

An Equitable Approach for Transferring Terminal Delay with En Route Speed Control

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 bi-criteria mixed integer programming model designed to assign delays to flights well in advance of the terminal. The model objective weights equity against the potential for scheduling conflicts between flights. We also present an additional integer programming model designed to work in succession with our bi-criteria model by maximizing the throughput of arriving flights. A series of trade studies is performed to evaluate our concept. First, our bi-criteria model is tuned by constructing a Pareto Frontier to identify best weights for our objective function. Both models were tested in a simulation based on historical data. The results demonstrate that our approach can effectively transfer a considerable portion of the delay in the terminal to the en route phase of flight. This delay transfer yields significant fuel savings benefits on a per flight basis.

Keywords— Speed Control, Equity, Traffic Management Initiatives, Integer Programming

I. INTRODUCTION

The National Airspace System (NAS) is continually challenged by heightened levels of congestion. When faced with inclement weather, the capacity of airports is often insufficient to accommodate demand. These capacity/demand imbalances will often lead Federal Aviation Administration (FAA) air traffic managers to impose Ground Delay Programs (GDPs) at airports or Airspace Flow Programs (AFPs) in the air. Such programs limit the influx of arriving flights into the region so that the flow is more compatible with the resource's reduced capacity. Specifically, flights receive a later departure time from their origin airport commensurate with the rate of arrival desired by the affected airport. These programs lower airline operation costs and reduce the workload imposed on air traffic controllers by mitigating the potential resulting airspace congestion.

When a GDP or AFP is implemented flights are assigned a controlled time of arrival (CTA) based on the order that they appear in the schedule. This assignment procedure, known as Ration-by-Schedule (RBS), has become widely accepted as a standard for equitable allocation and serves as an integral component of the Collaborative Decision Making process

(CDM) used to manage Traffic Flow Management Initiatives (TMI). [1][2]

To date GDPs and AFPs have composed the only major forms of TMI implemented within the NAS. They are typically managed at the strategic and pre-tactical levels, last for several hours and require significant advance planning to implement. In recent years, however, the development of the NextGen air traffic management system has led the air traffic management community to discuss new ways in which traffic management initiatives could be conducted to enhance performance within the NAS. Along these lines the thinking has led to a concept known as Trajectory Based Operations (TBOs) [3]. In this framework flights would fly negotiated 4D trajectories which allow for more precise active management of trajectories across all phases of flight. In order to facilitate this NextGen development the ATM community needs to define new ways to implement TMI over shorter time frames with ground delays, speed control and re-routing.

While RBS has served to promote the equitable assignment of delay in GDPs/AFPs, the challenges associated with implementing TMI in a more dynamic setting raise some questions about its applicability to en route flight management. Fluctuations in arrival times due to the influence of wind, gate availability, runway availability and airspace capacity can each cause flights to deviate from their estimated times of arrival (ETAs) predicted at the time of departure. These deviations can lead to scheduling conflicts between arriving flights. As controllers do not manage their sectors in a first scheduled first served manner, such conflicts can lead to increased inequity within the system and can result in increases to both fuel burn in airborne flights and controller workload through vectoring. In this context one could argue that en route TMI initiatives could be enhanced by including some sort of advanced scheduling conflict resolution to supplement an otherwise equitable RBS-like allocation scheme using both speed control and ground delays.

Speed control has been widely studied for a variety of air traffic management applications. At the tactical level Neuman and Erzberger [4] describe 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) has since been proposed. [5] The technology has also been used cooperatively in Traffic Flow Programs. [6] 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 [7].

In recent years the horizon for such air traffic management initiatives has also moved further away from the airport. Airservices Australia developed the ATM Long Range Optimal Flow Tool (ALOFT) to allow pilots to control speeds out to 1000 nmi away from the airport. In so doing, they achieved an estimated fuel savings of nearly 1 million kg in 2008 [8]. Since then they have also used additional metering fixes to better manage trajectory and arrival time uncertainty [9]. 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 [10]. At Schiphol a ground based planning system that interfaced with aircraft through datalink was used to remove vectoring in their nighttime operations [11]. Knorr et al. [12] 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. [13] 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 examined at the pre-tactical level. 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. 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 [14], [15], [16]. Jones and Lovell showed speed control could also be used to help curb the exemption bias in GDP slot assignments [17].

