all slots for for flight f at runway r

Ω_f –set of eligeble fixes for flight f

R –set of all runways

λ – weight assigned to each term in the objective criteria

β – A discount factor applied to the second assignment period at each iteration

b_{fk}^{i}- the fuel burned by flight f when assigned to slot k at fix i

b_{f}^{min} - the minimum possible fuel burn by flight f for all feasible slots

t_{ks}^{ir}- the slot corresponding to slot k at fix i and slot s at runway r

e_{fks}^{ir}- the earliest possible slot k at fix i and slot s at runway r that can be assigned to flight f

c_{fks}^{ir}- the system delay cost of assigning flight f to slot k at fix i and slot s *at* runway r

d_{fks}^{ir}- the fuel cost of assigning flight f to slot k at fix i and slot s *at* runway r

$$x_{fks}^{ir} = \begin{cases} 1 \text{ if flight } f \text{ is assigned to slot } k \text{ at fix } i \text{ and slot } s \text{ in runway } r \\ 0 \text{ otherwise} \end{cases}$$

$$\min \sum_{\substack{f\in F, k\in S_{if},\\ i\in\Omega_f, s\in S_{rf},\\ r\in R}} \begin{array}{l} \lambda c^{ir1}_{fks} x^{ir1}_{fks} + (1-\lambda) d^{ir1}_{fks} x^{ir1}_{fks} + \\ \beta\left(\lambda c^{ir2}_{fks} x^{ir2}_{fks} + (1-\lambda) d^{ir2}_{fks} x^{ir2}_{fks}\right) \end{array} \tag{2}$$

$$S.T. \quad \sum_{\substack{k\in S_{if}, i\in\Omega_f,\\ s\in S_{rf}, r\in R}} x^{ir2}_{fks} + x^{ir2}_{fks} = 1 \quad \forall f \in F \tag{3}$$

$$\sum_{\substack{f\in F, s\in S_{rf},\\ r\in R}} x^{ir1}_{fks} + x^{ir2}_{fks} \le 1 \quad \forall k \in S_{if}, \forall i \in \Omega_f \tag{4}$$

$$\sum_{\substack{f\in F, k\in S_{if},\\ i\in\Omega_f}} x^{ir1}_{fks} + x^{ir2}_{fks} \le 1 \quad \forall s \in S_{rf}, \forall r \in R \tag{5}$$

$$x^{ir1}_{fks}, x^{ir2}_{fks} \in \{0,1\} \tag{6}$$

Equation (3) states that in each time period every flight is assigned to one slot. Equation (4) states that each slot at each fix can be assigned to no more than one flight. Equation (5) states that each slot at each runway can be assigned to no more than one flight. Equation (6) states that our decision variables for each period are binary.

Equation (2) forms a two period weighted objective function of throughput and fuel costs with the second period discounted. Our throughput coefficients will vary based on the amount of time between their corresponding slot and earliest possible arrival time. Our fuel cost coefficients will also vary based on the fuel burn rate for each assigned flight (aircraft). A more explicit expression of each coefficient is shown in equations (7) and (8).

$$c^{irt}_{fks} = t^{irt}_{ks} - e^{irt}_{fks} \tag{7}$$

$$d^{irt}_{fks} = \frac{b^{i}_{fk}}{b^{\min}_{f}} \tag{8}$$

IV. EXPERIMENT

.

In this section we describe a computational experiment based on historical data and provide the results of that experiment.

A. Scenario Description

The basis of our experiment is data set collected from Atlanta Hartsfield-Jackson Airport (ATL) on May 1, 2012. The weather conditions were clear and sunny and all runways were active. The data set was obtained by merging data from two sources: an ADL file (obtained from the FAA's Traffic Flow Management System) and an ASDX file (surface surveillance data). The key fields included: flight number, collection time stamp, expected time of arrival (ETA), scheduled time of arrival (STA), origin airport, actual time of departure, current aircraft position, aircraft type, runway arrival time, STAR trajectory and last available fix.

We assumed airport acceptance rates of up to 100 flights per hour. This assumption is consistent with ATL operating practice under the weather conditions for the time in question (full use of 2 runways and partial use of a third). The experiment was run over a 4 hour period from 2:00-6:00 EST. A slot width of 90 seconds was used.