In this paper we propose two models designed to reduce the number of arrival time conflicts through speed control. The first model we present is a bi-criteria mixed integer program (MIP) that weights equity against potential for scheduling conflicts. It is designed to operate at a pre-tactical level 1000 nmi beyond the airport. The second model attempts to improve the throughput rate inside the terminal. It is designed to be used at a more a tactical level at 500 nmi away from airport. In Section II we provide some description of our fuel burn assumptions relevant to speed control and describe the issue of scheduling conflicts between arriving flights. 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 model to reduce airborne delay in the terminal due to holding and vectoring while maintain a high standard of equity. We also show the

ability of our two models to work in concert to manage the arriving flights and transfer delay within the system.

II. BACKGROUND

A. Fuel Burn 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 while maintaining a high standard of equity. This can be viewed as the process of transferring delay from the terminal area to the en route portion of a flight as well as the ground. 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 in terms of fuel burn.

Note also from Figure 2 that as altitude increases the specific range curves move markedly 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. Role of Uncertainty in Scheduling

Under the current CDM framework during TMIs arrival times are allocated to airlines using an algorithm known as Ration-by-Schedule (RBS). This algorithm is based on the consensus recognition among stakeholders that carriers have the right to capacity of the airport based on their original flight schedule. [2] Accordingly it assigns capacity to airlines based on the order that their flights appear in the schedule. While the algorithm does not guarantee an efficient allocation of slots it does ensure a fair and equitable assignment to all carriers.

The algorithm however, does not account for the uncertainty in meeting arrival times. While the algorithm provides equitable allocation in a deterministic setting the actual arrival times are generally affected by a variety of stochastic elements. When flights approach their destination they will often deviate from their scheduled arrival times. These deviations can lead to scheduling conflicts. Consider the example depicted in Figures 3a and 3b. Here 3 flights with scheduled times of arrival (STAs) within the immediate vicinity of flight 3 all approach the airport at the same time. When this happens 3 of the 4 flights will need to hold in the air until the airport can accommodate them. This additional holding leads to excess fuel burn on flights. Moreover since air traffic controllers are not involved in managing TMIs the actual order of precedence between the 4 flights could deviate from the assigned order.

As such situations are clearly undesirable it may be advantageous to incorporate their likelihood in our decision making. In the following section we present two models designed to reduce the likelihood of such scheduling conflicts. The first balances the need to issue an equitable allocation with the chance that the allocation will lead to potential conflict. The second model works to raise the level of throughput by

reducing scheduling conflicts between managed and unmanaged flights.




Figure 3a: An illustration of a scheduling conflict due to uncertainty in flight arrival times.




Figure 3b: The effect of scheduling conflicts on flight arrival times.

III. METHODOLOGY

In this section we develop the structure of our two models used to coordinate and assign arrival times to flight arrival banks. Each model assumes a multi-resource framework and iteratively resolves the problem to accommodate the changing conditions within the airspace. While both aim to transfer delay away from the terminal the first model seeks to incorporate equity in its decision making. The first model was designed to offer a relatively equitable assignment of resource capacity to carriers and provide a baseline that could be later solved to incorporate carrier preferences with assigned arrival times. The second model was designed to be used as a second layer to revise previously issued assignments once flights have traveled downstream and more information is available.

A. Equitable Allocation with Conflict Penalties

We wanted to develop a model that could be used assign flights to slots 2 hours prior to their arrival. One of the challenges in executing this scheme lies in accommodating the uncertainty in arrival times. Specifically, we did not want to issue allocations that would likely be undermined by changes in flight times that manifest closer to the airport and lead to airborne holding. Such holding is significantly more likely when flights are spaced closely together. Figure 4 shows a sparse and dense allocation around a given slot n . If a flight assigned to slot $n+2$ arrives 1 slot late in the dense allocation it

This work was sponsored in part by research grant number DTFAWA11D00017 from the Federal Aviation Administration. The opinions expressed herein are solely those of the authors.

could impact several other flights while the effect is not so prevalent in a sparse allocation. As such while we wanted to achieve a fair allocation it was critical that we looked to promote allocations that reduced the likelihood of conflict downstream. To that end we sought to define a model with the following properties:

- 1) The model should equitably assign delay when no potential scheduling conflicts are present.
- 2) The model objective should balance the desire to transfer delay away from the terminal with the need to maintain a relatively high standard of equity.
- 3) The model should be compatible with other phases of CDM. We view our model as a means to provide an initial allocation of capacity to carriers. Carriers will then be allowed to determine how this capacity will be used based on their own needs and priorities..




Figure 4: An illustration of two allocations of flights within the vicinity of a given slot n . A sparse allocation is less prone to conflicts than a dense allocation.

A bi-criteria mixed integer programming model was developed to balance our priorities of equity and delay transfer. This model considers flight assignments by dynamically resolving the problem in intervals of 15 minutes. A 15 minute look-ahead window was used to incorporate knowledge of future arrivals into our decision making. When the model was solved the flights assigned to the relevant period of interest were collected while those assigned in the look ahead window were discarded so that they could be assigned a CTA in the subsequent period. Our formulation of the problem is shown below.

Parameters:

- F – set of all flights
- S – set of all slots
- S_s – set of all slots within the neighborhood of slot s
- Ω_f – set of eligible fixes for flight f
- λ – coefficient of the convex combination in the objective function
- $a_{s'}$ – A coefficient designed to weight the proximity of slot s' to slot s
- N – Maximum number of slots that can be assigned within the neighborhood of a fix

δ – A nominally small quantity

t_{oag_f} – The scheduled arrival time flight f

t_s – The arrival time of flight f in slot s

Variables:

y_s – A continuous variable designed to assess the impact of assigning a flight within the neighborhood of a given flight assignment.

$$x_{fs}^k = \begin{cases} 1 & \text{if flight } f \text{ is assigned to slot } s \text{ through fix } k \\ 0 & \text{otherwise} \end{cases}$$

$$\min \sum_{\substack{f \in F, s \in S \\ k \in \Omega_f}} (1 - \lambda) (t_s - t_{oag_f})^{1+\epsilon} x_{fs}^k + \sum_{s \in S} \lambda y_s \quad (1)$$

$$S.T. \sum_{s \in S} x_{fs}^k = 1 \quad \forall f \in F, k \in \Omega_f \quad (2)$$

$$\sum_{f \in F} x_{fs}^k \leq 1 \quad \forall s \in S, k \in \Omega_f \quad (3)$$

$$\sum_{\substack{f \in F, k \in \Omega_f \\ s \in S_s}} x_{fs}^k \leq N \quad \forall s \in S_s \quad (4)$$

$$\sum_{\substack{f \in F, q \in Q, \\ s \in S_s}} \frac{a_{s'}}{(t_s - t_{oag_f})^{1+\epsilon} + \delta} x_{fs}^k = y_s \quad \forall s \in S \quad (5)$$

$$x_{fs}^k \in \{0, 1\}, y_s \geq 0 \quad (6)$$

Equation (2) assures that every flight is assigned to one slot. Equation (3) ensures that each slot can be assigned to no more than one flight. Equation (4) forces the number of flights assigned within a neighborhood of a given slot to not exceed a threshold value N . Equation (5) dictates the value of our continuous variable used in the objective function. It ensures that the variable value is proportional to the number of flights assigned to the slots within the neighborhood of its corresponding slot. It features three terms used to weight the variable. $a_{s'}$ is a coefficient that varies based on its proximity to the central slot. The variable is also inversely proportional to the cost of equitably assigning the flight to the slot. Since it is possible for this cost to be zero, we have added a term δ , which takes on a nominal value, to ensure that the variables stay finite. The objective function in equation (1) features a term for equity and a penalty term for flight density. Since the density term is inversely proportional to the cost coefficient in the equity term, it remains small when an equitable allocation is achieved at high cost and when an equitable allocation of flights is not within the vicinity of a given slot. When the equity term is small and there is a potential allocation that requires a number of flights be assigned within the proximity of a given slot, the density term penalizes the objective by making such an assignment substantially more expensive. Thus the model promotes balance between the equity and conflict potential.