ATL has 4 corner posts at the northeast, northwest, southeast and southwest corners of the airport. Arriving flights commonly fly through one of these corner posts and are sent to one of 3 runways: 2 runways are used full time and another runway is partially used. The runway capacity was bound by the wake vortex separation requirement between classes of aircraft. Based on the general fleet mix present in the data we found that we could assign uniform slot sizes that could be later adjusted to achieve tighter spacing.

To study the problem we developed a simulation intended to model the basic effects of TMA and more generally the manner in which the CTAs produced by our model would impact operations in practice. The simulation assumes that each flight is assigned an arrival fix and a CTA assigned by the IP. The flight then proceeds to the metering fix and attempts to arrive at the assigned CTA. Randomization is applied to the travel times between the boundary and the metering fix so that flights arrive within the near vicinity of their CTA but not necessarily at that specific time. The simulation then accepts flights for vacant runway slots on a first-come-first-served basis. Since not all flights receive a CTA from our model, e.g. those whose departure airport is within 500 nm of ATL, the simulation may have to assign delay to flights after they reach the arrival fix in order to properly manage the runways.

A baseline run was used to evaluate the delay performance with no intervention. This trial used flight ETAs and projected them backward to get the approximate arrival time at the metering fix. The travel times between each fix and runway were modeled by fitting flight data with separate normal distributions and sampling from these distributions.

It is also worth noting that there were a significant number of short haul flights that originated from airports within 500 nmi of the airport en route. These flights composed 54.35% of all flights arriving at the terminal. The short haul flights were merged with the stream of CTA assigned flights outside 500 nmi to create an integrated sampling of flights. Since these flights do not have assigned CTAs we cannot actively manage their approaches until they reach the metering fix where they are accepted on a first-come-first-serve basis. As such these flights impose some limit on our ability to assign delay.

B. Analysis of Objectives

Our model incorporates two objectives that each play a critical and distinct role in the performance of our concept. Yet despite their importance it is not immediately clear what the relative magnitudes of their respective coefficients should take. To better understand the performance trade-off between each objective the coefficients were weighted iteratively to generate a Pareto Frontier. The model was run deterministically over a 4 hour time window in our dataset in 15 minute periods. Averages of both the system delay and fuel burn per flight were taken over the resulting outputs. Eleven different weighting combinations were used to create the frontier. This frontier captures the relative impact of each objective against performance goals. The resulting curve from this trade is shown in Figure 5.

We focus on the second point of the curve where the drop in throughput begins to level off. The presence of this point along the curve suggests that a great deal of the benefit of system throughput can be achieved with very little weight. We are therefore able to prioritize fuel conservation while having little impact on our throughput performance. We therefore select the weights of 0.1 and 0.9 for throughput and fuel respectively and will adopt this weighting for the remainder of our experiments.

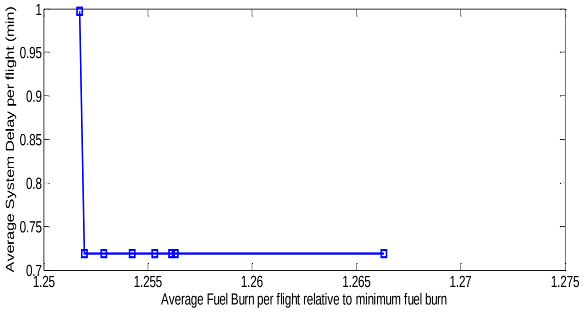

Figure 5: A Pareto Frontier illustrates the trade-off between fuel savings and throughput in our IP objective function.

C. *Delay Savings*

A simulation was run to evaluate the ability of our model to transfer delay away from the terminal. A baseline measurement was performed to gauge the amount of delay present in the terminal without our intervention. This run measured the amount of time that each flight spent in terminal airspace while waiting for a runway. If a flight arrived in the queue and could not receive a runway slot when it was within the allotted travel time it then waited until a space opened up. This waiting time was measured and averaged to calculate the average delay. We then configured our model to assign CTAs to flights near 500 nmi of the airport in 15 min intervals. We repeated the run with the assigned CTAs and measured the average delay per flight. This delay was compared relative to the average delay without intervention. For clarity an expression for calculating transferred delay is shown in equation (9).

$$Avg\_D_{transferred} = Avg\_D^{baseline} - Avg\_D^{CTA} \qquad (9)$$

Figure 6 shows the delay transfer of our model relative to total delay. The figure suggests that the model is able to transfer the majority of the delay away from the terminal. The figure exhibits reasonable tracking although in some instances the improvement is lacking. One potential reason for this deviation is the presence of unmanaged short haul flights that appear at the beginning and end of assignment periods. When they occur such flights can push back successive long haul controlled flights away from their initial assignment times. Over the 4 hours of flights samples the proportion of short haul (less than 500 nmi from the airport) unmanaged flights exceeded managed long haul flights by a 19%. These flights impose some limitation on our overall ability to transfer delay. It might be worth considering coupling such an initiative with a ground delay to maximize the potential benefit. Yet despite this limitation substantial benefits in fuel savings can be realized with even modest transfers.