B. Improving Terminal Throughput

While the model described in the previous section offers a first step towards allocating flights to carriers, it occurs well in advance of the terminal. There is a good deal of uncertainty in the problem when the relevant decisions are being made. To deal with this uncertainty the previous model left a number of slots unavailable for use to reduce the potential for scheduling conflict. As the flights get closer the airport the decision makers have significantly more information at their disposal and are still in a position to increase the delay transfer. In this context flights can move up into some of the previously unused spaces to improve throughput. To that end an integer programming model was developed to minimize system delay. This model considers flight assignments over a two-period rolling 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 -10 minutes. We define our variables and parameters as follows:

Parameters:

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 eligible fixes for flight f

R – set of all runways

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

t_{kfs}^{ir} – the runway arrival time of flight f assigned through slot k at fix i and slot s at runway r

e_j^{kr} – the earliest arrival time of flight f assigned through fix i and runway r

Variables:

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

$$y_{ik} = \begin{cases} 1 & \text{if a short haul flight is assigned to slot } k \text{ at fix } i \\ 0 & \text{otherwise} \end{cases}$$

$$\min \sum_{\substack{f \in F, k \in S_{if}, \\ i \in \Omega_f, s \in S_{if}, \\ r \in R}} (t_{kfs}^{ir1} - e_f^{ir1}) x_{fks}^{ir1} + \beta (t_{kfs}^{ir2} - e_f^{ir2}) x_{fks}^{ir2} \quad (7)$$

$$S.T. \sum_{\substack{k \in S_{if}, i \in \Omega_f, \\ s \in S_{if}, r \in R}} x_{fks}^{ir1} + x_{fks}^{ir2} = 1 \quad \forall f \in F \quad (8)$$

$$\sum_{\substack{f \in F, s \in S_{if}, \\ r \in R}} x_{fks}^{ir1} + x_{fks}^{ir2} \leq 1 \quad \forall k \in S_{if}, \forall i \in \Omega_f \quad (9)$$

$$\sum_{\substack{f \in F, k \in S_{if}, \\ i \in \Omega_f}} x_{fks}^{ir1} + x_{fks}^{ir2} \leq 1 \quad \forall s \in S_{if}, \forall r \in R \quad (10)$$

$$x_{fks}^{ir1} + x_{fks}^{ir2} < 1 - y_{ik} \quad \forall f \in F, \forall k \in S_{if}, \forall i \in \Omega_f, \forall s \in S_{if}, \forall r \in R \quad (11)$$

$$x_{fks}^{ir1}, x_{fks}^{ir2}, y_{ik} \in \{0,1\} \quad (12)$$

Equation (8) states that in each time period every flight is assigned to one slot. Equation (9) states that each slot at each fix can be assigned to no more than one flight. Equation (10) states that each slot at each runway can be assigned to no more than one flight. Equation (11) states that if a slot at a fix is filled by a flight originating within 500 nmi of the airport it cannot be assigned to a flight outside 500 nmi. Equation (12) notes that our variables for each period are binary. Our objective function equation (7) minimizes system delay and serves as a proxy for maximizing throughput.

IV. RESULTS AND DISCUSSION

A computational experiment was performed to evaluate the performance of our model. A scenario was composed using historical data was used to study the effect of speed control measures at a single airport. In this section we describe the scenario and associated assumptions, we present our experimental results, and we provide some analysis.

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 was obtained from an ADL file in conjunction with an ASDX file, the combination of which listed flight numbers, collection time, ETA, scheduled time of arrival (STA), the origin airport, actual time of departure, aircraft position, aircraft type, runway arrival time, STAR routes and last available fix.

The airport has 4 corner posts at the northeast, northwest, southeast and southwest corners of the airport. Arriving flights commonly fly through these corner post fixes and are sent to one of 3 runways, 2 primary runways that are used full time and another runway that is partially used.

We assumed an airport acceptance rate of 80 flights per hour. This assumption was based on full use of 2 runways and partial use of a third. The data was tested over a 4 hour period from 3:00-7:00 EST. CTAs were assigned using slot window sizes designed to accommodate the planned airport capacity at the time of arrival.

To model the problem we developed a simulation intended to mimic the basic effects of TMA. The simulation assumes

flights proceed on their trajectories with the goal of meeting their CTAs. Once a CTA is issued flights proceed to their assigned metering fix. When the flights reach their fixes the simulation accepts flights for vacant runway slots on a first-come-first-served basis.

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. Additional uncertainty was imposed to model the variability of flights in arriving at their metering fix on time. Flights were grouped into 4 pools: Airborne flights beyond 1000 nmi, airborne flights within 1000 nmi, grounded flights between 500 and 1000 nmi and grounded flights within 500 nmi. Each pool was perturbed by sampling from a different distribution to represent the variation in travel time and received a different range of permissible arrival times. Flights beyond 500 nmi were allowed to take any arrival time that could be realized solely through speed control. Flights inside of 500 nmi were allowed a ground delay of up to 30 minutes.

B. The Trade-off between Equity and Delay Transfer

The objective function of our first model features two terms, one for equity and one to promote delay transfer through conflict reduction. In order to determine the best balance between them for our model we ran our simulation at a variety of weights and calculated the resulting delay transfer and deviation from STA. The results of these runs are shown in Figure 5. The figure points to an envelope of dominant points which we have traced in red. We would like our model to produce pairings that have small deviations from the schedule that transfer large amounts of delay per flight or equivalently small schedule deviations with low numbers of flights per minute of delay transfer. This feature is evident in the figure. The knee of the curve occurs at (0.4186, 363) while the pure equity weighting resulted at (0.6608, 273). This difference corresponds to values of 2.39 vs. 1.51 minutes of delay transfer per flight when averaged over a 4 hour period and results in a total additional schedule deviation of 90 minutes over the 4 hour period. While this difference in delay transfer is small on a scale of minutes, most of the benefit of our intervention is realized when demand exceeds capacity over short intervals, which only occurred occasionally in our test case. During these periods, delay transfer performance, when using the weighting at the knee, can exceed that of the equity weighting by several minutes per flight. The number of minutes transferred would be significantly higher if our intervention were used with present TMI efforts. This point is further illustrated in the next section.




Figure 5: A Pareto Frontier illustrates the trade-off between Terminal Delay Transfer and Equity in our IP objective function.

C. Transferring Terminal Delay

A set of simulations was performed to study the ability of our models to transfer delay away from the terminal. A baseline measurement was used to determine the delay present in the terminal prior to model intervention. This run measured the holding time of each flight in terminal airspace before runway assignment. This waiting time was measured and averaged to find the mean holding delay. Our first model was then configured to assign CTAs to flights 2 hours prior to their STA in 15 min intervals. We reconfigured our simulation to use both models in a sequential fashion. In each set of runs, the delay was compared relative to the average delay without intervention. An expression for the delay transfer is provided in equation (13). The results of these runs are shown in Figure 6.

$$Avg_D_{transferred} = Avg_D^{baseline} - Avg_D^{CTA} \quad (13)$$




Figure 6: Delay Transferred Away from the Terminal using various models.

The figure suggests that our models are fairly effective in transferring delay away from the terminal. In most instances

the delay transfer with the addition of the throughput model exceeds the level achieved when only the first model was used. We transferred 58.75% of the delay using just model one while transferring 70.09% using both models. Both curves exhibit reasonable tracking although in some instances the improvement is lacking. This most commonly occurs when the capacity of the airport meets or exceeds the demand. Since our model tries to aggressively resolve scheduling conflicts, performance declines slightly when there is no conflict to be found; however, these instances are few and far between and the amount of lost delay when they happen is minimal and pales in comparison to the amount of delay transferred. The approach is also reasonably efficient in all instances; the model construction and solution time for both models never exceeded 30 seconds when used sequentially.

D. Fuel Savings

While we have demonstrated the ability of our model to transfer delay we have thus far ignored one of the primary benefits of our intervention, fuel conservation. In order to measure the fuel savings gained by our approach we need to conceptualize how that savings occurs. By transferring the delay away from the terminal we are able to move the delay to a higher altitude where the fuel efficiency of the aircraft is greater. More importantly, we are also able to eliminate a considerable amount of vectoring in the terminal and reduce the distance traveled on flight routes.

Our fuel burn savings was measured by assessing the extent to which vectoring in the terminal contributes to additional fuel burn. We assumed that aircraft vectoring inside the terminal would occur over a range of FL100 to FL250. A set of altitudes in this range were sampled from our dataset using an empirical inverse CDF. These altitudes were then used to measure the average fuel burn at a given speed based on values obtained from the BADA database. The results of these computations can be seen in Figure 7 below.