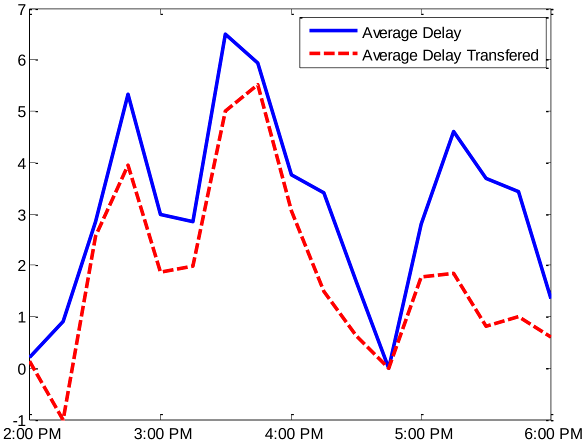

Figure 6: Delay Transferred Away from the Terminal using our IP model at 500 nmi.

D. *Effect of Compliance*

In section II we described two methods for transferring delay away from the terminal. The first involved issuing CTAs through air traffic controllers while the second used an AOC. This choice will likely have a strong impact on the degree of compliance within the system. As noted due to strict requirements on pilots to obey directives from controllers, ATC issued CTAs are likely to be significantly more well adhered to relative to those issued by the AOC. While this might not completely undermine its viability it is important to understand the ramifications behind the approach.

To that end we sought to evaluate the impact of compliance on system performance. A compliance threshold was introduced into our model to study the problem. A random number generator was used in conjunction with the threshold to identify whether a flight had chosen to comply with the assigned CTA. When the random number draw exceeded the threshold the associated flight was then rescheduled to arrive at its original ETA. These flights were then modeled to reach the metering fix at the assigned CTA/ETAs and assigned to a runway. The resulting performance of at 8 levels of compliance is shown in TABLE I.

The table shows modest delay transfer at compliance levels as low as 40%. This relatively robust performance suggests that the system can operate at levels well below strict compliance. This is encouraging as it is probably unrealistic to expect the system to consistently operate at levels near 100% due to the complexities within the airspace. It still is questionable, however, whether AOC issued assignments could even achieve these levels of conformity. Given this uncertainty and potential need for monitoring of the carriers this evidence might slightly point us in favor of issuing CTAs through ATC.

TABLE I: THE EFFECT OF COMPLIANCE ON DELAY TRANSFER

Compliance Level %	Average Delay Transfer (min)	Percentage of Total Delay Transferred
100.00	1.69	58
90.00	1.26	43
80.00	1.31	45
70.00	1.41	48
60.00	1.01	34
50.00	1.02	35
40.00	0.95	33
30.00	-0.40	-14

E. *Fuel Savings Impact*

While we have established that en route speed control can transfer delay away from the terminal our primary objective is to save fuel. We would like to understand how such delay savings translates into fuel conservation. In order to measure the average fuel savings we needed to conceptualize the way that the savings occurs. Substantial fuel savings can result when delay on a given flight is transferred from the terminal area to the en route phase of flight. While some of this savings results from transferring the site of delay from a lower to a higher altitude, the majority of the benefit is attributable to the reduction in miles traveled. As we discussed in section 2, terminal delay is applied largely by path extension of flights. By transferring the delay en route we are able to eliminate a considerable portion of the extended path. Since the fuel burn rates en route are nearly equivalent for the standard and speed controlled flights the conservation of fuel achieved through the reduction in path extension in terminal airspace is essentially free. In the case of short haul flights this benefit is even more apparent. As these flights are not involved in the CTA assignment process they do not incur any en route delay in this scheme, they do however, reap the benefits of a less congested terminal airspace. Therefore the benefit is also free for those flights and can be calculated by computing the resulting reduction in fuel burn inside the terminal.