Figure 7: Average Fuel Burn rates Savings (kg/min) from the total fleet mix.

Given our fuel burn curve we can now demonstrate how the delay transfer curves shown in part c translate into fuel savings. Figure 8 shows the fuel burn savings made possible by the delay transfer relative to five different vectoring speeds

at the terminal when using both models. A comparison of the plots shows that considerable savings can be achieved regardless of vectoring speed. Although the savings is largest when vectoring at 300 knots, in every case in this example flights are able to save upwards of 87 kg of fuel per hour over the 4 hour period.




Figure 8: Average Fuel Burn rates Savings vs Time over a 4 hour period.

V. CONCLUSIONS AND FUTURE WORK

In this paper we presented two models to transfer delay away from the terminal. The first was a bi-criteria mixed integer program that balanced equity and the potential for scheduling conflict. The second model was designed to be used sequentially with the first model and sought to improve throughput. Each approach demonstrated a strong ability to transfer delay within the terminal; however, the use of both models proved most effective. An analysis of fleet fuel burn showed that the delay transfer yields significant fuel savings on a per flight basis.

The work presented in this paper opens the door to a number of areas of future exploration. In this paper we assumed a static capacity profile. In GDPs and AFP the airport/sector arrival rates often vary with time. A stochastic programming model could be built to handle this variation in capacity. As mentioned earlier the intent of our first model was to provide a baseline to assign arrival capacity to airlines in lieu of RBS during en route flight management. Additional models could be built to incorporate airline preferences for slot assignments. In this framework airlines could openly substitute between grounded and airborne slots. By allowing airlines to express their interests prior to assignment we hope to minimize the impact that any potential change in arrival time could have on future flights.

REFERENCES

- [1] T. W. Vossen and M. O. Ball, "Optimization and mediated bartering models for Ground Delay Programs," *Naval Research Logistics*, vol. 53, 2006, pp. 75 – 90.
- [2] T. W. Vossen and M. O. Ball, "Slot trading opportunities in collaborative ground delay programs," *Transportation Science*, vol. 40(1), 2006, pp. 29-43.

- [3] Joint Planning and Development Office: Concept of Operations for the Next Generation Air Transportation System, Version 2.0, June 2007. Available for public download from http://www.jpdo.gov/library/NextGen_v2.0.pdf.
- [4] F. Neuman and H. Erzberger, "Analysis of delay reducing and fuel saving sequencing and spacing algorithms for arrival traffic," NASA Technical Memorandum, 103880, 1991.
- [5] H. N. Swenson, J. Thipphavong, A. Sadosky, 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.
- [6] 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), 2012, pp. 203-224.
- [7] 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.
- [8] Airservices Australia, "Annual Report: 2007-2008," Airservices Australia, Tech. Rep., 2008.
- [9] 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.
- [10] J. Leib, "Flights flow gets innovative fix," Denver Post, Denver, CO, May 26, 2008.
- [11] D. Nieuwenhuisen, and N. de Gelder, "Optimizing nightly Schiphol traffic through time based operations," *Air Transport and Operations: Proceedings of the Third International Air Transport and Operations Symposium*, 2012: pp. 205-216.
- [12] 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 (ATM2011), Berlin, Jun. 2011.
- [13] 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 (ATM2013), Chicago, Jun. 2013.
- [14] L. Delgado and X. Prats, "En-route speed reduction concept for absorbing air traffic flow management delays," *Journal of Aircraft*, vol. 49, no. 1, Jan. 2012, pp. 214-224.
- [15] X. Prats and M. Hansen, "Green delay programs, absorbing ATFM delay by flying at minimum fuel speed," in 9th USA/Europe Air Traffic Management Research and Development Seminar (ATM2011), Berlin, Jun. 2011.
- [16] 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 (ATM2013), Chicago, Jun. 2013.
- [17] J. C. Jones, and D. J. Lovell, "Methods for curbing the exemption bias in Ground Delay Programs through speed control," presented at the 93rd Annual Meeting of the Transportation Research Board, Washington DC, USA, January, 2014.