In order to explicitly calculate the average savings rate incurred on a per flight basis we measured the fuel burn rate near the terminal at various altitudes. We assumed that the aircraft vectoring inside the terminal would occur over a range of FL100 to FL250. With this range we sampled over a set of altitudes from an empirical inverse CDF. These altitudes were then used to measure the average fuel burn rate at a given speed based on values obtained from the BADA database. Separate values were computed for short and long haul flights. The results of these computations can be seen in Figure 7 below.

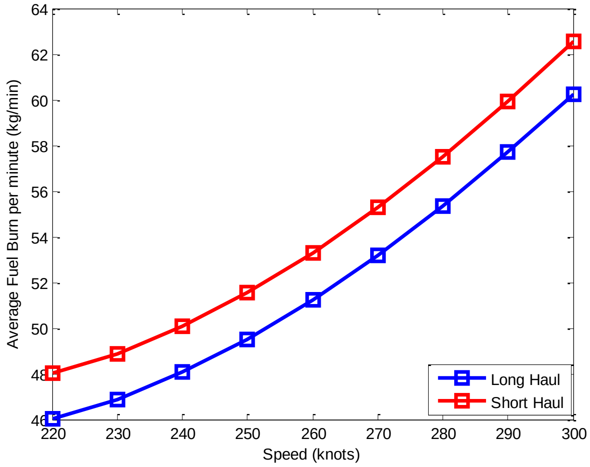

Figure 7: Average Fuel Burn Rates of long haul (500 nmi) and short haul flights at various speeds (CAS).

Not surprisingly the graph suggests that there is a slight nonlinear relationship between fuel burn rate and speed. Using the fuel burn rate values obtained from Figure 7 we can express the cumulative fuel burn savings rate by taking a weighted average of the respective savings rates for the long and short haul flights. The resulting calculations are shown in TABLE II and displayed graphically in Figure 8.

TABLE II: THE FUEL SAVINGS RESULTING FROM DELAY TRANSFER

Speed (knots)	Fuel Burn Savings rate kg/min (Long Haul)	Fuel Burn Savings rate kg/min (Short Haul)	Fuel Burn Savings rate kg/min
220.00	46.02	48.03	47.11
240.00	48.07	50.08	49.16
260.00	51.23	53.30	52.35
280.00	55.33	57.50	56.51
300.00	60.25	62.55	61.50

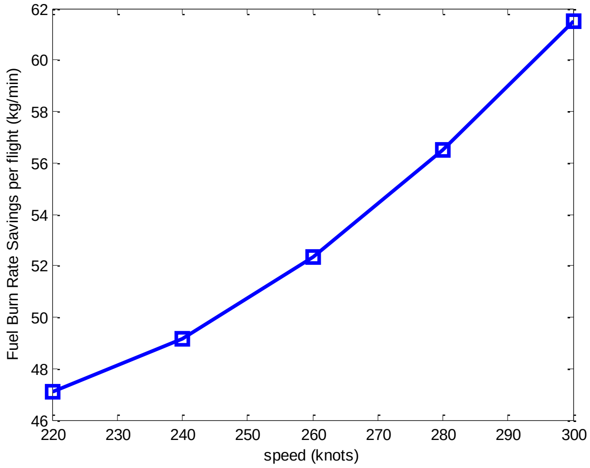

Figure 8: Average Fuel Burn Rates Savings (kg/min) from the total fleet mix vs. Speed (CAS).

Given the inherent fuel burn savings rate associated with moving small amounts of delay away from the terminal it is illustrative to examine how the delay transfer curve shown in part c translates to direct fuel savings. Figure 9 shows the fuel burn savings made possible by the delay transfer relative to five different vectoring speeds at the terminal. A comparison of the plots shows that the savings is considerable regardless of vectoring speed. Although the savings is largest when vectoring at 300 knots, in every case in this example we are able to save an average of 86 kg per flight over the 4 hour period.

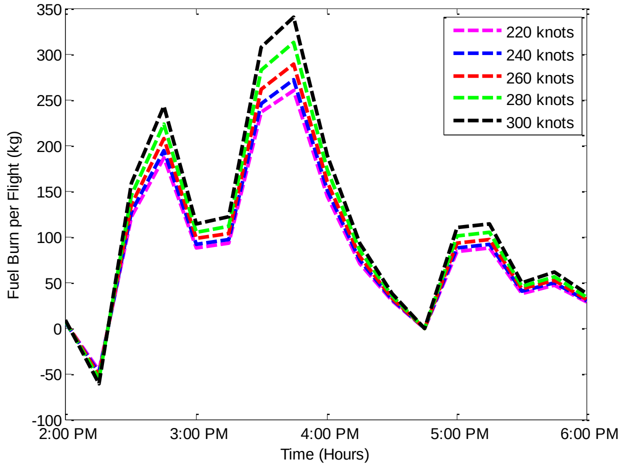

Figure 9: Average Fuel Burn Savings vs. time over a 4 hour period.

V. CONCLUSIONS

In this paper we have explored a method for transferring delay away from the terminal using en route speed control. We provided an operational overview of the key aspects of system implementation. We then developed a bi-criteria integer programming model to facilitate delay transfer through CTA assignment. A Pareto Frontier was constructed to identify the best balance of weights of fuel savings and throughput in our objective function. Using these weights the model was wrapped in a simulation which demonstrated the model's ability to transfer delay. The capacity of the model to transfer delay remained relatively stable in the face of lower compliance. An analysis of the fuel burn rate of the fleet showed that the delay transfer yielded considerable savings with regard to fuel.

This study raises a number of interesting questions that could prove the subject of future work. Assigning CTAs at a 500 nmi boundary (and not any closer) resulted in a considerable number of unmanaged flights. While it is difficult to issue speed control guidelines to these short haul flights due to the short distance between their origin and destination airports the issuance of ground delays could prove effective at limiting the uncertainty that they inject into the assignment process.

There are currently no mechanisms in our scheme to enforce compliance amongst airlines. While we have demonstrated some resiliency of our model it would be preferable to create a scheme that discourages carriers from ignoring directives.

We recognize that from an implementation standpoint it is often easier to gain acceptance for simple allocation rules rather than using integer programming models. We are currently working to develop a greedy algorithm to perform many of the same functions provided by the IP and we hope to explore the operational feasibility going forward.

VI. ACKNOWLEDGEMENTS

This work was supported by the FAA through the NEXTOR-II Consortium. The authors wish to thank John Gulding and Dave Knorr of the FAA for guidance in the initiation and early concept development of this project and J.P. Clarke and Jim Barnes of Georgia Tech for helpful discussions and for providing data.

REFERENCES

[1] F. Neuman and H. Erzberger, "Analysis of delay reducing and fuel saving sequencing and spacing algorithms for arrival traffic," NASA Technical Memorandum, 103880, 1991.

[2] G. C. Carr, H. Erzberger, and F. Neuman, "Airline arrival prioritization in sequencing and scheduling," 2nd USA/Europe Air Traffic Management R&D Seminar, Orlando, FL, December 1998.

[3] R. G. Dear, "The dynamic scheduling of aircraft in the near terminal area," Research Report R76-9, MIT Flight Transportation Laboratory, Cambridge, MA, USA, 1976.

[4] J. E. Beasley, M. Krishnamoorthy, Y. M. Sharaiha, and D. Abramson, "Scheduling aircraft landings – the static case," Transportation Science, vol. 34(2), pp. 180–197, 2000.

[5] H. Balakrishnan and B. Chandran, "Algorithms for scheduling runway operations under constrained position shifting," Operations Research, vol. 58, pp. 1650-1665, 2010.

[6] Bianco, L., Rinaldi, G., and Sassano, A. "A combinatorial optimization approach to aircraft sequencing problem," Flow Control of Congested Networks. Vol. 38, Springer-Verlag, Berlin, 1987, pp. 323-339.

[7] United Airlines Pilot Briefings, " London Heathrow Trials"

[8] Air Services Australia 2008 Annual Report, Canberra, Australia, 2008.

J. Leib, "Flights flow gets innovative fix", Denver Post, Denver, CO, May 26, 2008.

[9] 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.

[10] 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.

[11] D. Knorr, X. Chen, M. Rose, J. Gulding, P. Enaud, and H. Hegendoerfer, "Estimating ATM efficiency pools in the descent phase of flight: Potential saving in both time and fuel," 9th USA/Europe Air Traffic Management R&D Seminar, Berlin, Germany, June 2011.

[12] J. C. Jones, D. J. Lovell, and M. O. Ball, "Algorithms for dynamic resquencing of en route flights to relieve terminal congestion," 5th International Conference on Research in Air Transportation, Berkeley, CA, May, 2012.

AUTHOR BIOGRAPHY

James C. Jones is a PhD candidate 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.

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.

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